OPSWARE INC
Automating IT™

Opsware® SAS 5 User’s Guide

Corporate Headquarters

599 North Mathilda Avenue Sunnyvale, California 94085 U.S.A.
T + 1 408.744.7300 F +1 408.744.7383 www.opsware.com

Opsware SAS Version 5.5.1

Table of Contents

Preface 41

Contents of this Guide 41

Conventions in this Guide 44

Icons in this Guide 45

Guides in the Documentation Set and Associated Users. 46

Opsware, Inc. Contact Information 46

Chapter 1: Introduction to Opsware SAS 49

Overview of Opsware SAS 49

Types of Opsware Users 51

Opsware SAS Interfaces and Tools 52

Opsware SAS Features 53

Software Provisioning 56

Operating System Provisioning 58

Patch Management for Unix 59

Code Deployment & Rollback 59

Configuration Tracking 60

Script Execution 60

Data Center Intelligence Reporting 62

Discovery and Agent Deployment 64

Server Explorer 64

Server Compliance 64

Patch Management for Windows 65

Visual Packager 65

Application Configuration Management 66

Global Shell 66

Opsware SAS Terms and Concepts 66

Agent-Server Architecture of Opsware Technology 67

Opsware SAS Model-Based Approach 68

Opsware SAS Model and Server Management 70

Software Tree, Nodes, Packages, and Templates 72

Server Management in Multiple Facilities 76

Multimaster Support 77

Chapter 2: Getting Started with the Opsware Command Center 79

Getting Started with the Opsware Command Center 79

Supported Browsers for the Opsware Command Center 79

Browser Configuration Requirements 80

Access to Features in the Opsware Command Center.................80

System Management of Opsware User Information..................80

Best Practices for Selecting Passwords.............................81

Updating User Profiles and Passwords 81

Opsware Command Center User Interface 82

Requirements for Logging In 82

Overview of the Opsware Command Center User Interface...........82

My Profile...90

Search...91

My Servers...91

Mouseover Icon Tooltips......................................91

Chapter 3: Getting Started with OCC Client 93

Overview of the OCC Client...93

Discovery and Agent Deployment...................................94

Server Explorer 94

Server Compliance...94

Patch Management for Windows 95

Visual Packager 95

Application Configuration Management............................96

Global Shell 96

Requirements for Running the OCC Client...........................97

Installing and Launching the OCC Client 97

OCC Client User Interface 99

Menus 100

Navigation Pane 101

Content Pane 101

Preview Pane. 103

Status Bar 105

OCC Client Options 105

Accessing OCC Client Options 106

General Options 106

Unmanaged Servers. 106

Packaging Servers 109

Terminal and Shell 110

Patch Policies 111

Chapter 4: Server Tracking in the OCC 113

Server Tracking. 113

Ways to Locate, List, and Display Servers. 113

Tracked Server Properties 114

Supported Operating Systems for Managed Servers 115

Server Lists. 117

Types of Server Lists 117

Server Pool 118

Manage Servers List 119

Filters on the Manage Servers List 120

My Servers 122

Overview of My Servers 122

Adding Servers to My Servers 122

Removing Servers from My Servers 123

Server Search 124

Searching for a Server By Using the Search Box. 124

Ways to Use Search. 125

Using the Search Feature 125

Rules for Server Search and for Creating Dynamic Groups 130

Line Break Workaround for Server Search 139

Conditions for Searching with Multiple Rules 140

Server Searching by IP Address 141

Example: Server Search 142

Searching for a Server Group 143

Server Identification. 144

Overview of Server Identification 144

Ways that Servers are Identified by Opsware SAS 145

Customer Accounts in Opsware SAS. 146

Associating Servers with Customers 148

Server Histories and Reports 151

Overview of Server History. 151

Viewing Server History 153

Time Stamp for Server Operations. 153

DCI Reporting 154

Hardware Information for Managed Servers 155

Viewing Managed Server Hardware Information 155

Chapter 5: Exploring Servers and Jobs in OCC Client 157

Server Explorer and Groups Browser 157

Accessing the Server Explorer 158

Server Explorer Interface 159

Actions Menu in the Server Explorer 159

Server Information 160

Server Summary 160

Server Properties 161

Server Explorer File System 163

Creating a Configuration Template from a File. 166

Creating a Package from a File 166

Opening a Remote Terminal 166

Viewing Installed Packages 167

Configured Applications 167

Configuration History. 167

Services 168

Patches. 168

Patch Policies 171

Windows Registry. 171

Windows IIS Metabase 172

Windows COM+ Objects 172

Compliance 172

Server History Window 174

Server Groups Browser 174

Accessing the Server Groups Browser 175

Server Group Members. 175

Configured Applications for Server Groups 175

Server Group Patches 176

Server Group Patch Policies 177

Server Group History Properties 178

Browsing Job Logs 179

Chapter 6: Opsware Agent Management 181

Opsware Agent on Managed Servers 181

Overview of the Opsware Agent on Managed Servers 181

Security for Opsware Agents Running on Managed Servers. 183

Opsware Agent Functionality on Managed Servers. 184

Server Information that the Opsware Agent Tracks 185

Opsware Discovery and Agent Deployment. 189

Discovering Servers for Installing an Opsware Agent 190

Setting Deployment Actions for Each Server 191

Specifying Login Settings 191

Opsware Agent Installer Options. 192

Reports on Server Status 194

Opsware Agent Installation Using ODAD 195

Prerequisite Setup for Opsware Discovery and Agent Deployment. . . 196

Permissions Required for Opsware Discovery and Agent Deployment 196

Installing Opsware Agents Using ODAD 196

Viewing Unmanaged Server Information 199

Summary Information for an Unmanaged Server. 199
History Information for an Unmanaged Server 201
Remote Terminal Sessions on Unmanaged Servers 202
Reports on Opsware Agent Installation 202
Agent Reachability Communication Tests 204
Overview of Communication Tests 205
Communication Tests and Unreachable Opsware Agents. 205
Communication Test Types 206
Communication Test Errors 207
Additional Information on Communication Tests. 212
Running a Communication Test on an Individual Server 212
Running a Communication Test on Multiple Servers. 213
Viewing Servers by Communication Status 215
Searching for Unreachable Servers 216
Creating Communication Test DCI Reports 218
Viewing My Jobs Communication Tests. 218
Exporting the Unreachable Server Status List to CSV 218

Chapter 7: Server Management in Opsware Command Center 219

Overview of Server Management 219
Communication Between Managed Servers and Opsware SAS 221

Network Address Translation (NAT) for Managed Servers. 221

Key Terms of Opsware SAS Managed Server Communication. 221

Locating the Management IP Address of a Managed Server 222

Code Deployment and Static NAT 224

Setting the Primary IP Address of a Server. 224

NAT Table Mapping and Managed Servers 225

Server Groups. 226

Overview of Server Groups 226

Types of Server Groups 229

Public Group Modeling 230

Using Server Groups Tabs 233

Ways to Create Server Groups 235

Adding a Server to Static Groups, Both Public and Private 249

Adding a Server to Dynamic Groups, Both Public and Private 249

Removing Servers from Static Groups, Both Public and Private. 250

Removing Servers from Dynamic Groups, Both Public and Private . . . 251

Moving Servers from One Static Group to Another 252

Duplicating Server Groups 252

Rules for Deleting Server Groups 254

Server Life Cycle 255

OS Provisioning and the Server Life Cycle 255

Server Properties 261

Server Management Tasks Related to the Server Life Cycle 263

Changing the Use and Stage Values for Servers 265

Editing the Properties of a Server 265

Tasks Associated with Deactivating a Server 267

Deactivating a Server. 267

Deleting a Server from Opsware SAS 268

Overview of Cloning a Server. 268

Server Management Jobs 270

My Jobs 270

My Jobs Display Information 271

My Jobs in the OCC. 272

My Jobs in OCC Client 272

Viewing Job Details in the OCC. 273

Viewing My Job Details from inside the OCC Client 275

Server Management Scheduling and Notification 276

Time-Outs for Server Management Jobs. 279

Custom Fields for Servers. 280

Overview of Custom Fields for Servers 280

Changing the Value for a Custom Field. 281

Custom Attributes for Servers 282

Overview of Custom Attributes for Servers. 283

Managing Custom Attributes 283

Adding Server Custom Attributes 284

Editing Server Custom Attributes 285

Deleting Server Custom Attributes 286

Service Levels. 286

Overview of Service Levels. 287

Adding a Service Level to the Opsware Command Center 287

Editing a Service Level 288

Ways to View the Service Level for Servers 289

Assigning a Server to a Service Level 290

Removing a Server from a Service Level. 291

Network Configuration 292

Overview of the Server Network Configuration 293

Configuring Networking for an Opsware Managed Server. 294

Chapter 8: Server Compliance 297

Overview of Server Compliance 297

Server Objects 299

Selection Criteria 300

Inclusion and Exclusion Rules 302

Snapshot Selection Criteria Overlap 306

Snapshots 307

Ways to Create a Snapshot. 310

Snapshot Templates 311

Creating a Snapshot Template 312

Setting Up Snapshot Selection Criteria 313

Creating a Snapshot 315

Saving Snapshot Selection Criteria for Reuse 316

Loading Snapshot Selection Criteria 317

Managing Your Snapshot Selection Criteria 317

Editing a Snapshot Template 318

Saving Snapshot Selection Criteria for Reuse 318

Copying a Snapshot Template for Reuse 319

Deleting a Snapshot Template. 319

Snapshot Scheduling 320

Scheduling a Snapshot 320

Viewing a Snapshot Schedule 323

Editing a Snapshot Schedule 324

Deleting a Snapshot Schedule. 325

Uses of Snapshots 325

Ways to Find a Snapshot. 326

Browsing Contents of a Snapshot. 327

Browsing File Properties, Permissions, and Content 329

Converting a *Snapshot* to a *General Snapshot*. *330*

Overview of Copying Objects from a Snapshot to a Server 331

Viewing a Snapshot Job 333

Deleting a Snapshot. 333

Auditing 334

Ways to Perform an Audit 336

Performing an Audit from a Managed Server 336

Performing an Audit form a Server Group. 336

Performing an Audit from a Snapshot 337

Performing an Audit from an Audit Template 337

Re-running an Audit from Audit Results. 337

Audit Templates 337

Creating an Audit Template 338

Specifying Source and Targets in the Audit Template 339

Setting Up Audit Selection Criteria 340

Performing an Audit. 342

Tracking an Audit's Progress 343

Saving Audit Selection Criteria for Reuse 344

Loading Audit Selection Criteria 344

Managing Your Audit Selection Criteria. 345

Editing an Audit Template. 346

Copying an Audit Template for Reuse 346

Deleting an Audit Template 347

Audit Schedules 347

Scheduling an Audit. 348

Viewing an Audit Schedule. 351

Editing an Audit Schedule. 352

Deleting an Audit Schedule 352

Audit Results. 352

Ways to View Audit Results. 353

Viewing Audit Results on a Server by Server Basis 354

Viewing Contents of Objects that Differ 356

Viewing OLE/COM Objects that Differ. 357

Viewing Metabase Objects that Differ 359

Overview of Copying Objects to a Destination Server 360

Exporting Audit Results to HTML or CSV 361

Deleting Audit Results 361

Chapter 9: Visual Packager 363

Overview of Visual Packager. 363

Packaging Server Setup 364

Installing the IDK on a Packaging Server. 365

Configuring Options for a Packaging Server. 367

Overview of Packages 368

Packaging Process. 370

Ways to Create a Package 371

Creating a Package from a Managed Server 371

Creating a Package from Audit Results 372

Creating a Package from a Snapshot. 372

Creating a Package 373

Adding New Package Content. 377

Specifying Options for New Package Content 380

Viewing Package Details 384

Chapter 10: Global Shell 385

Overview of Global Shell 385

Remote Terminal 386

Opening a Remote Terminal 386

Global Shell Session 387

Opening a Global Shell Session 387

Copying Files To and From the OGFS 388

Opsware Global File System (OGFS) 389

Directories in the OGFS 390

Diagram Conventions 390

root Directory 391

opsw Directory 391

Facility Directory. 392

Group Directory 393

Server Directory 394

Script Directory. 396

Application Directory 397

Customer Directory 398

OS Directory 398

Realm Directory 399

ServiceLevel Directory 399

Server Filtering in the OGFS 400

Remote Opsware Shell (rosh) 401

rosh Operations 403

Character Encoding for OGFS. 404

Terminal Application Configuration. 405

Data that Cannot Be Displayed 405

swenc Command (Switching Character Encoding) 406

LANG and LC_CTYPE Environment Variables 407

Transcoded Data in a Managed Server 408

Disabling the Transcoding of Managed Server Data 409

Error Messages. 410

Chapter 11: Distributed Script Execution 411

Distributed Script Execution 411

Overview of Distributed Script Execution 411

Scripts Types – My Scripts, Shared Scripts, and Ad-Hoc Scripts. 412

Script Execution Functionality 413

Ways to Initiate Script Operations 415

Script Management Tasks. 416

Viewing the Scripts List. 420

Editing, Deleting, and Downloading a Script 420

Script Version History 424

Script Execution 424

Overview of Script Execution. 425

Script Execution Guidelines 426

Executing My Script or Shared Scripts 426

Creating and Executing an Ad-Hoc Script. 431

Script Execution Results 435

Viewing Execution Results Immediately After Script Execution 437

Viewing Execution Results Stored in Opsware SAS 438

Error Resolution for Script Execution. 439

Resolving Script Uploading Errors. 439

Resolving Script Time-out Events 439

Investigating Script Non-Zero Return Codes 440

Investigating Server Authentication Errors 440

Investigating Partial Executions. 440

Opsware SAS Custom Extensions 440

Running a Custom Extension. 442

Control Scripts and Intelligent Software Modules 446

Parameters and Custom Attributes. 449

Control Scripts 449

Opening the Control Window. 449

Running a Control Script on an ISM 450

Viewing the Results of a Control Script 452

Chapter 12: Patch Management for Windows 453

Overview of Patch Management for Windows. 453

Features of Patch Management for Windows. 455

Software Library 456

Patch Management for Windows Prerequisites 456

Microsoft Patch Database 457

Opsware SAS Integration 458

Support for Patch Testing and Installation Standardization. 458

Supported Windows Patch Types 459

Supporting Technologies for Patch Management 459

Windows Hotfixes. 460

Patch Management Roles for Windows 461

Patch Management Process 462

Patch Details. 465

Patch Dependencies and Supersedence. 467

Viewing Windows Patches 467

Editing Patch Properties 468

Importing Custom Documentation for a Patch. 468

Finding Vendor-Recommended Patches 469

Finding Servers That Have a Patch Installed. 469

Finding Servers That Do Not Have a Patch Installed. 470

Importing a Patch 470

Exporting a Patch 471

Exporting Patch Information. 471

Deleting a Patch 472

Policy Management 473

Patch Policy 473

Patch Policy Exception 475

Precedence Rules for Applying Policies 476

Reconcile Patches Process 477

Reconciling Patches. 478

Setting Reconcile Options 480

Setting Reboot Options for a Patch Reconcile. 480

Specifying Pre and Post Install Scripts for a Patch Reconcile. 481

Scheduling a Patch Installation for a Patch Reconcile 483

Setting Up Email Notifications for a Patch Reconcile 483

Previewing a Patch Reconcile 484

Verifying Patch Policy Compliance 486

Creating a Patch Policy. 487

Deleting a Patch Policy 487

Adding a Patch to a Patch Policy 488

Removing a Patch from a Patch Policy 488

Attaching a Patch Policy to a Server 489

Detaching a Patch Policy from a Server 489

Setting a Patch Policy Exception. 490

Finding an Existing Patch Policy Exception. 490

Copying a Patch Policy Exception 491

Removing a Patch Policy Exception 492

Patch Compliance 492

Patch Compliance Scans 492

Patch Compliance Levels 494

Patch Compliance Rules. 494

Patch Administration for Windows 495

Setting the Default Patch Availability 495

Importing the Microsoft Patch Database 496

Selecting Windows Products to Track for Patching 496

Scheduling a Patch Policy Compliance Scan 497

Setting the Patch Policy Compliance Level. 498

Editing the Customized Patch Policy Compliance Level 498

Patch Installation 499

Installation Flags 500

Service Packs, Update Rollups, and Hotfixes 501

Installing a Patch. 501

Setting Install Options 503

Setting Reboot Options for a Patch Install 504

Specifying Pre and Post Install Scripts for a Patch Install 505

Scheduling a Patch Installation 506

Setting Up Email Notifications for a Patch Install. 507

Previewing a Patch Installation 507

Viewing Job Progress of a Patch Install. 508

Patch Uninstallation 509

Uninstallation Flags 510

Uninstalling a Patch 511

Setting Uninstall Options 512

Setting Reboot Options for a Patch Uninstall. 512

Specifying Pre and Post Install Scripts for a Patch Uninstall. 514

Scheduling a Patch Uninstallation. 515

Setting Up Email Notifications for a Patch Uninstall 515

Previewing a Patch Uninstallation 515

Viewing Job Progress of a Patch Uninstall 516

Chapter 13: Patch Management for Unix 519

Overview of Patch Management for Unix. 519

Features of Patch Management for Unix 520

Opsware SAS Integration 520

Support for Unix Patch Testing and Installation Standardization 521

Patch Management Roles for Unix 521

Patch Administrator 522

System Administrator 522

Patch Management for Specific Unix Operating Systems 523

Supported Unix Versions and Patch Types. 523

Underlying Technologies for Patch Management on Unix 524

AIX Patches 525

Solaris Patches 526

HP-UX Patches 527

Patch Uploads for Unix 527

Patch Uploads for Specific Unix Versions 527
Methods for Uploading a Unix Patch 528
Preparation for Uploading a Patch for Unix. 528
Uploading a Patch for Unix. 529
Patch Testing for Unix 531
Patch Administration for Unix 531
Setting the Patch Status for Unix 531
Editing Patch Options for Unix. 532
Setting Patch Installation Order Dependencies for Unix 533
Patch Installation and Uninstallation for Unix 535
Application Patches for Unix 536
Installation Scripts for Patches on Unix 536
Installation and Uninstallation Flags for Patches on Unix. 537
Installing a Unix Patch 538
Installing an Application Patch on Unix 542
Uninstalling a Unix Patch. 545
Uninstalling an Application Patch on Unix. 546

Chapter 14: Software Provisioning 549

Installing and Uninstalling Software 549
Overview of Installing and Uninstalling Software 549
Software Package Management in the Model Repository 550
Ways to Install Software 550
Opsware SAS Categories for Software 551
Software Provisioning Functionality 551
Platform-Specific Reconcile. 552
Software Installation and Uninstallation Options .552

Script Error Conditions 552

Inheritance and the Install and Uninstall Software Wizards. 553

Installing Software with the Install Software Wizard 553

Package Installation on Servers with Low Disk Space 558

Uninstalling Software with the Uninstall Software Wizard 560

Template Installation 563

Overview of Installing Templates. 564

Installing Templates with the Install Templates Wizard. 564

Chapter 15: Operating System Provisioning 567

Supported Operating Systems for OS Provisioning 567

OS Provisioning 569

Overview of OS Provisioning 569

Server Life Cycle for OS Provisioning 570

OS Provisioning 571

Overview of OS Provisioning 572

Solaris OS Provisioning. 573

Linux OS Provisioning 574

Windows OS Provisioning. 574

Hardware Preparation 575

New Server Booting 576

Booting New Servers with Different Operating Systems 576
OS Build Agent 577
Booting a Windows or Linux Server with PXE 577
Booting a Windows or Linux Server 579
Booting a Solaris Server Over the Network 580
Installation of OS Build Agents 581
Verifying Installation of an OS Build Agent 581
Recovering when an OS Build Agent Fails to Install 582
OS Installation with the Opsware Command Center 582
Servers and the Server Pool. 583
Ways to Install Operating Systems on Servers. 584
Installing an OS by Using a Template. 585
Installing an OS by Using a Custom Installation. 590
Recovering when an OS Installation Fails 592
Network Configuration for Servers after OS Provisioning 595
Requirements for Reprovisioning Solaris and Linux Servers. 595
Reprovisioning a Solaris or Linux Server 596

Chapter 16: Application Configuration Management 599

Overview of Application Configuration Management (ACM) 599
Application Configuration Creation and Use 600
ACM Components 602
Configuration Template. 603
Application Configuration 603
Value Set Editor 603
Configuration Markup Language (CML). 607
Application Configuration Inheritance 607

Application Configuration Default Values 607

Application Instance Values 609

Sequence Merging and Inheritance 612

Sequence Replace. 613

Sequence Append 613

Sequence Prepend. 615

Using ACM. 616

Creating an Application Configuration 617
Creating a Configuration Template. 618
Viewing Application Configuration Template Sources 620
Adding or Removing Configuration Templates. 621
Deleting Application Configurations 621
Loading a Template File 622
Setting a Configuration Template to Run as a Script 624
Specifying Template Order. 625
Editing Default Values for an Application Configuration. 626
Attaching an Application Configuration to a Server or Group. 629
Setting Application Configuration Values on a Server or Group. 630
Loading Existing Values into a Configuration Template 633
Pushing Changes to a Server or Group 634
Scheduling an Application Configuration Push. 635
Comparing Two Configuration Templates 637
Comparing a Template Against an Actual Configuration File 638
Auditing an Application Configuration 638
Scheduling an Application Configuration Audit 640
Rolling Back to a Previous State 641

Chapter 17: Configuration Tracking 643

Overview of Configuration Tracking 643
File Types Supported. 644
Configuration Tracking Operations. 644

Change Detection 645

Types of Actions Performed. 645

Types of Backups Performed. 645

Email Automated Configuration Tracking and Logging Actions 646

Creating the Email Notification List 647

Customized Configuration Tracking Policies. 648

Node-Based Entries and Server-Based Entries 649

Policies for Customizing Multiple Servers 650

Adding or Editing Customized Tracking Policy Entries. 651

Disabling Customized Tracking Policy Entries 654

Enabling Customized Tracking Policy Entries. 654

Viewing a Server’s Tracking Policy 655

Reconciling Customized Tracking Policies 655

Enabling and Disabling Configuration Tracking on Servers. 656

Manual Backups 657

Performing Manual Backups 658

Backup History 658

Viewing the Backup History 659

Viewing the List of Backup Events 660

Types of Backup Events 660

Backup History Search Options 661

Backup Info and Backup Manifest 661

File Info and File Versions. 664

Deleting Backups 665

Restoration of Backed Up Files. 666

Overview of Procedures Used to Restore Backed Up Files. 666

Restore Queue 667

Storing files in the Restore Queue 667

Incremental Backups for Restoration and Rollbacks 667

Entries for Directories in the Backup History 668

File Not Found Entries. 668

Restoring Backups. 669

Rolling Back Restored Files 669

Chapter 18: Code Deployment and Rollback 671

Opsware Code Deployment Process 671

Deploying Code 671

Uploading Code and Content to Staging. 673

CDR Operations and Directories. 674

CDR Features 675

CDR Permissions 676

Accessing CDR 677

Services, Synchronizations, and Sequences 678

Overview of Performing Services, Synchronizations, and Sequences . 679
Synchronization of Site Code and Content. 679
Performing Synchronizations 681
CDR Cutover Operation 683
CDR Service Operations 685
Starting and Stopping Host Services 685
Backing Up Code and Content 685
Restoring Code and Content from a Previous Version. 687
Rolling Back Code and Content to the Previous Version 687
Accessing Service Operations in CDR 688
Performing Service Operations by Service Name 689
Performing Service Operations by Host Name. 691
Performing Sequences 693
Processing Code Deployment Requests from Users 695
Status View of Previous Operations 696

Appendix A: Reconcile 699

Overview of Reconcile 699
Ways to Perform Reconcile 700
Reconcile Operations and Guidelines. 701

Reconcile Operations 701
Reconcile and Package Metadata. 703
Installation and Uninstallation Order. 703
Software Installation Order for Adopted Software. 704
Patches and Reconcile 704
Preview Reconcile 705
Types of Reconcile. 705
Reconcile on Supported Operating Systems. 706
AIX Reconcile 707
HP-UX Reconcile 708
Solaris Reconcile 708
Linux Reconcile. 708
Specifying Scripts During Reconcile. 709
Reconcile Output 709
Assigning to and Removing Servers from Nodes 710
Assigning Servers to Nodes 710
Removing Servers from Nodes 710
Reconcile Software Wizard 711
Directly Reconciling Servers 711

Appendix B: Opsware Agent CLI Utilities 715

Agent Installation Using the CLI 715

Overview of Agent Installation Using the CLI 716

Preparation for Opsware Agent Installation 718

Preassimilation Checklist 718

Installing an Opsware Agent by using the Agent Installer CLI. 721

Opsware Agent Installer Options. 722

Starting an Opsware Agent on a Server 727

Verifying Opsware Agent Functionality 727

Augmenting the Information for a Managed Server 728

Uninstalling an Opsware Agent on Unix and Windows 729

Uninstalling Earlier Versions of Opsware Agents on Unix 731

Uninstalling Earlier Versions of Opsware Agents on Windows 732

Opsware Agent Upgrade Tool 732

Ways to Upgrade Opsware Agents. 733

Prerequisites for Using the Opsware Agent Upgrade Tool. 734

Upgrading the Opsware Agent on Managed Servers 734

Commands for the Opsware Agent Upgrade Tool. 735

Options for the Opsware Agent Upgrade Tool 738

Example: Options for the Opsware Agent Upgrade Tool 740

Example: Commands and Output for Agent Upgrade Tool 740

Appendix C: Communication Test Troubleshooting 743

Overview of Agent Communication Tests 743

Command Engine to Agent (AGT). 744

AGT – OK. 745
AGT – Untested 745
AGT – Unexpected error. 745
AGT – Connection refused. 745
AGT – Connection time-out 746
AGT – Request time-out 746
AGT – Server never registered 747
AGT – Realm is unreachable 747
AGT — Tunnel setup error. 748
AGT — Gateway denied access. 748
AGT — Internal Gateway error 748
AGT — Gateway could not connect to server. 748
AGT — Gateway time-out 749
Crypto Match (CRP) 749
CRP – OK. 749
CRP – Untested 749
CRP – Unexpected error. 749
CRP – Agent certificate mismatch 750
CRP – SSL negotiation failure 750
Agent to Command Engine (CE) 751

CE – OK 751

CE – Untested. 751

CE – Unexpected error 752

CE – Connection refused 752

CE – Connection time-out 752

CE – DNS does not resolve 753

CE – Old Agent version. 753

CE – Realm is unreachable 753

CE – No Gateway defined 754

CE – Tunnel setup error 754

CE – Gateway denied access 754

CE – Gateway name resolution error 754

CE – Internal Gateway error. 755

CE – Gateway could not connect to server 755

CE – Gateway time-out 755

CE – No callback from Agent 756

Agent to Data Access Engine (DAE). 756

DAE – OK. 757
DAE – Untested 757
DAE – Unexpected error. 757
DAE – Connection refused. 757
DAE – Connection time-out 758
DAE – DNS does not resolve 758
DAE – Old Agent version 758
DAE – Realm is unreachable 759
DAE – No Gateway defined 759
DAE – Tunnel setup error 759
DAE – Gateway denied access 760
DAE – Gateway name resolution error 760
DAE – Internal Gateway error 760
DAE – Gateway could not connect to server 760
DAE – Gateway time-out. 761

Agent to Software Repository (SWR) 761

SWR – OK 762
SWR – Untested 762
SWR – Unexpected error 762
SWR – Connection refused 762
SWR – Connection time-out 763
SWR – DNS does not resolve 763
SWR – Old Agent version 763
SWR - Server identification error 764
SWR – Realm is unreachable. 764
SWR – No Gateway defined. 765
SWR – Tunnel setup error 765
SWR – Gateway denied access. 765
SWR – Gateway name resolution error 765
SWR – Internal Gateway error. 766
SWR – Gateway Could not connect to server 766
SWR – Gateway time-out 766

Machine ID Match (MID). 767

MID – OK 767
MID – Untested 767
MID – Unexpected error 767
MID – MID mismatch 767

Common Troubleshooting Tasks 768

Verifying that an Agent is Running 768
Verifying that a Port is Open on a Managed Server 769
Restarting an Opsware Agent 770
Checking Management IP of a Managed Server 770
Checking Network Gateway Configuration 770
Resolving Host Name 771

Appendix D: Content Pack 773

Overview of the Content Pack. 773
Application Configurations 773
Application Configuration Description Key 774
System Application Configurations. 775
Service Application Configurations. 784
"Application" application configurations 788
Selection Criteria for Snapshots. 789
Linux-FilesystemRuleset 790
Solaris-FilesystemRuleset. 790
Unix/Linux-Apache2Ruleset 791
Unix/Linux-SunOne6.1RuleSet. 791
Windows 2000 SecurityRuleSet(WindowsSecureDC). 791
Windows 2003 SecurityRuleSet(WindowsSecureDC). 792
Global Shell Scripts 794

FindByProcessName.sh 794

FindProcessMoreThanXCpu.sh 795

FindSpaceAvailable.sh 795

FindSpaceLessThanX.sh 795

FindTopNCpu.sh 796

Appendix E: Glossary 797

Index 813

Preface

Welcome to the Opsware Server Automation System (SAS) – an enterprise-class software solution that enables customers to get all the benefits of the Opsware data center automation platform and support services. Opsware SAS provides a core foundation for automating formerly manual tasks associated with the deployment, support, and growth of server and server application infrastructure.

This guide describes how to use Opsware SAS, starting with an introduction to the system and how to navigate the user interface. It provides information about managing servers, operating system provisioning, managing software packages, provisioning applications, managing patches, reconciling servers, script execution, configuration tracking, and deploying and rolling back code. This guide is intended for system administrators who are responsible for all aspects of managing and provisioning the servers in an operational environment.

Contents of this Guide

This guide contains the following chapters and appendices:

Chapter 1: Introduction to Opsware SAS: Provides a high-level overview of Opsware SAS, including the system features, Web Service APIs, and multimaster.

Chapter 2: Getting Started with the Opsware Command Center: Includes information about supported operating systems and browsers, navigation of the user interface, and an explanation of each of the features found on the Opsware Command Center Home page.

Chapter 3: Getting Started with OCC Client: Includes information about how to get started using the OCC client, the user interface, the client installation and launch, and OCC Client main features.

Chapter 4: Server Tracking in the OCC: Provides information about server asset tracking, server lists, server search, server histories and reports.

Chapter 5: Exploring Servers and Jobs in OCC Client: Provides information browsing servers, server groups, and jobs, also includes information about copying files on a managed server's file system.

Chapter 6: Opsware Agent Management: Includes information about Opsware Agents on managed servers, the Opsware Discovery and Agent Deployment feature, and Opsware Agent reachability communication test.

Chapter 7: Server Management in Opsware Command Center: Provides information about all aspects of server management including server groups, server life cycle, server locking, and service levels.

Chapter 8: Server Compliance: Includes information about the server compliance and auditing features, and explains how to create snapshots, differencing servers, and auditing servers.

Chapter 9: Visual Packager: Provides information about creating software packages from managed server information, and includes such topics as how to configure the OCC Client to access a packaging server, how to create a package, and explains different methods of software packaging.

Chapter 10: Global Shell: Provides information about the Global Shell, the Opsware Global File System, setting user and group permissions with the Global Shell, accessing the Global Shell and remote terminals on servers, and explains Global Shell commands.

Chapter 11: Distributed Script Execution: Describes script types, permissions required for scripting tasks, script version history, and how to create, edit, and delete scripts. It also discusses script execution for all script types, script execution results, and scripting error resolution.

Chapter 12: Patch Management for Windows: Provides information about managing patches for the Windows 2000, Windows 2003, and Windows NT4.0 operating systems. These patches include Service Packs, Update Rollups, and hotfixes. It describes the user roles: a policy setter, a patch administrator, and a system administrator. It also describes reconciling, previewing (an install), installing, and uninstalling patches.

Chapter 13: Patch Management for Unix: Provides information about managing patches for Unix operating systems. It discusses patch types, testing, and installing and uninstalling patches. It review the roles of the patch administrator and system administrator in applying patches, and the permissions required for performing patch management.

Chapter 14: Software Provisioning: Provides information about installing and uninstalling software and using templates for installation.

Chapter 15: Operating System Provisioning: Provides information about supported environments for OS provisioning and an overview of the permissions and server life cycles associated with OS provisioning. It also describes the process for provisioning, an overview of the hardware preparation, information about booting new servers, and using the Opsware Command Center to install operating systems.

Chapter 16: Application Configuration Management: Provides information about managing application configurations through the OCC Client, and includes such topics as creating Application Configurations, Application Configuration inheritance, editing value sets, and applying Application Configurations to a server.

Chapter 17: Configuration Tracking: Provides information about configuration tracking policies, the supported types of configuration files and databases, how changes are detected, node-based tracking policies, and the tracking policies for a specific server. It also discusses reconciling customized tracking policies, performing manual backups, viewing backup history, restoring backed up files, and enabling and disabling configuration tracking.

Chapter 18: Code Deployment and Rollback: Provides information about uploading code and content to staging, and performing services, synchronizations, and sequences to deploy code and content to managed servers.

Appendix A: Reconcile: Reviews the server reconciliation process and how it works, with an in-depth discussion of how to perform a reconcile, reconcile and package metadata, installation and uninstallation order, *adopted* software, patches and reconcile, and reconcile preview. Types of reconcile are also discussed, along with reconcile on supported operating systems, reconcile and scripts, reconcile output, assigning servers to and removing servers from nodes, and the Reconcile Software Wizard.

Appendix B: Opsware Agent CLI Utilities: Provides information about installing Opsware Agent using the Opsware Command Line Interface and the Agent Upgrade Tool.

Appendix C: Communication Test Troubleshooting: Provides troubleshooting information to diagnose Opsware Agent unreachability problems.

Appendix D: Content Pack: Describes content for Opsware SAS application configurations, selection criteria for snapshots, and Global Shell scripts.

Appendix E: Glossary: Defines terminology and acronyms that are unique to Opsware SAS.

Conventions in this Guide

This guide uses the following typographical and formatting conventions.

NOTATION	DESCRIPTION
Bold	Identifies field menu names, menu items, button names, and inline terms that begin with a bullet.
`Courier`	Identifies text that is entered or displayed at the command-line prompt, such as Unix commands, Opsware SAS commands, file names, paths, directories, environment variable names, contents of text files that are viewed or edited with a text editor, source code in a programming language, and SQL (database) commands.
Italics	Identifies document titles, DVD titles, web site addresses. Used to introduce new terms when they are first defined in a document and for emphasis.

Icons in this Guide

This guide uses the following iconographic conventions.

ICON	DESCRIPTION
📄	This icon represents a note. It identifies especially important concepts that warrant added emphasis.
☑	This icon represents a requirement. It identifies a task that must be performed before an action under discussion can be performed.
💡	This icon represents a tip. It identifies information that can help simplify or clarify tasks.
⚠	This icon represents a warning. It is used to identify significant information that must be read before proceeding.

Guides in the Documentation Set and Associated Users

- The *Opsware® SAS User's Guide* is intended to be read by the system administrator who is responsible for performing the day-to-day functions of managing servers, provisioning operating systems, uploading packages, setting up the Software Tree and node hierarchies, attaching software applications and installing them on servers, managing patches, reconciling servers with software, creating and executing scripts, tracking configuration, and deploying and rolling back code and content. It also documents the day-to-day functions of managing servers, such as server compliance and auditing, software packaging, application configuration, agent deployment, and global shell remote data center management.
- The *Opsware® SAS Administration Guide* is intended to be read by Opsware administrators who will be responsible for monitoring and diagnosing the health of the Opsware SAS components.
- The *Opsware® SAS Deployment and Installation Guide* is intended to be used by system administrators who are responsible for the installation of Opsware SAS in a facility. It documents how to run the Opsware Installer and how to configure each of the components.
- The *Planning Deployments for Opsware® SAS* is intended to be used by advanced system administrators who will be responsible for planning all facets of an Opsware SAS installation and deployment. It documents all the main features of Opsware SAS and scopes out the planning tasks necessary to successfully deploy Opsware SAS. Sections include: planning the Opsware SAS design for a core, types of installations, and discusses business goals that can be achieved using the software. It also includes information on system sizing, checklists, and best practices.
- The *Opsware® SAS Configuration Guide* is intended to be used by system administrators who are responsible for all facets of configuring the Opsware Command Center. It documents how to set up users and groups, how to configure Opsware server management, and how to set up the main Opsware Command Center features, such as patch management, configuration tracking, software repository replicator setup, code deployment, and software provisioning.

Opsware, Inc. Contact Information

The main web site and phone number for Opsware, Inc. are as follows:

- *http://www.opsware.com/index.htm*
- +1 (408) 744-7300

For links to the latest product documentation and software downloads, see the Opsware Customer Support site:

- *https://download.opsware.com/opsw/main.htm*

For troubleshooting information, you can search the Opsware Knowledge Base at:

- *https://download.opsware.com/kb/kbindex.jspa*

The Opsware Customer Support email address and phone number follow:

- support@opsware.com
- +1 (877) 677-9273

Chapter 1: Introduction to Opsware SAS

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Opsware SAS
- Types of Opsware Users
- Opsware SAS Features
- Opsware SAS Terms and Concepts

Overview of Opsware SAS

Opsware SAS provides a core set of features that automate critical areas of server and application operations – including the provisioning, deployment, patching, and change management of servers – across major operating systems and a wide range of software infrastructure and application products.

Opsware SAS does not just automate your operations, it also allows you to make changes more safely and consistently, because you can model and validate changes before you actually commit the changes to a server. Opsware SAS helps ensure that modifications to your servers work on your first attempt, thereby reducing the risk of downtime.

Using Opsware SAS, you can coordinate many operations tasks, across many IT groups with everyone working with the same understanding of the state of servers, applications, and configurations. This coordination ensures that all IT administrators have full knowledge of the current state of the environment before further changes are made.

Opsware SAS allows you to incorporate and maintain operational knowledge gained through long hours of trial-and-error processes. After an administrator has found and tested a procedure or configuration, that knowledge can be translated into a model that is stored in a central repository. This allows you to continue to benefit from the operational knowledge gained by your system administrators, even if they are no longer working in your organization.

The following figure provides an overview of how Opsware SAS automates server and application operations across all major platforms and a wide range of applications. Each feature that is shown in the diagram is discussed in the following sections.

Figure 1-1: Overview of Opsware SAS Features

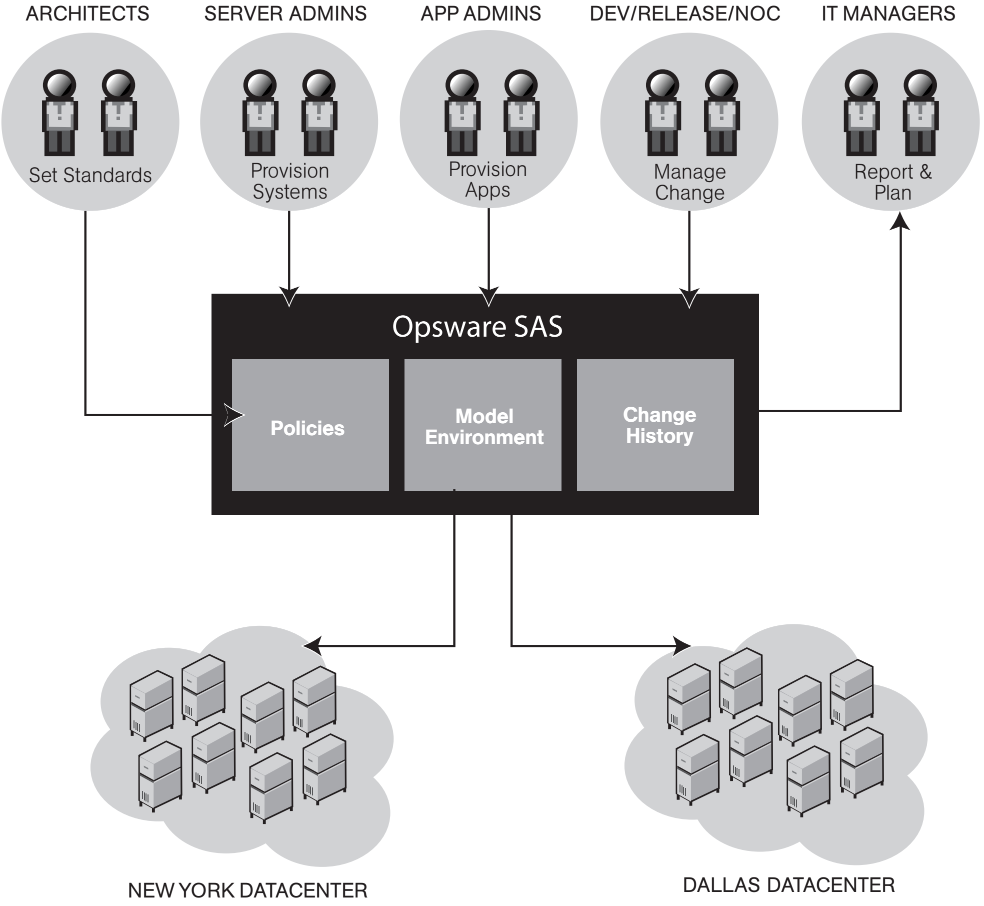

Types of Opsware Users

The following table identifies the types of Opsware users and their responsibilities.

Table 1-1: Types of Opsware Users

OPSWARE USER	RESPONSIBILITIES
Data Center and Operations Personnel	After manually racking and stacking servers, manage customer facilities and boot bare-metal servers over the network or from an Opsware boot image.
System Administrators	Install operating systems and applications (for example, Solaris 5.7 or WebLogic 6.0 Web Server), upgrade servers, create operating system definitions, and set up software provisioning.
Site Engineers and Customer Project Managers	Deploy custom code on servers.

In addition to the Opsware users listed above, this guide describes the following three types of users:

- **End Users** are responsible for all aspects of managing and provisioning the servers in an operational environment. In the Opsware SAS documentation, these users are referred to as Opsware users or system administrators. These users log into the Opsware Command Center and OCC Client and use these interfaces to manage servers in their IT environment.
- **Opsware Administrators** are the users, with special training and information, who are responsible for installing and maintaining Opsware SAS. In the Opsware SAS documentation, these users are referred to as Opsware administrators. They use the Administration features in the Opsware Command Center to manage Opsware SAS and Opsware users (by adding user accounts and assigning permissions for different levels of operation and access), to add customers and facilities, and to change Opsware SAS configurations. They monitor and diagnose the health of Opsware SAS components. Opsware administrators need to understand how Opsware SAS features operate to support users and Opsware SAS.
- **Policy Setters** are the power users who are responsible for architecting what Opsware SAS will do in the managed environment; for example, they determine which operating systems can be installed on your managed servers and how those operating systems

will be configured during installation. Policy setters, for example, prepare specific features in Opsware SAS by defining the Software Tree, preparing Operating System Definitions, and acting as Patch Administrators to approve patches for installation in the operational environment.

Opsware SAS Interfaces and Tools

Depending on the type of operation you need to perform with Opsware SAS, you select the appropriate user interface, as Figure 1-2 shows.

Figure 1-2: Interfaces in Opsware SAS

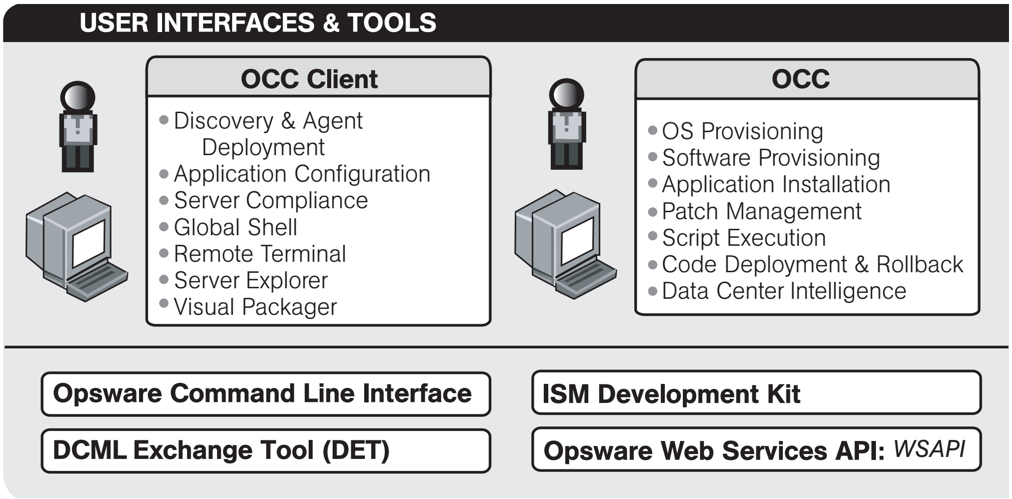

- **Opsware Command Center**: The web-based user interface to Opsware SAS through which users can manage servers, provision applications and operation systems onto servers, run distributed scripts on servers, and deploy code and content to servers, among other things.
- **OCC Client:** A Java Web-Start application that extends the Opsware Command Center (OCC) features and provides the following new features:
    - Discovery and Agent Deployment
    - Server Explorer
    - Server Compliance
    - Visual Packager

  - Application Configuration Management
  - Global Shell

- **Opsware Command Line Interface (OCLI)**: A command line interface that users can use to upload packages to the Opsware Software Repository (version 1), and perform many other Opsware SAS operations (version 2).
- **DCML Exchange Tool (DET)**: A utility that enables users to export almost all server management content from any Opsware core – standalone or multimaster mesh – and import it into any other Opsware core. Opsware SAS can also provide pre-packaged server management content appropriate for new installations that can be imported into a core after initial setup. See the *Opsware® SAS DCML Exchange Tool (DET) Reference Guide* for information about using this utility.
- **ISM Development Kit**: A development kit that consists of command-line tools and libraries for creating, building, and uploading ISMs. An ISM is a set of files and directories that include application bits, installation scripts, and control scripts. See the *Opsware® SAS ISM Development Kit Guide* for information about using the ISM Development Kit.
- **Opsware Web Services APIs**: A standards-based programming environment built on open industry standards such as SOAP (Simple Object Access Protocol) and WSDL (Web Services Definition Language). The Web Services API enables the integration of applications and other systems with Opsware SAS. This broadens the scope of how IT can use Opsware SAS to achieve operational goals

Opsware SAS Features

Opsware SAS is made up of a set of Opsware SAS features. Opsware SAS features are the components that automate particular IT processes.

The features are designed to replace ad hoc, error-prone, manual processes. For example, by using the OS Provisioning feature, users can set standards for different types of servers and automatically provision the servers, saving time and ensuring that operating system builds are consistent. By using the Patch Management feature, users can establish polices about how patches are installed. Opsware SAS uniformly enforces those polices.

The following features are currently available as part of Opsware SAS:

- Software Provisioning
- Operating System Provisioning

- Patch Management for Unix
- Patch Management for Windows
- Code Deployment & Rollback
- Configuration Tracking
- Script Execution
- Data Center Intelligence Reporting
- Discovery and Agent Deployment
- Server Explorer
- Server Compliance
- Visual Packager
- Application Configuration Management
- Global Shell

All Opsware SAS features support cross-platform environments and are designed to automate both new and existing data center environments. See the following figure.

Figure 1-3: Opsware SAS Features

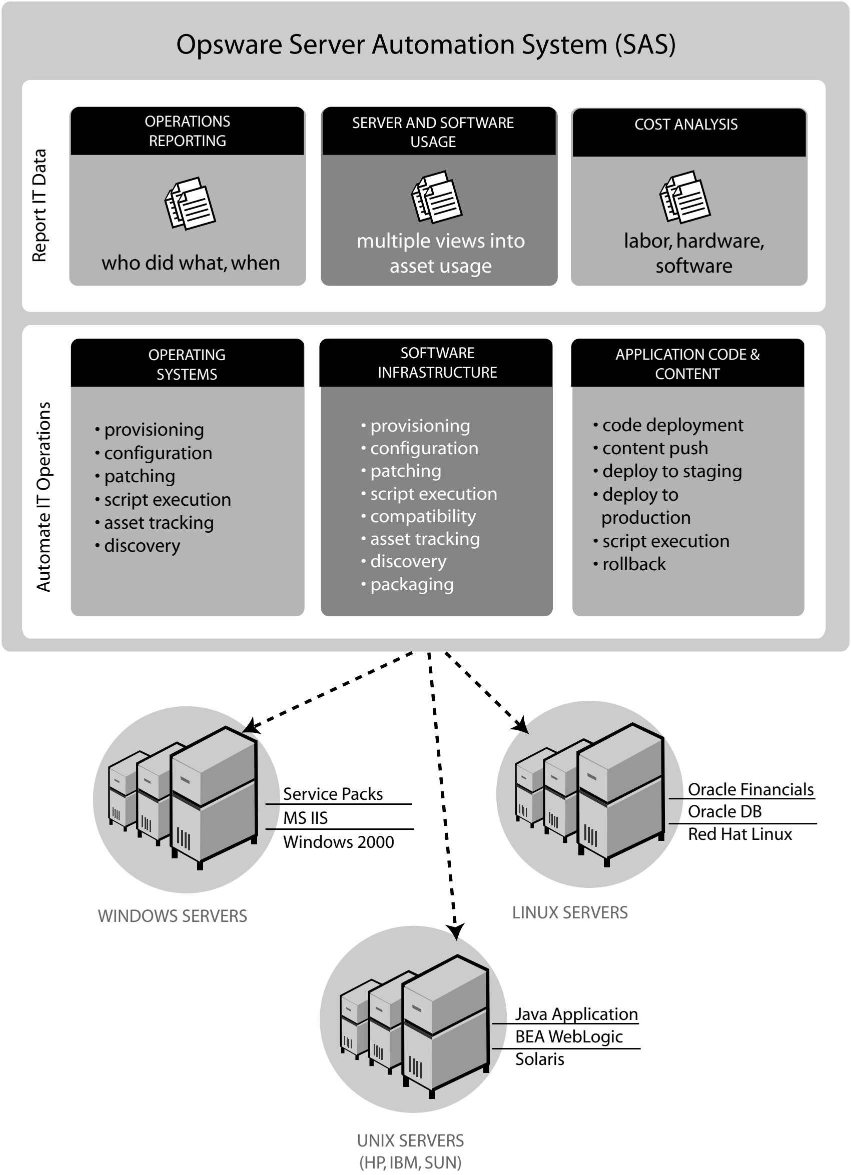

Software Provisioning

The Software Provisioning feature provides a systematic way to install, configure, and remove packaged software across Windows, Unix, and Linux servers distributed across different data centers. The Opsware SAS unique model-based approach enables many different teams, such as the system administration team, the database team, and the application development team, to manage the same set of servers. Each of these teams has a common view of the environment.

The Software Provisioning feature leverages the Opsware SAS model-based approach, which provides the following unique capabilities and benefits:

- **Detailed information about the latest system state and configurations**

  The Software Provisioning feature automatically creates and updates two lists: the list of software that users indicate should be installed on a server and the list of software that is actually installed on a server. By maintaining this detailed model of the server's current state, Opsware SAS helps keep different IT groups managing the same server in sync and ensures that all groups making server changes are working with the same knowledge of the current state of the environment.

  Using this model, Opsware SAS enables multiple groups to manage the same server without stepping over each other's changes. An accurate model of the software installed on a server, granular role-based access control, a unified audit trial, and the ability to rollback changes, all contribute to the Opsware SAS ability to coordinate the activities of many different administrators managing the same server.

- **Integration with other automation functions**

  The Software Provisioning feature is fully integrated with other Opsware SAS features, enabling software provisioning to be performed automatically with other tasks, such as operating system provisioning. Because software provisioning shares the same environment model as the other functions, the state of the environment is always known. This means that different groups, such as OS administrators, application administrators, security administrators, and others, can work together and communicate more effectively.

- **Simulation of software installation and removal**

  The Opsware SAS provisioning engine simulates installation and uninstallation actions before it applies changes to production servers. Users can view the list of software packages to be added or removed before they authorize Opsware SAS to execute the change. This ensures that all changes are pre-tested and validated before propagating changes to the production environment.

- **An up-to-date model of the actual server environment**

  Opsware SAS regularly refreshes its view of what is installed on a server, including both hardware and software. This real-time understanding of server state and configurations ensures that administrators provision the right software to the right servers at the right time. It also ensures that dependencies and prerequisites are checked and installed as needed.

- **Sophisticated role-based access control**

  Opsware SAS enforces a security policy that allows only authorized users to install or remove particular types of software on a particular server. For example, companies can define an access control rule that permits only DBAs to add or remove database software from a server.

- **A unified audit trail**

  Opsware SAS maintains a comprehensive audit trail of the software that Opsware users install, configure, and remove from a server. When combined with the additional events that Opsware SAS tracks – including configuration updates, business application pushes and rollbacks, hardware upgrades, and executed scripts – organizations gain a complete view of server activity over time.

- **The ability to rollback to a last known good state**

  The Software Provisioning feature allows users to back out of software provisioning operations. In the event an upgrade or installation goes awry, administrators can back out the change to return to the last known good state.

- **Ability to store powerful name-value pairs**

  Opsware SAS helps organizations increase software package re-use by enabling administrators to install the same software package on different servers. Server-specific configuration values are fetched from Opsware SAS (or calculated based on those values).

Operating System Provisioning

The OS Provisioning feature gives administrators the ability to provision operating system baselines onto bare metal servers quickly, consistently, and with minimal manual intervention. Bare metal OS provisioning is a key part of the overall process of getting a server into production.

Benefits of the OS Provisioning feature include the following items:

- **Integration with the other features of Opsware SAS**

  Because the OS Provisioning feature is integrated with the suite of Opsware SAS automation capabilities, including patch management, software provisioning, and distributed script execution, handoffs between IT groups are seamless. Opsware SAS ensures that all IT groups are working with a shared understanding of the current state of the environment, which is an essential element of delivering high-quality operations and reliable change management.

- **The ability to easily update server baselines without re-imaging servers**

  Unlike many other OS provisioning solutions, systems provisioned with Opsware SAS can be easily changed after provisioning to adapt to new requirements. The key to this benefit is the Opsware SAS use of templates and its installation-based approach to provisioning.

- **Flexible architecture designed to work in many environments**

  Opsware engineers carefully designed the OS Provisioning feature to handle many different types of servers, networks, security architectures, and operational processes. Opsware SAS works well in floppy (Windows provisioning), CD (Linux provisioning), or network-boot environments, with scheduled or on-demand workflows, and across a large variety of hardware models. This flexibility ensures that you can provision operating systems to suit your organization's needs.

Opsware SAS automates the entire process of provisioning a comprehensive server baseline, which typically consists of the following tasks:

- Preparing the hardware for OS installation
- Installing a base operating system and default OS configuration
- Applying the latest set of OS patches, the exact list depends on the applications running on the server

- Installing system agents and utilities such as SSH, PC Anywhere, backup agents, monitoring agents, or anti-virus software
- Installing widely-shared system software such as Java Virtual Machines
- Executing pre-installation or post-installation scripts that configure the system with values such as a root password

Patch Management for Unix

The Patch Management for Unix feature provides two features critical to patch management: the ability to react quickly to newly-discovered threats and the degree of control required to ensure that a new patch has been properly tested and installed in a uniform way.

Opsware SAS has a deep understanding of native patch formats and structure. System administrators upload patches directly into Opsware SAS, which understands and respects the structure of those patches in their native forms. It treats Solaris patch clusters, for example, differently from AIX APARs. Native patch support greatly increases both the flexibility and reliability of patch installation.

The Patch Management for Unix feature provides the following functionality:

- Scalable, cross platform patch deployment
- Reduced risk throughout automated patch rollback
- A central, shared patch repository to improve access
- Secure access control
- The ability to install patches on one server, or simultaneously on many servers
- The ability to schedule automated future installation (for example, to take advantage of maintenance windows)
- The inclusion of patches in the template for an operating system, so all newly provisioned servers receive the most up-to-date set of recommended patches

Code Deployment & Rollback

Opsware SAS automates code and content deployment to reduce the risk and time requirements associated with pushing new code to production. The Code Deployment & Rollback (CDR) feature provides an automated system for deploying code (such as, ASP, JSP, JAR, Java, C++, and Perl files) and content (such as, HTML, JPEG, GIF, and PDF files).

Specifically, CDR enables you to perform the following actions:

- Push code from staging or development environments to production environments.
- Synchronize code and content across multiple servers and locations.
- Automatically rollback to the previous version of code or content.
- Sequence multiple, complex deployment steps into repeatable workflows.
- Manage changes across heterogeneous operating systems.

Configuration Tracking

The Configuration Tracking feature tracks, backs up, and recovers critical software and system configuration information across Unix and Windows servers.

System administrators set up policies that describe the configuration files and databases to track, and the actions to take when a change in configuration is detected. Policies can be assigned to software, individual servers, groups of servers, and customers, and applied either locally or globally across data centers.

When Opsware SAS notices a server configuration change, it can log the change, notify administrators about the change with email, or back up the configuration, depending on the policy set by the administrator.

When a bad configuration change forces administrators to rollback to a previous version, they can use Opsware SAS to restore the configuration file to the saved version of the configuration. By notifying users about configuration changes – and maintaining a version history of those changes – organizations can quickly diagnose problems related to configuration errors and rollback to a known good state. In addition, this capability helps teams plug security holes inadvertently created by bad server configurations.

Typically, system administrators define configuration-tracking policies on a per-application basis. So for example, a policy for BEA WebLogic might specify, "Monitor the `weblogic.conf` file, notify app-server-admins@company.com of any changes, and maintain a version history of any changes that occur for 30 days." After a policy is defined in this fashion, administrators can apply the policy to all the WebLogic servers running in their environment or to specific servers.

Script Execution

The Script Execution feature enables you to share and run ad-hoc or custom scripts across an entire farm of Opsware-managed servers.

By executing scripts with Opsware SAS instead of manually, administrators benefit by using the following features:

- Parallel script execution across many Unix and/or Windows servers, saving time and ensuring consistency.
- Role-based access control, ensuring only authorized administrators can execute scripts on hosts to which they have access.
- The ability to control access to scripts by storing them in private or in public libraries.
- The ability to see and download script output one server at a time or in a consolidated report, which captures output from all servers in a single place.
- The ability for scripts to be mass-customized. Administrators can access information in Opsware SAS about the environment and the state of servers. This is critical to ensuring that the right scripts are executed on the right servers.
- A comprehensive audit trail that reports who, what, when, and where a particular script was executed.
- The ability to rollback changes (when used in conjunction with the Configuration Tracking feature).
- Automatic backup of all private and shared scripts to all other Opsware-managed data centers (when used in conjunction with an Opsware Multimaster Mesh).

Because the Script Execution feature is an integrated part of Opsware SAS, administrators enjoy unique benefits when compared to standalone script execution tools:

- Using known system state and configuration information to customize script execution, users can tailor each script by referencing and accessing the rich store of information in Opsware SAS, such as the customer or business that owns the server, whether the server is a staging or production server, which facility the server is located in, and custom name-value pairs.
- By sharing scripts without compromised security, users can share scripts with each other without compromising security because Opsware SAS maintains strict controls on who can execute scripts on which servers and generates a comprehensive audit trail of script execution.

Data Center Intelligence Reporting

Every change made to your managed servers is recorded in the Opsware SAS Model Repository. The Model Repository maintains precise information about the state and configuration of every server under your management.

You can now take advantage of this information though the Opsware SAS Data Center Intelligence Reporting (DCI) component. The DCI provides accurate, detailed, and up-to-date information about your operational environment. The DCI provides a new level of visibility into your operational environment that can help organizations make better decisions.

DCI reporting provides the following features and benefits:

- **Exact information about the latest system state and configurations**

  DCI reports display the most accurate and up-to-date information available, even during periods of frequent and substantial change. This level of accuracy reduces your risk of making the wrong decisions because of old data.

- **Visibility across the data center environment**

  Opsware SAS provides a comprehensive view across all operating systems and locations, allowing IT managers to generate on-demand snapshots driven from a single, high-quality data source. The ad-hoc capability allows you to view a variety of report types, filter by specific criteria, and display summary graphics or list views. In addition, a set of Quick Reports are pre-designed for one-click access to real-time information from the Reports Home page.

- **Accurate and detailed change history information**

  When a server's performance suddenly degrades, the best way to diagnose the cause is to learn the changes made to the server and who made the changes. Often, talking with the people who made the changes can help you understand the cause of the performance degradation.

  In most facilities, however, it's often difficult, if not impossible, to find out a server's exact change history, since records are not accurately kept. But Opsware SAS maintains a detailed record of each change: who made the change, what was the nature of the change, and when it occurred. This record is presented in a comprehensive series of reports; these reports can significantly reduce the time and effort in debugging server and software problems.

- **A comprehensive set of patch reports**

  One of the most time-consuming aspects of patching servers is identifying the vulnerable servers. Data collection for this task typically involves manually logging in to each server to see if it contains a particular version of software, what patches are already installed on the server, and what patches are *not* installed on the server. Opsware SAS helps administrators avoid this up-front effort by offering a comprehensive set of patch management reports.

- **The ability to extend the DCI reports**

  You can create new reports or modify the reports that ship with Opsware SAS. Opsware SAS provides the database necessary for creating reports.

  The Reports Home page checks for any new custom reports that you create, and presents them on the Reports Home page for easy access to all users. These reports are created by using the readily available Crystal Reports Designer 9.

  New reports can be extended to integrate with your own data sources (databases, spreadsheets, XML, and so forth), creating a powerful tool for more advanced data intelligence.

  See the *Opsware System DCI Administrator's Guide* for information about how to set up the DCI Reporting component.

  See the online Data Center Intelligence help and tutorial documentation for information about how to use and run the reports.

The Opsware Data Center Intelligence Reporting component is an optional component. By default, it is not installed with Opsware SAS. If this reporting component is not available for your organization, contact your Opsware Support Representative for information about how to obtain it so that you can generate reports. The DCI component must be installed and running in order to access the online documentation.

Discovery and Agent Deployment

The Opsware Discovery and Agent Deployment (ODAD) feature allows you to deploy Opsware Agents to a large number of servers in your facility and place them under Opsware management.

Using the ODAD features, you can perform the following tasks:

- Scan your network for servers.
- Select servers for Opsware Agent installation.
- Select a communication tool and provide user/password combinations.
- Choose agent installation options and deploy agents.

Server Explorer

The Server Explorer lets you view information about servers in your managed environment.

From the Server Explorer, you can perform the following tasks:

- Create a server snapshot, perform a server audit, audit application configurations, create a package, and open a remote terminal session on a remote server.
- Browse a server's file system, registry, hardware inventory, software and patch lists, and services.
- Browse Opsware information such as properties, configurable applications, and even server history.

From the Server Groups Browser, you can perform the following tasks:

- Audit system information, take a server snapshot, and configure applications.
- View and access group members (servers and other groups).
- View group summary and history information.

Server Compliance

The Opsware Server Compliance feature enables you to compare existing servers to a known working server and perform an audit that shows the difference between the two. With the audit results, you can investigate and identify servers that are not performing well. troubleshoot them and fix the malfunctioning servers.

Using Opsware Server Compliance, you can perform the following tasks:

- Compare servers or snapshots to reference servers or snapshots.
- Create compliance audits for repeated use.
- Associate audits with individual servers or dynamic server groups.
- Remediate problems at multiple levels, including files, directories, patches, registry keys, and packages.

Patch Management for Windows

The Patch Management for Windows feature enables you to identify, install, and remove Microsoft® Windows patches and maintain a high level of security across managed servers in your organization. With OCC Client user interface, you can identify and install patches that support security vulnerabilities for the Windows 2000, Windows 2003, and Windows NT4.0 operating systems. These patches include Service Packs, Update Rollups, and hotfixes.

Visual Packager

The Create Package feature allows you to create an installable software package from a managed server and from server compliance information, such as server snapshots and audit results. File system objects that are recorded in snapshots and compliance information produced by audits help you define the content of a package. In turn, you can use that package to update a server with new or missing server objects.

Create Package enables you to selectively package server objects according to the operating system of the servers that you want to distribute the package to, supporting both Unix and Windows operating systems.

It allows packages to contain the following objects:

- A Unix package can contain files (including attributes), directories, packages, patches, and patch clusters.
- A Windows package can contain files (including attributes), directories, packages, patches, Windows registry, and Windows services.

Application Configuration Management

Opsware Application Configuration Management (ACM) allows you to create templates so you can modify and manage application configurations associated with server applications. ACM enables you to manage, update, and modify those configurations from a central location, ensuring that applications in your facility are accurately and consistently configured.

Using ACM, you can perform the following tasks:

- Manage configurations based on files and objects, such as Windows registry, IIS metabase, WebSphere, COM+, and more.
- Preview configuration changes before applying them.
- Edit and push configuration changes to individual servers or server groups.
- Use information in the Opsware data model to set configuration values.
- Manage configurations of any application by building configuration templates.
- Audit the Application Configurations on a server to determine if any of the configuration files on the server are out of sync with the values stored in your templates.

Global Shell

The Opsware Global Shell feature enables you to manage servers by using a command-line interface. You can remotely perform the following tasks:

- Complete routine maintenance tasks on managed servers.
- Troubleshoot, identify, and remediate problems on managed servers.

Global Shell consists of a file system and a command-line interface to that file system for managing servers in Opsware SAS. The file system is known as the Opsware Global File System (OGFS). All object types in the OGFS (such as servers, customers, and facilities) are represented as directory structures in this file system.

The Global Shell feature also manages user permissions for accessing the file system, Windows Registry, and Windows Services objects on managed servers.

Opsware SAS Terms and Concepts

This section discusses the following topics:

- Agent-Server Architecture of Opsware Technology

- Opsware SAS Model-Based Approach
- Opsware SAS Model and Server Management
- Software Tree, Nodes, Packages, and Templates
- Software Tree Nodes and Server Management
- Packages and Server Management
- Templates and Server Management
- Server Management in Multiple Facilities
- Multimaster Support

Agent-Server Architecture of Opsware Technology

The agent-server architecture of Opsware SAS enables server management. The server portion of Opsware SAS consists of multiple, integrated components, each with a unique purpose. Each server managed by Opsware SAS runs an intelligent agent (the Opsware Agent).

Figure 1-4: Opsware SAS *Agent-Server Architecture*

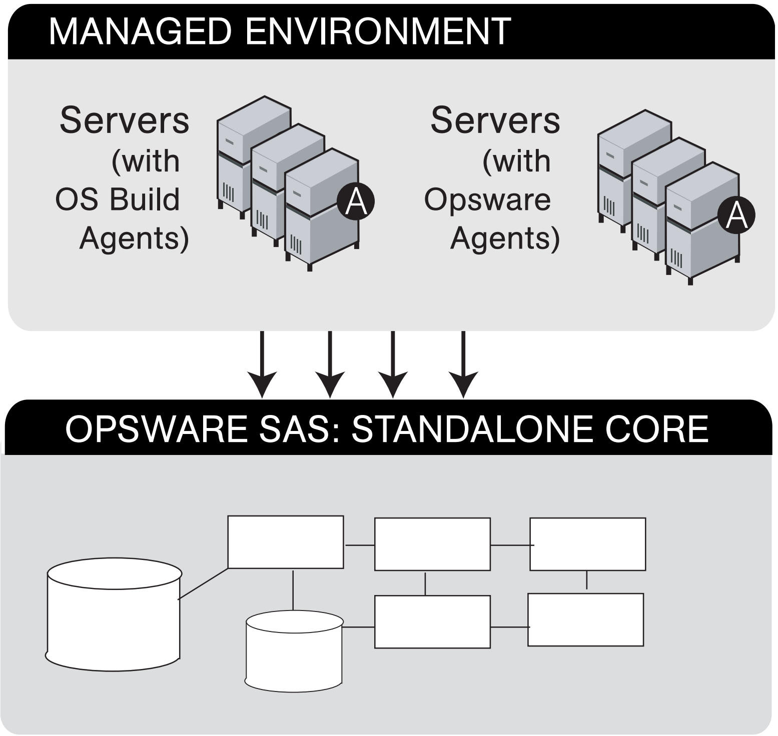
MANAGED ENVIRONMENT
Servers (with OS Build Agents)
A
Servers (with Opsware Agents)
A
OPSWARE SAS: STANDALONE CORE

See the *Opsware® SAS Administration Guide* for a detailed description of the components that make up the Opsware SAS server portion of the architecture.

The Opsware Agent is the agent of change on a server. Whenever Opsware SAS needs to make changes to servers, it does so by sending requests to the Opsware Agents. Depending on the request, the Opsware Agent on a server might use global Opsware SAS services in order to fulfill the request. For example, the Opsware Agent might often make requests to the Model Repository, the central database for all Opsware SAS components, and the Software Repository, the central repository for all software that Opsware SAS manages.

Some functions that the Opsware Agent supports are:

- Software installation and removal
- Configuration of software and hardware
- Periodically reporting server status
- Auditing of the server

An Opsware Agent is idle unless Opsware SAS is trying to perform some change on the server. In addition, each Opsware Agent periodically contacts the Data Access Engine and registers itself. The Data Access Engine is an XML-RPC interface to the model repository. The Data Access Engine sends this data to the Model Repository, which allows the Model Repository to keep track of server status, and know when particular servers are disconnected from or reconnected to the network.

After you install an Opsware Agent on a server, users can manage the server by installing or upgrading software, patching the OS software, removing software, changing server properties, or decommissioning the server.

See "Agent Reachability Communication Tests" on page 204 in Chapter 6 for information about how to install an Opsware Agent on a server so that Opsware SAS can manage it.

Opsware SAS Model-Based Approach

Opsware SAS employs a methodical, model-based approach, using a centralized information system as the starting point for any changes made to the physical environment. This information system approach to initiating change into the environment means that when team members need to make a change, they model the change first to verify and validate the change. Next, they let Opsware SAS propagate that change out to the environment in a secure, consistent, quick, and reliable manner.

Figure 1-5 gives a high-level view of the steps involved when users interact with the model.

Figure 1-5: High-Level View of User-Model Interaction

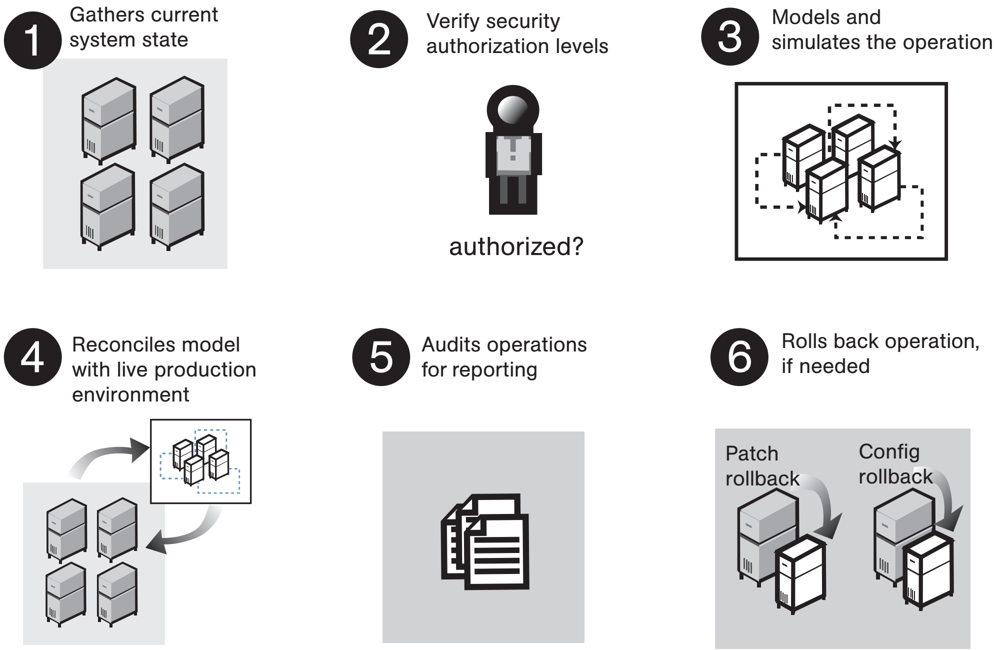
1
Gathers current system state
2
Verify security authorization levels
authorized?
3
Models and simulates the operation
4
Reconciles model with live production environment
5
Audits operations for reporting
6
Rolls back operation, if needed
Patch rollback
Config rollback

Opsware SAS includes three additional components to automate core functions. The role of these components is to provide policies, environment state information, and a comprehensive audit log for activity and asset reporting.

- **Software Tree**

  The Software Tree records a variety of information for software applications and operating systems, including, for example, data about how changes to a given software application might impact other existing applications.

- **Environment Tree**

  The Environment Tree manages characteristics about a customer's unique facility environment, including hardware, location of servers, network infrastructure, application names, business units, and service levels assigned to servers and applications. The information contained in the Environment Tree, combined with the data contained in the Software Tree, is used by Opsware SAS to model and simulate operational changes before they are executed in the production environment.

- **Modeling and Change Simulation Engine**

  Opsware SAS enables users to first model and simulate proposed operational changes to their environment before propagating these changes to production servers and applications. Utilizing the information contained in the Software and Environment Trees, the Modeling and Change Simulation Engine maintains a model of the various hardware and software configurations and other customer characteristics associated with each of the production servers under Opsware SAS control.

  Before committing any proposed changes to the production servers, this engine first conducts an impact analysis of the requested operation, enabling customers to test and validate changes prior to executing them in the production environment. This feature of Opsware SAS is designed to increase the success rates associated with initial deployments, improve the accuracy and security of these changes, and reduce application downtime.

Opsware SAS Model and Server Management

Opsware SAS employs a model-based approach to managing the operational environment. Users interact with the Opsware Command Center to accomplish OS provisioning, software provisioning, patch management, and server asset tracking for the operational environment. When users work in Opsware SAS, they work in the model, and Opsware SAS pushes the changes to the managed environment by reconciling the managed environment with the model. See Figure 1-6.

See “Opsware SAS Model-Based Approach” on page 68 in this chapter for information about how Opsware SAS is model-based.

Figure 1-6: Opsware SAS Model-Based Approach to Server Management

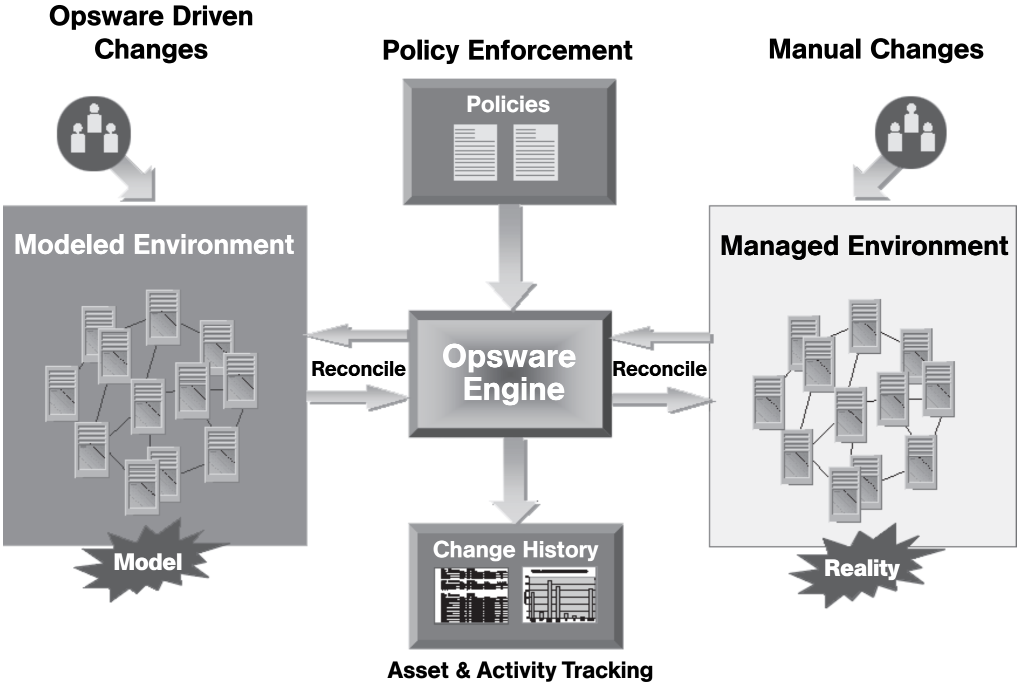

The Opsware Command Center visualizes the model as a tree called the Software Tree. By using the Software Tree, Opsware SAS enables users to embed technical best practices by specifying software installation order and by creating templates.

Before system administrators can use Opsware SAS to manage servers in the operational environment, the operational environment must be modeled in the Opsware Command Center.

A policy setter for your organization performs this goal by completing the following setup tasks:

- Defining the Software Tree applications that are available to install on the servers running in your environment

  See the *Opsware® SAS Configuration Guide* for more information about software provisioning setup.

- Uploading the packages that you will use in the operational environment

  See the *Opsware® SAS Configuration Guide* for more information about package management.

- Creating templates so that users can quickly install a complete software baseline on servers in the environment

  See the *Opsware® SAS Configuration Guide* for more information about creating templates.

- Setting up OS provisioning by preparing OS definitions for the operating systems that you need to install on your servers

  See the *Opsware® SAS Configuration Guide* for more information about OS provisioning setup.

- Setting up the patch management by uploading the required patches for your environment.

  See the *Opsware® SAS Configuration Guide* for more information about setting up patch management.

After a policy setter completes these setup tasks, end-users are ready to manage the servers running in the operational environment. These setup tasks do not need to be repeated by end users for them to manage servers.

Software Tree, Nodes, Packages, and Templates

To understand the way that Opsware SAS manages servers through the Opsware model, you need to understand the distinctions between the following elements in the model:

- The Software Tree and its nodes
- Packages uploaded to the Software Repository and added to nodes
- Templates that model complete software baselines

Software Tree Nodes and Server Management

The Software Tree organizes all the technology building blocks of your organization's environment, including the software and hardware. The Software Tree is made up of nodes and subnodes, modeling the interrelationships and dependencies among the software and customer accounts in the operational environment. Users navigate the Software Tree to perform specific operations.

Using nodes simplifies server management within Opsware SAS and the management of the software applications and configurations associated with those servers. Nodes are a hierarchical set of categories or types that classify software, configuration, and other components running in the operational environment.

The Software Tree provides the model for the operational environment, as shown in Figure 1-7.

Figure 1-7: Illustration of the Software Tree

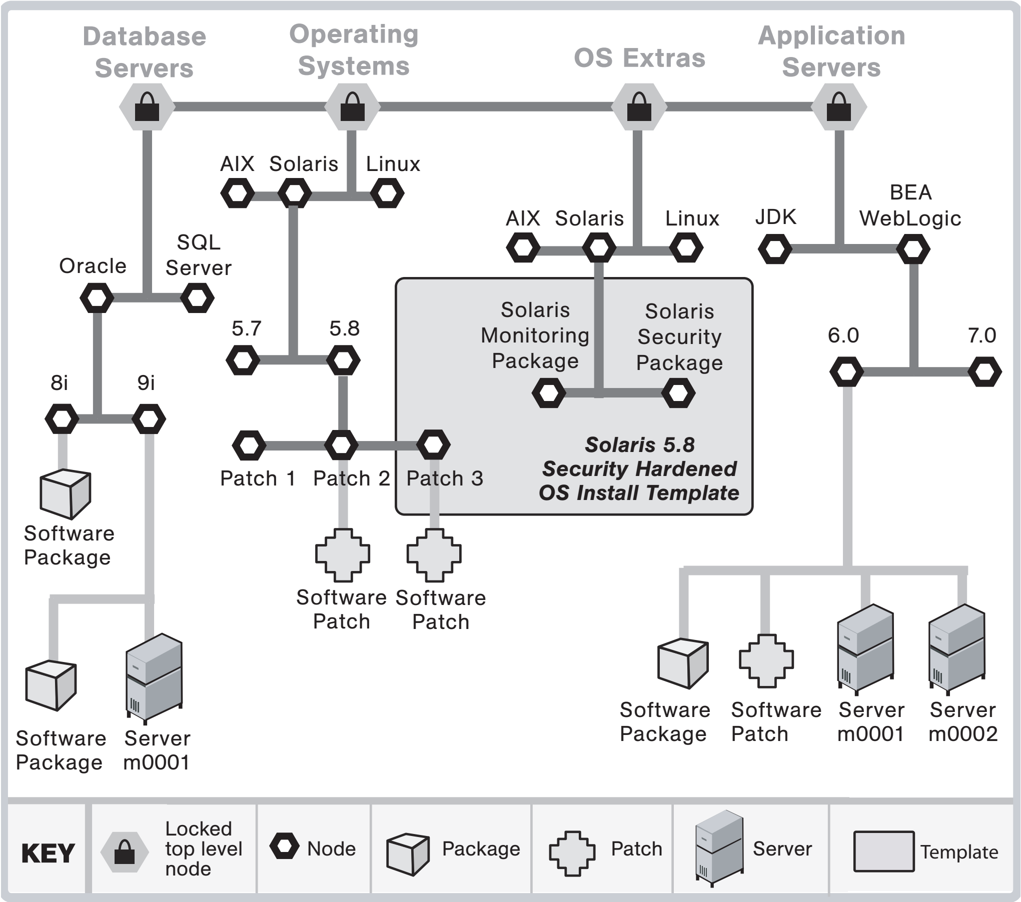

In the Opsware Command Center, you create nodes in the Software Tree. From the navigation panel, click Software ➤ Applications. The Applications page appears. As you navigate the Software Tree, you see the hierarchy of nodes in each category of applications, as shown in Figure 1-8.

Figure 1-8: Hierarchy of Nodes within an Application Category

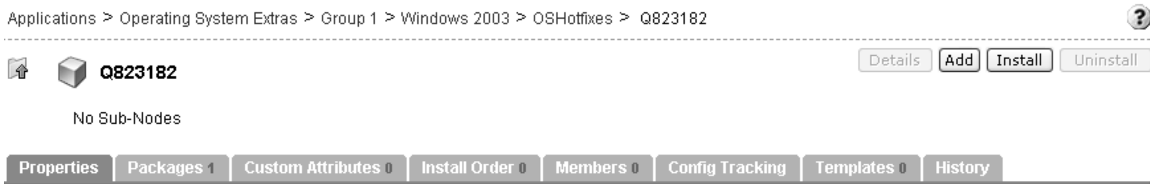

The Software Tree has the following characteristics:

- Each point in the Software Tree is called a node.
- The node information at the top of the Software Tree is more general, and as you travel farther down the tree, each successive subnode contains more detailed and specific information that relates to the node above it.
- A node inherits properties or software from the nodes above it.
- Users can add nodes and subnodes within each category.
- Software or a server might be assigned to a node depending on its location in the Software Tree.

  See the *Opsware® SAS Configuration Guide* for more information about setting up the Software Tree for software provisioning.

  The user's guide primarily documents how servers are automatically assigned to and removed from nodes when you use one of the Opsware wizards to install or uninstall software. Using an Opsware wizard is an efficient and easy method to assign or to remove servers from nodes.

  In certain situations, you might want to use the Software Tree and the reconcile operation to install or uninstall applications from managed servers.

  See "Assigning to and Removing Servers from Nodes" on page 710 in Appendix for more information.
- Assigning a server to a node determines the software that is installed on the server and its configuration.

Users perform the following tasks for servers by using nodes:

- Locating servers and viewing the nodes to which they are assigned
- Copying the configuration (assigned to nodes) of a server to other servers
- Assigning servers to and removing servers from nodes (so that you can install or uninstall software by running the reconcile operation).

  See “Reconcile” on page 699 in Appendix A for more information.

  See the *Opsware® SAS Configuration Guide* for more information about how to set up the Software Tree in the Opsware Command Center so that end users can install applications on managed servers in Opsware SAS.

Packages and Server Management

In Opsware SAS, packages contain software that is registered in the Software Repository. Packages contain software for operating systems, applications (for example BEA WebLogic, IBM WebSphere), databases, customer code, and software configuration information.

Packages are made available in Opsware SAS by uploading the packages to the Software Repository with the Opsware Command Center or by using the Opsware Command Line Interface.

The term package describes the collection of executables, configuration, or script files that are associated with an Opsware SAS-installable application or program.

See the *Opsware® SAS Configuration Guide* for more information about managing packages in Opsware SAS.

Templates and Server Management

Opsware templates allow you to group related sets of software so that they can be installed in a single operation by using the Opsware wizards. The two basic types of Opsware templates are:

- Templates that include the installation of an operating system
- Templates that do not include an operating system

You can, for example, use Opsware templates to quickly bring new servers into production. A template might include an operating system for a new server, the latest security patches for the operating system, plus all the applications required to run a full-fledged web service. Or, you can use templates to install a new service, made up of a set of applications and patches, on servers that are already in production.

Templates require little set up, and you can create them and deploy them quickly. To create a template, select packages, patches, or operating systems that are already configured and tested for installation, and add them to the template. You can later edit the template. For example, if a new patch is released, you can edit the template and add the patch to the template.

See the *Opsware® SAS Configuration Guide* for more information about creating templates.

Server Management in Multiple Facilities

The managed environment might span several facilities. A facility refers to the collection of servers that a single Opsware Model Repository manages, and the database that stores information about the managed environment. For example, one facility might be dedicated to an organization's Intranet, while another facility might be dedicated to the web services offered to the public. Your Opsware SAS can contain facilities (a full Opsware SAS is installed) and Satellite facilities. See Figure 1-9.

Figure 1-9: Server Management in Multiple Facilities

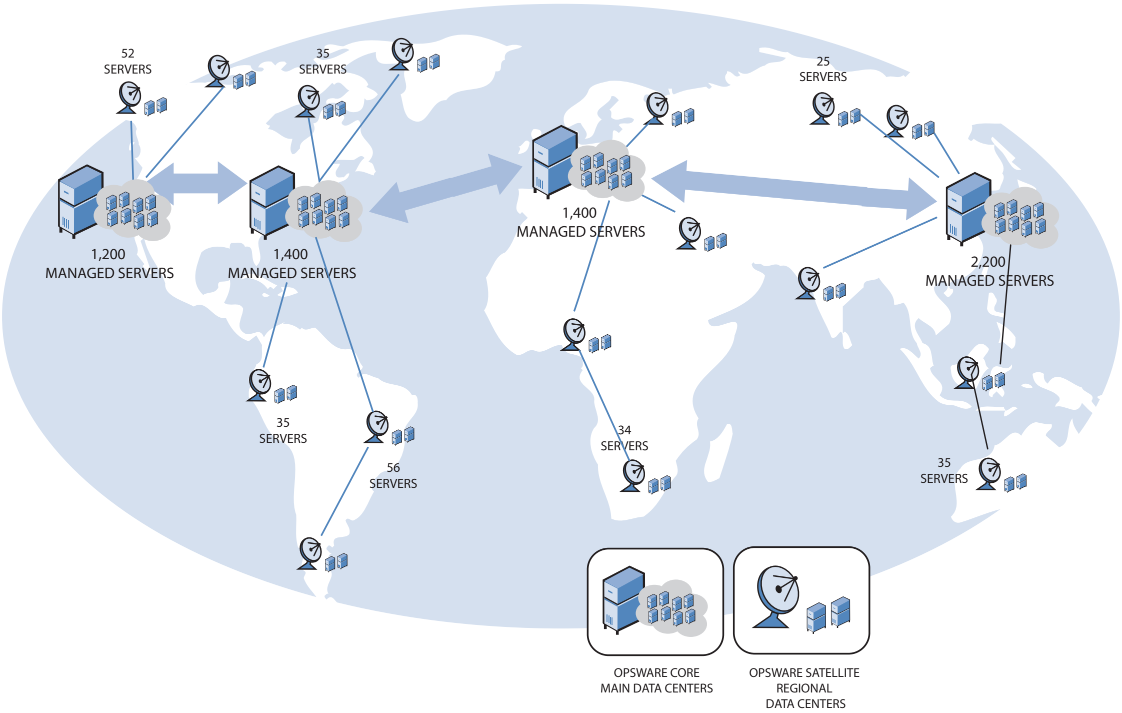

See the *Planning Deployments for Opsware® SAS* for more information about the types of installations that Opsware SAS supports.

Users can manage servers in any facility from an Opsware Command Center in any facility. When a user updates data in a facility, the Model Repository for that facility is synchronized with the Model Repositories located in all remote facilities.

When using Opsware technology in multiple facilities, users should follow these work process rules to reduce the chance of data conflicts between facilities:

- Users should not change data in one facility and then make the same change in another facility.
- More than one user should not change the same object in different facilities at the same time. For example, two users should not manage the same server from different facilities.

Multimaster Support

With the Opsware Model Repository Multimaster Component, customers can store and maintain a blueprint of software and environment characteristics of each data center (referred to as a facility in the Opsware Command Center) in multiple locations so the infrastructure can be easily rebuilt in the event of a disaster. The Multimaster Replication Engine not only provides the ability to replicate an environment in case of a disaster, but can also assist in facility migration activities as well as knowledge sharing across the enterprise.

Through the Model Repository Multimaster Component, Opsware SAS provides the ability to easily rebuild server and application environments, provision additional capacity, distribute updates, and share software builds, templates and dependencies – across multiple facilities and from one user interface. See Figure 1-10.

Figure 1-10: Multimaster Support

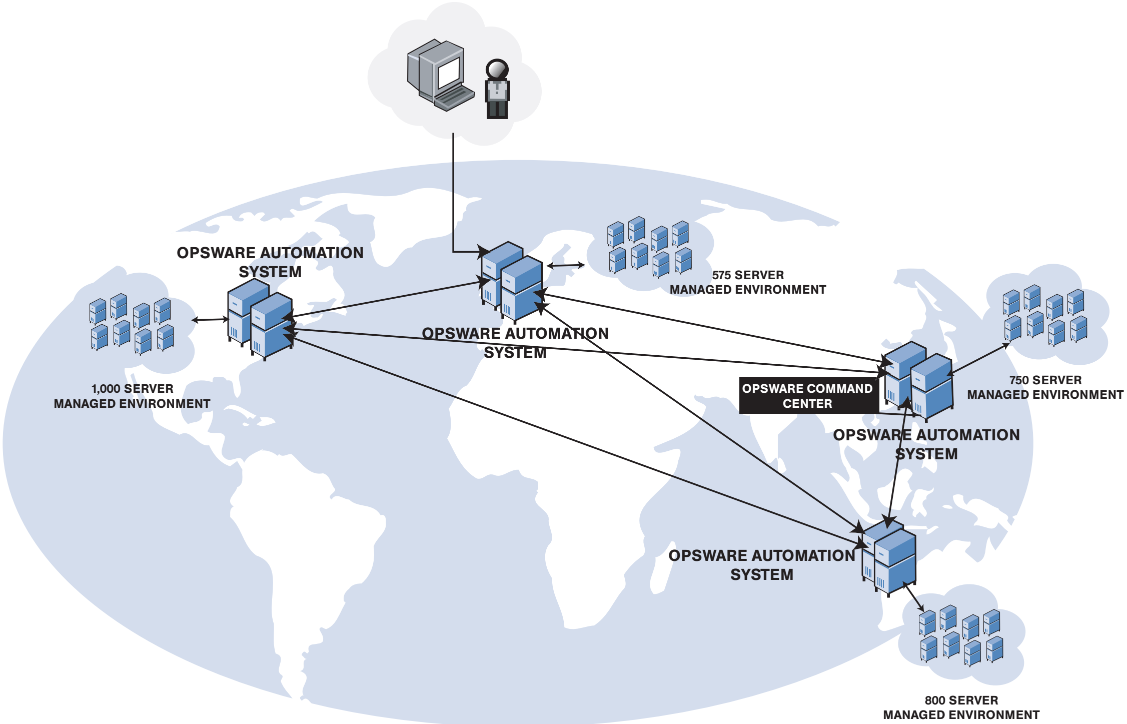
OPSWARE AUTOMATION SYSTEM
575 SERVER MANAGED ENVIRONMENT
1,000 SERVER MANAGED ENVIRONMENT
OPSWARE AUTOMATION SYSTEM
OPSWARE AUTOMATION SYSTEM
OPSWARE COMMAND CENTER
750 SERVER MANAGED ENVIRONMENT
OPSWARE AUTOMATION SYSTEM
OPSWARE AUTOMATION SYSTEM
800 SERVER MANAGED ENVIRONMENT

Chapter 2: Getting Started with the Opsware Command Center

IN THIS CHAPTER

This section provides an overview of the Opsware Command Center and discusses the following topics:

- Getting Started with the Opsware Command Center
- Access to Features in the Opsware Command Center
- Opsware Command Center User Interface

Getting Started with the Opsware Command Center

The following section describes Supported Browsers for the Opsware Command Center (a web application) and Browser Configuration Requirements.

Supported Browsers for the Opsware Command Center

The following table lists the supported browsers for the Opware Command Center.

Table 2-1: Supported Browsers for the Opsware Command Center

BROWSER	WINDOWS 2000	WINDOWS 2003	WINDOWS XP	LINUX 6.2+	SOLARIS 6 +	MAC OS X
Microsoft Internet Explorer 5.5	X					
Microsoft Internet Explorer 6.0	X	X	X			
Mozilla 1.6	X	X	X			

Table 2-1: Supported Browsers for the Opsware Command Center

BROWSER	WINDOWS 2000	WINDOWS 2003	WINDOWS XP	LINUX 6.2+	SOLARIS 6 +	MAC OS X
Firefox 1.0	X	X	X			

Browser Configuration Requirements

To run the Opsware Command Center, your browser must be configured in the following manner:

- The browser must accept cookies and be able to use Java.
- The browser must support SSL and should provide 128-bit encryption (recommended).
- Using a pop-up blocker might prevent some functions from working correctly. Either disable the pop-up blocker completely or use the supported browser's native pop-up blocking function instead of a third-party product.

Access to Features in the Opsware Command Center

This section explains how to configure your user profile to access features within the Opsware System. This section contains the following topics:

- System Management of Opsware User Information
- Best Practices for Selecting Passwords
- Updating User Profiles and Passwords

System Management of Opsware User Information

An Opsware administrator creates additional users who can use the Opsware Command Center, and the Opsware administrator assigns them temporary passwords. Once added to the system, an Opsware user can update their personal information, password, and time zone and date display preferences by using the My Profile link.

An Opsware user cannot change their access permissions by using the My Profile link. If an Opsware user needs additional access permissions, they contact their Opsware administrator.

Best Practices for Selecting Passwords

The Opsware Server Automation System (SAS) enforces a security policy that allows only authorized users to log into the Opsware Command Center. Opsware users are advised to select a password based on the following guidelines:

- Change your password frequently to ensure that your account information is secure.
- Select a password that is easy to remember so that you don't have to write it down.
- Use a mixture of upper and lower case letters, numbers, and punctuation in your password.
- Do not share your password.
- Do use a password that you can type quickly, without having to look at the keyboard.

Updating User Profiles and Passwords

As an Opsware user you can change your name, contact information, password, and preferences such as time zone and date display. You cannot change your access permissions. Contact your Opsware administrator to change your access permissions.

Perform the following steps to change your profile and password.

1. Log into the Opsware Command Center. The Opsware Command Center Home Page appears.
2. Click the My Profile link located at the top of the page. The Edit My Profile Page appears.
3. To change your profile information, enter new information in the Edit My Profile Page.
4. Click **Save**.
5. To change your password, click the Change Password Link. The Password Change page appears.
6. Enter your old password.
7. Enter your new password.
8. Confirm your new password in the Confirm Password field.
9. Click **Save**. A confirmation page appears indicating that your password was successfully changed.
10. Click **Okay**. The Edit Profile page appears.

Opsware Command Center User Interface

The following section discusses getting started with the Opsware System and contains the following topics:

- Requirements for Logging In
- Overview of the Opsware Command Center User Interface
- My Profile
- Search
- My Servers
- Mouseover Icon Tooltips

Requirements for Logging In

In order to log in and access Opsware Command Center features, your Opsware administrator must have created a login ID and password for you, and assigned user permissions that control the features you can access and the actions you can take when you use them.

Overview of the Opsware Command Center User Interface

The Opsware Command Center user interface consists of the five following sections:

- Home Page
- Tasks
- My Jobs
- My Customers
- Navigation Panel

Home Page

Figure 2-1 shows the Home page as it appears when you log in or when you click the Home link in the navigation bar.

Figure 2-1: Opsware System Home Page

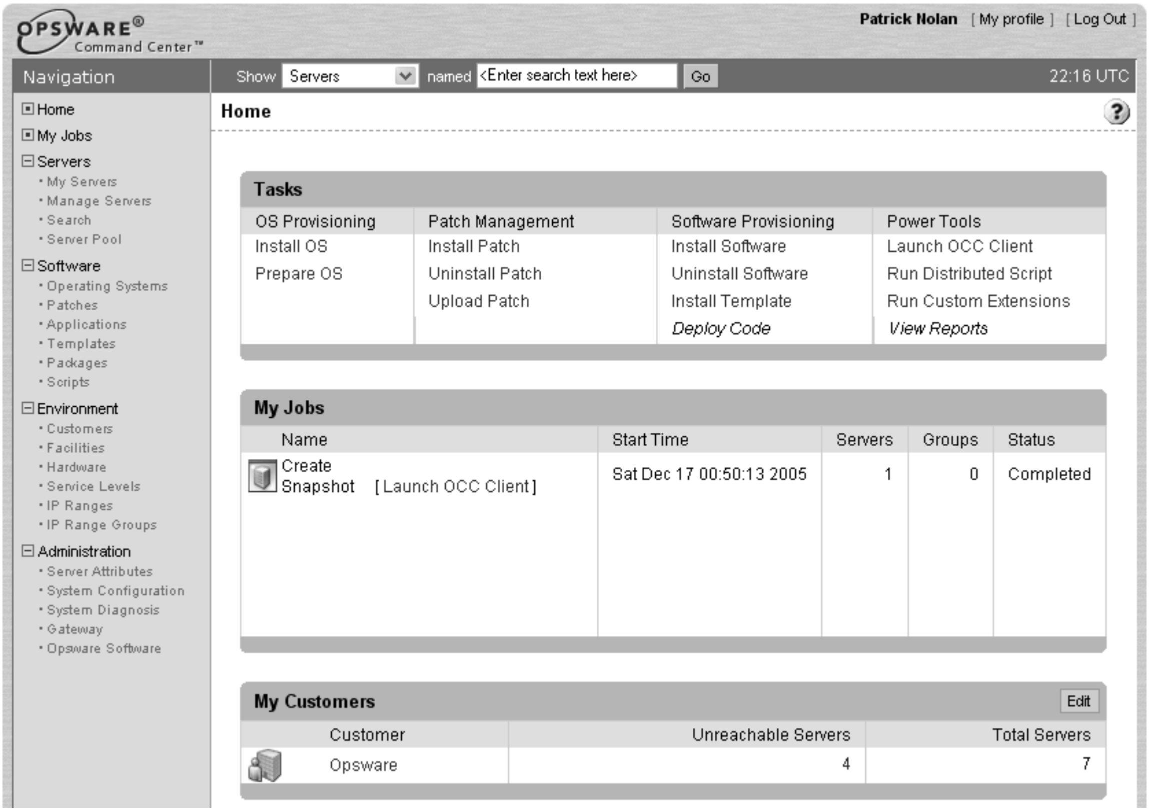

Show Servers named <Enter search text here>

Tasks

OS Provisioning	Patch Management	Software Provisioning	Power Tools
Install OS	Install Patch	Install Software	Launch OCC Client
Prepare OS	Uninstall Patch	Uninstall Software	Run Distributed Script
	Upload Patch	Install Template	Run Custom Extensions
		Deploy Code	*View Reports*

My Jobs

Name	Start Time	Servers	Groups	Status
Create Snapshot [Launch OCC Client]	Sat Dec 17 00:50:13 2005	1	0	Completed

My Customers

Customer	Unreachable Servers	Total Servers
Opsware	4	7

The time zone that appears in the upper right corner of the Home page is taken from the time zone preference that was defined for you when your profile was created. Consequently, the date and time information that displays throughout the Opsware Command Center is for that time zone. The occasional exceptions however are always labeled GMT.

Tasks

The Tasks area of the Home page displays links to the wizards that you have permissions to access, a link to the Data Center Intelligence reports if you have that module installed, and a link to the Code Deployment page if you have that permission. If you do not have permissions to a task in this area, the task name still displays, but it is italicized and it is not an active link. Figure 2-2 shows the Tasks area with all permissions enabled.

Figure 2-2: Tasks Area of the Home Page

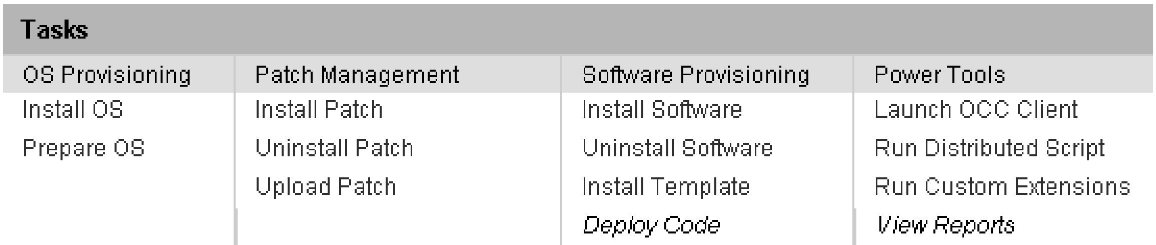

Tasks

OS Provisioning	Patch Management	Software Provisioning	Power Tools
Install OS	Install Patch	Install Software	Launch OCC Client
Prepare OS	Uninstall Patch	Uninstall Software	Run Distributed Script
	Upload Patch	Install Template	Run Custom Extensions
		Deploy Code	*View Reports*

My Jobs

The My Jobs area of the Home page is populated with details of the jobs that you have run, jobs that are currently in progress, or jobs that you have scheduled to run, including the name of the job, the start time, the number of servers affected by the job, and the status of the job. If there are more than six jobs, you can see the rest of them by clicking the See All link, which also shows the total number of jobs in parentheses, as Figure 2-3 shows.

If the job was run in the OCC Client, then there will be a link next to the job that (when clicked) will launch the OCC Client. You can view the more detailed information about the job in the OCC Client.

Figure 2-3: My Jobs Area of the Home Page

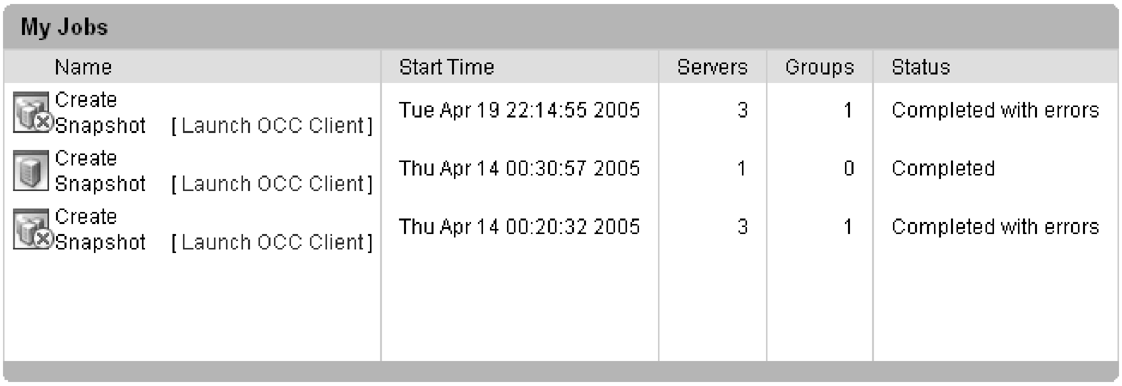

My Jobs

Name	Start Time	Servers	Groups	Status
Create Snapshot [Launch OCC Client]	Tue Apr 19 22:14:55 2005	3	1	Completed with errors
Create Snapshot [Launch OCC Client]	Thu Apr 14 00:30:57 2005	1	0	Completed
Create Snapshot [Launch OCC Client]	Thu Apr 14 00:20:32 2005	3	1	Completed with errors

My Customers

The My Customers area of the Home page is populated with customer information, including unreachable servers associated with a customer and the total number of servers associated with that customer. To select the customers to display in the My Customers area of the Home page, click **Edit** and select the check box next to the customer name. See Figure 2-4.

Figure 2-4: My Customers Area of the Home Page

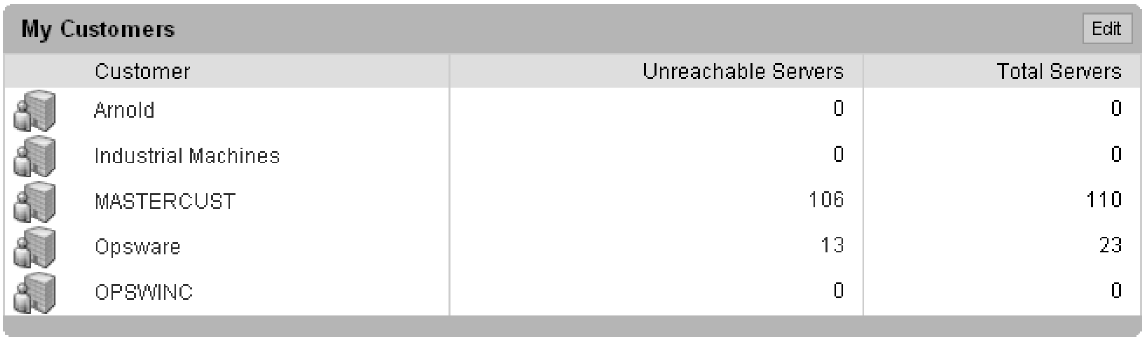

My Customers

Customer	Unreachable Servers	Total Servers
Arnold	0	0
Industrial Machines	0	0
MASTERCUST	106	110
Opsware	13	23
OPSWINC	0	0

Navigation Panel

The navigation panel on the left side of the Opsware Command Center shows all possible features, as shown in Figure 2-5. The features you can access and the actions you can perform depend upon your user profile, as defined by the Opsware administrator.

Figure 2-5: Navigation Panel, All Permissions View

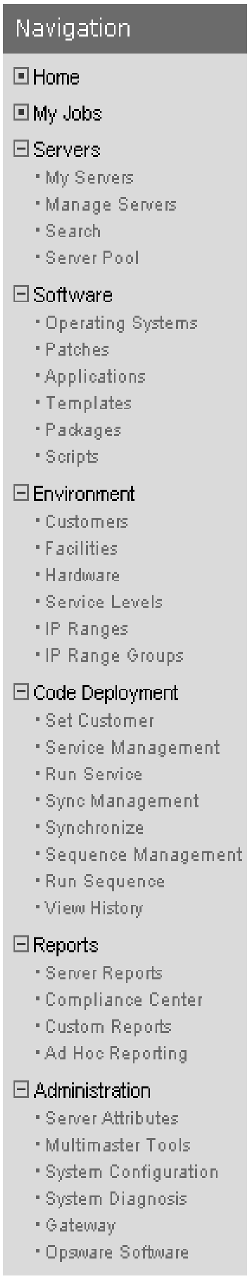

The items that appear in the navigation panel depend on the permissions the user has. Clicking an item on the navigation panel displays that feature in the main part of the Home page. For a user with all permissions, the following links appear:

Home: Displays the top level of the Opsware Command Center. The Home page is described in this section of the guide. Wizards are documented in their respective functional areas of the system.

My Jobs: Displays the My Jobs page, showing jobs completed during the previous 30 days, jobs currently in progress, and currently scheduled jobs. This page is an expanded view of the contents of the My Jobs area of the Home page and has the same effect as clicking **Show All** in the My Jobs area of the Home page.

Servers: Expands to display these selections:

- **My Servers**: Use to add any server or server group to your own personal view of servers. My Servers provides an efficient way to manage servers when your operational environment contains hundreds or thousands of servers.
- **Manage Servers**: Use filters to display a list of servers and perform operations on them such as edit server values, assign, reconcile, run scripts, and add servers to My Servers. See “Server Management in Opsware Command Center” on page 219 in Chapter 7 for more information.
- **Search**: Find specific servers using default criteria or user-defined criteria. See “Server Search” on page 124 in Chapter 4 for more information.
- **Server Pool**: Use filters to display a list of servers, install operating systems on the servers, and delete the servers. See “Operating System Provisioning” on page 567 in Chapter 15 for more information.

Software expands to display these selections:

- **Operating Systems**: Prepare operating systems for installation and delete existing operating systems. See “Operating System Provisioning” on page 567 in Chapter 15 for more information.
- **Patches**: Prepare Unix patches for installation and define Unix patch preferences. See See “Patch Management for Unix” on page 519 in Chapter 13 for more information.
- **Applications**: Define nodes in the Software Tree, identify packages to attach to nodes, and assign servers to nodes. See the *Opsware® SAS Configuration Guide* for more information about software provisioning setup.

- **Templates**: Define templates and associate operating systems, patches, applications and service levels. See the *Opsware® SAS Configuration Guide* for more information about software provisioning setup.
- **Packages**: Upload packages, browse uploaded packages, and search for packages. See the *Opsware® SAS Configuration Guide* for more information about Package Management.
- **Scripts**: Run scripts, upload scripts, and create new scripts. See "Distributed Script Execution" on page 411 in Chapter 11 for more information.

Environment expands to display these selections:

- **Customers**: Create new customers and edit or delete existing customers. See the *Opsware® SAS Configuration Guide*.
- **Facilities**: Create new facilities, edit facility properties, and assign and edit custom attributes for facilities. See the *Opsware® SAS Deployment and Installation Guide*.
- **Hardware**: A read-only view of servers categorized by the hardware manufacturer and model, and their related information. See "Hardware Information for Managed Servers" on page 155 in Chapter 4 for more information.
- **Service Levels**: Define service levels and custom attributes. See the *Opsware® SAS Configuration Guide* for more information.
- **IP Ranges**: Identify and create IP ranges and IP range types. See the *Opsware® SAS Configuration Guide* for more information.
- **IP Range Groups**: Create IP Range Groups. See the *Opsware® SAS Configuration Guide* for more information.

Code Deployment expands to display these selections:

- **Deployment Home**: The exact CDR links that you see in the Code Deployment area are based on the permissions that you have for the customer you want to work with.
- **Service Management**: Create, modify, and delete service definitions. Services define the location and commands to manipulate applications on hosts.
- **Run Service**: Perform service operations on one or more hosts, or request that a service operation be performed on your behalf. Service operations include starting or stopping applications, cutting over or rolling back code, and backing up or restoring code.

- **Sync Management**: Create, modify, and delete synchronization definitions. Synchronizations define the path for pushing code from a source host to one or more destination hosts.
- **Synchronize**: Perform a synchronization to one or more hosts, or request that a synchronization be performed on your behalf.
- **Sequence Management**: Create, modify, and delete sequence definitions. Sequences allow the grouping of service operations and synchronization operations to define higher level code deployment operations.
- **Run Sequence**: Perform a pre-defined sequence of service operations and synchronizations on one or more hosts, or request that a sequence be performed on your behalf.
- **View History**: Get information about previously run Code Deployment operations. See “Code Deployment and Rollback” on page 671 in Chapter 18 for more information.

Reports: Displays the Data Center Intelligence top-level page (if this module is installed). See *Opsware System DCI Administrator’s Guide* for more information.

Administration: Expands to display the following features. For more information about these features, see the *Opsware® SAS Administration Guide* and *Opsware® SAS Configuration Guide*.

- **Users & Groups**: Administrators use this feature to create user groups, define permissions for those groups, create new administrators, and add users to groups.
- **Server Attributes**: Define and edit server use attributes, enable them for code deployment, and define deployment stages. Also define and edit server deployment stage attributes.
- **System Configuration**: Contains the configuration parameters that define how the Opsware System works in your environment. This selection is only used at the direction of Opsware Inc.
- **System Diagnosis**: Runs a series of tests on Opsware components to make sure that they are functioning correctly.
- **Gateway**: Allows you to connect Satellites with this or other cores.
- **Opsware Software**: Provides a view of the properties, custom attributes, installation order, and history of the software attached to the Opsware nodes in the system.

The Administration set of features is only available if you are logged in as an Opsware administrator.

Opsware Command Center Navigation

Top-level navigation from the Home page is simple. To access any of the features in the navigation panel, click the feature name. To access any of the wizards or other features in the Tasks area of the Home page, click the name of the task.

After you select a task or a feature, other pages appear, which might have one or more of the following means of navigation:

- Clicking a hyperlinked name to display a page

  For example, if you select Software ➤ Applications (assuming that you have the correct permissions), the page that appears shows all of the applications that have been uploaded so far. Each application name is a hyperlink that takes you to another page with information about that application already displayed.

- Selecting a tab to display a page

  For example, when you select Applications and click one of the hyperlinked application names, the resulting page shows a row of tabs, like Figure 2-6 shows.

Figure 2-6: Example of Tabs

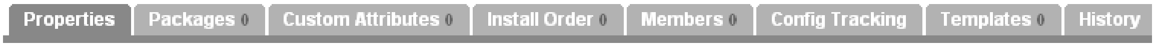

  Each tab displays a page, each with its own buttons and functionality.

- Clicking a button to display a page

  For example, when you select the Custom Attributes tab, a page appears with several buttons: **Add**, **Delete**, and **Copy**, which are common to each page called by these tabs, and **Add Custom Attribute**, which is unique to this particular tab. You will find similar functionality on all tabbed pages in the Opsware System.

My Profile

You can update your own personal user information without the assistance of the Opsware administrator with the My Profile link. You can change your first and last name, your contact information, your password, and your time zone and date display preferences.

Search

At the top of the Home page open the dialog box as Figure 2-7 shows.

Figure 2-7: Opsware Command Center Search Function

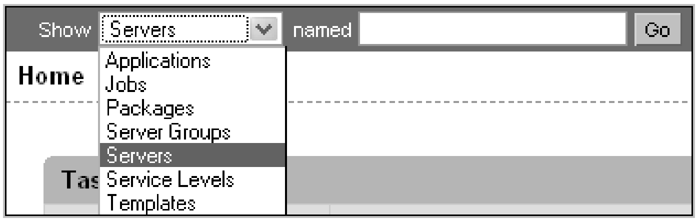

You can search for servers, applications (nodes), packages, jobs, server groups, service levels, and templates by making the selection from the drop-down list, and then entering an identifying string in the text box.

My Servers

The My Servers feature provides a convenient place to store a set of servers that have been selected and stored using the Add to My Servers function in Manage Servers. You might use it as a shortcut to the servers you work with most often, or as a way to gather a group of servers when you want to apply the same changes to all of them. All functions that are available in the Manage Servers page are also available from within My Servers.

Mouseover Icon Tooltips

When an icon appears on a page in the Opsware Command Center, a tooltip displays information about the icon when your mouse pointer hovers over it. For example, server icons display messages such as “Available or Build Failed” to describe the state of the server. Packages and patches display messages such as Available, Managed, Unmanaged, and so forth.

Chapter 3: Getting Started with OCC Client

IN THIS CHAPTER

This section discusses the following topics:

- Overview of the OCC Client
- Requirements for Running the OCC Client
- Installing and Launching the OCC Client
- OCC Client User Interface
- Accessing OCC Client Options

Overview of the OCC Client

The OCC Client is a powerful, speedy Java client for the Opsware Server Automation System which provides the responsiveness and look-and-feel of a Microsoft Windows desktop application with the cross-platform flexibility of Java. Over time, more and more of the Opsware Server Automation System's features will be available via the OCC Client.

The OCC Client provides the following features:

- Discovery and Agent Deployment
- Server Explorer
- Server Compliance
- Patch Management for Windows
- Visual Packager
- Application Configuration Management
- Global Shell

Discovery and Agent Deployment

The Opsware Discovery and Agent Deployment (ODAD) feature allows you to deploy Opsware Agents to a large number of servers in your facility and place them under Opsware management.

Using the ODAD features, you can perform the following tasks:

- Scan your network for servers.
- Select servers for Opsware Agent installation.
- Select a communication tool and provide user/password combinations.
- Choose agent installation options and deploy agents.

Server Explorer

The Server Explorer lets you view information about servers in your managed environment.

From the Server Explorer, you can perform the following tasks:

- Create a server snapshot, perform a server audit, audit application configurations, create a package, and open a remote terminal session on a remote server.
- Browse a server's file system, registry, hardware inventory, software and patch lists, and services.
- Browse Opsware information such as properties, configurable applications, and even server history.

From the Server Groups Browser, you can perform the following tasks:

- Audit system information, take a server snapshot, and configure applications.
- View and access group members (servers and other groups).
- View group summary and history information.

Server Compliance

The Opsware Server Compliance feature enables you to compare existing servers to a known working server and perform an audit that shows the difference between the two. With the audit results, you can investigate and identify servers that are not performing well. troubleshoot them and fix the malfunctioning servers.

Using Opsware Server Compliance, you can perform the following tasks:

- Compare servers or snapshots to reference servers or snapshots.
- Create compliance audits for repeated use.
- Associate audits with individual servers or dynamic server groups.
- Remediate problems at multiple levels, including files, directories, patches, registry keys, and packages.

Patch Management for Windows

The Patch Management for Windows feature enables you to identify, install, and remove Microsoft® Windows patches and maintain a high level of security across managed servers in your organization. With OCC Client user interface, you can identify and install patches that support security vulnerabilities for the Windows 2000, Windows 2003, and Windows NT4.0 operating systems. These patches include Service Packs, Update Rollups, and hotfixes.

Visual Packager

The Create Package feature allows you to create an installable software package from a managed server and from server compliance information, such as server snapshots and audit results. File system objects that are recorded in snapshots and compliance information produced by audits help you define the content of a package. In turn, you can use that package to update a server with new or missing server objects.

Create Package enables you to selectively package server objects according to the operating system of the servers that you want to distribute the package to, supporting both Unix and Windows operating systems.

It allows packages to contain the following objects:

- A Unix package can contain files (including attributes), directories, packages, patches, and patch clusters.
- A Windows package can contain files (including attributes), directories, packages, patches, Windows registry, and Windows services.

Application Configuration Management

Opsware Application Configuration Management (ACM) allows you to create templates so you can modify and manage application configurations associated with server applications. ACM enables you to manage, update, and modify those configurations from a central location, ensuring that applications in your facility are accurately and consistently configured.

Using ACM, you can perform the following tasks:

- Manage configurations based on files and objects, such as Windows registry, IIS metabase, WebSphere, COM+, and more.
- Preview configuration changes before applying them.
- Edit and push configuration changes to individual servers or server groups.
- Use information in the Opsware data model to set configuration values.
- Manage configurations of any application by building configuration templates.
- Audit the Application Configurations on a server to determine if any of the configuration files on the server are out of sync with the values stored in your templates.

Global Shell

The Opsware Global Shell feature enables you to manage servers by using a command-line interface. You can remotely perform the following tasks:

- Complete routine maintenance tasks on managed servers.
- Troubleshoot, identify, and remediate problems on managed servers.

Global Shell consists of a file system and a command-line interface to that file system for managing servers in Opsware SAS. The file system is known as the Opsware Global File System (OGFS). All object types in the OGFS (such as servers, customers, and facilities) are represented as directory structures in this file system.

The Global Shell feature also manages user permissions for accessing the file system, Windows Registry, and Windows Services objects on managed servers.

Requirements for Running the OCC Client

Java J2SE v 1.4.2 JRE must be installed on the system that runs the OCC Client. To download this version of Java, go to http://java.sun.com/j2se/1.4.2/download.html. The OCC Client is supported on the following operating systems:

- Windows 2003
- Windows 2000
- Windows XP

When using Internet Explorer, if you try to launch the OCC Client without installing Java J2SE v 1.4.2 JRE, you are prompted to install JRE and run the OCC Client in one step. When you select this option, Java v 1.4.2 JRE is installed and you get the following error message: `Java Web Start -unexpected Error. Unable to launch OCC Client.` To launch the OCC Client, ignore the error message and click **Retry**.

Installing and Launching the OCC Client

To access the OCC Client for the first time, you must launch the OCC Client installer from the Opsware Command Center. You only need to install the OCC Client once.

Once you have downloaded and installed the OCC Client, you can start it from the following locations:

- In the Power Tools section of the Opsware Command Center home page
- On your desktop (an OCC Client program icon is installed to your desktop if you choose to install one on your desktop)
- From the **Start** menu select ➤ **All Programs** ➤ **OCC Client**

By default, the core you log into is the core of the Opsware Command Center where you launched the OCC Client download. You can change the default core by configuring your general options. See “OCC Client Options” on page 105 in this chapter for more information.

To install and launch the OCC Client, perform the following steps:

1. Launch the OCC and log in.
2. From the Home page of the OCC, in the Power Tools section, click the Launch OCC Client link.
3. If you are a Firefox user, you will be able to open the JSP file to launch the OCC Client installation. (Internet Explorer users should go to the next step.)
4. After the OCC Client has downloaded to your computer, the Login window displays, as seen in Figure 3-1.

Figure 3-1: Login the OCC Client Dialog Box.

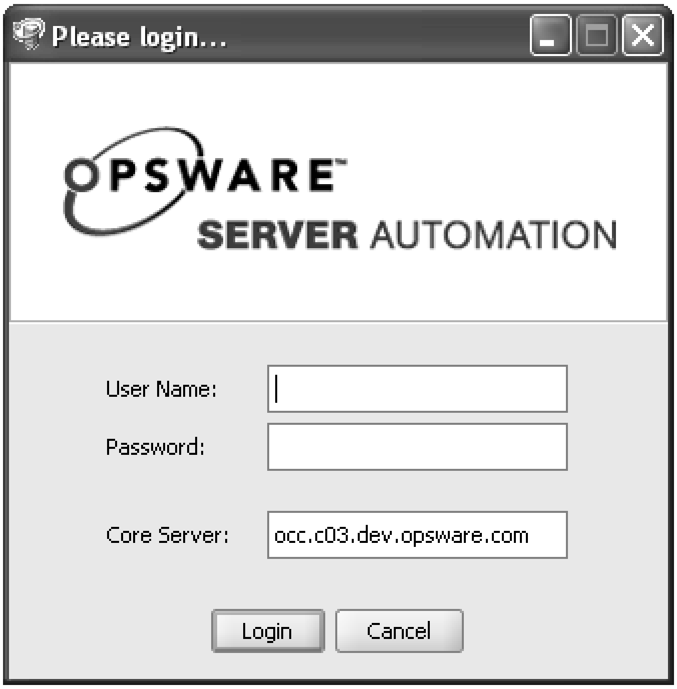

User Name:

Password:

Core Server: occ.c03.dev.opsware.com

Enter your user name and password, and click **Login**. If you have access to more than one core server, you can enter the core server's IP address or name in the core server field. If you don't specify a port with the host:port notation, port 443 is used.

5. If you are asked to accept the certificate from the core server, click **Yes**. The OCC Client now appears.

OCC Client User Interface

The OCC Client user interface provides easy access to all of the OCC Client features and functionality. The OCC Client user interface has six main areas:

- Menus
- Navigation Pane
- Content Pane
- Preview Pane
- Status Bar

Figure 3-2: OCC Client User Interface

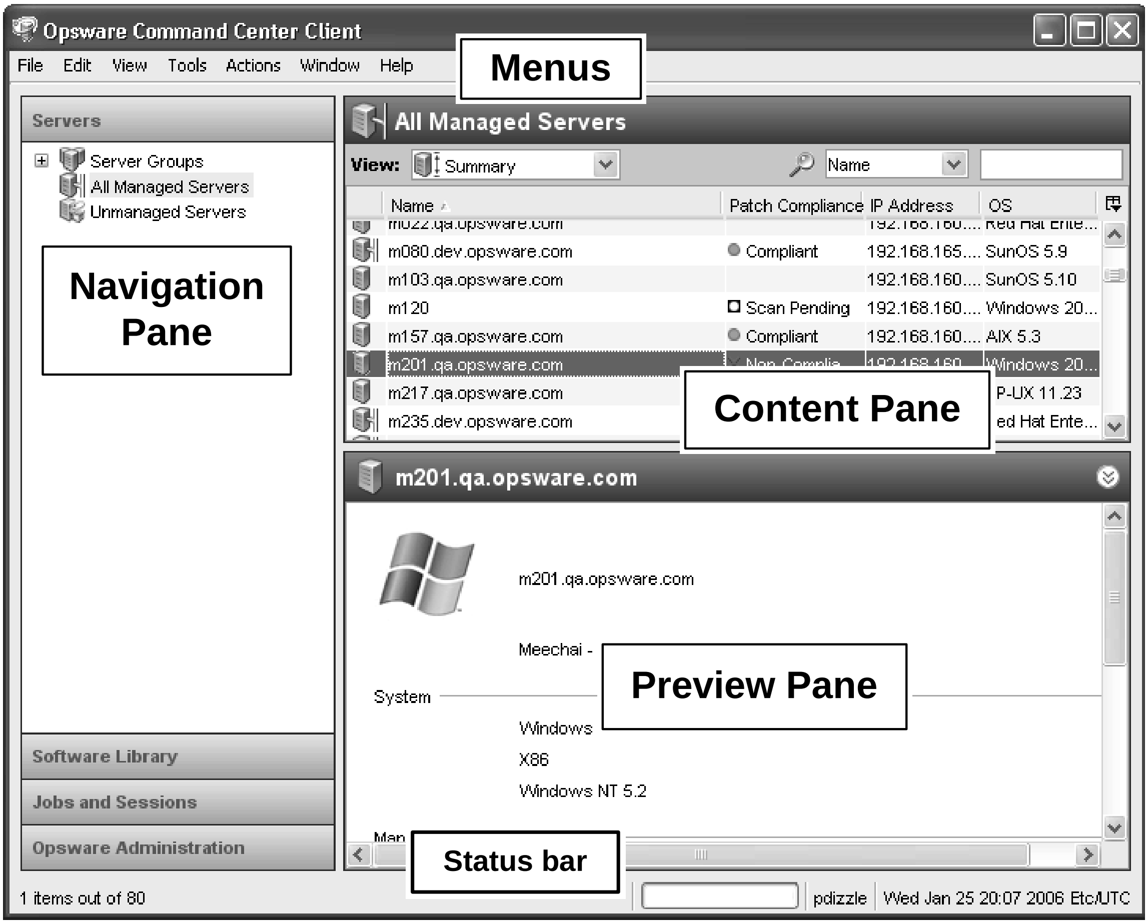

Name	Patch Compliance	IP Address	OS
m022.qa.opsware.com		192.168.160....	Red Hat Ente...
m080.dev.opsware.com	● Compliant	192.168.165....	SunOS 5.9
m103.qa.opsware.com		192.168.160....	SunOS 5.10
m120	□ Scan Pending	192.168.160....	Windows 20...
m157.qa.opsware.com	● Compliant	192.168.160....	AIX 5.3
m201.qa.opsware.com	✓ Non Complia...	192.168.160...	Windows 20...
m217.qa.opsware.com			P-UX 11.23
m235.dev.opsware.com			ed Hat Ente...

m201.qa.opsware.com

m201.qa.opsware.com

Meechai -

System

Windows

X86

Windows NT 5.2

Man

1 items out of 80 | pdizzle | Wed Jan 25 20:07 2006 Etc/UTC

Menus

OCC Client menus include the following menus:

- **File**: This allows you to open a new OCC Client window, or close the current window, or exit all open OCC Client windows.
- **Edit**: This allows you to cut, copy, paste and delete text.
- **View**: This refreshes the current view and shows the latest information from the core that you are currently logged into. From here, you can also access OCC Client features in the Navigation pane, such as Servers (server groups, managed and unmanaged servers), Software Library (application configuration, server compliance, patch management, Jobs and Sessions (job logs and shell sessions), and Opsware Administration (patch configuration and compliance rules). This also allows you to show or hide the Action pane and the Preview pane.
- **Tools**: This allows you to open a Global Shell session or access the OCC Client options.
- **Actions**: Depending upon the feature that you have selected in the Navigation panel, this menu allows you to perform numerous functions related to all main OCC Client features.
- **Window**: This allows you to access multiple instances of OCC Client windows, if more than one window is open.
- **Help**: This menu provides help for the OCC Client. Help F1 provides context-sensitive help relevant to the current feature window or dialog box selected or opened (same as F1). The contents and index will open the OCC Client help system to the main table of contents. (About OCC Client provides version and system information.)

Navigation Pane

To access OCC Client features, select a feature in the Navigation pane, as shown in Figure 3-3. When you select a feature, the contents of it appear in the Contents pane. You can access functions related to it through the Actions menu or the Actions pane.

Figure 3-3: Navigation Pane

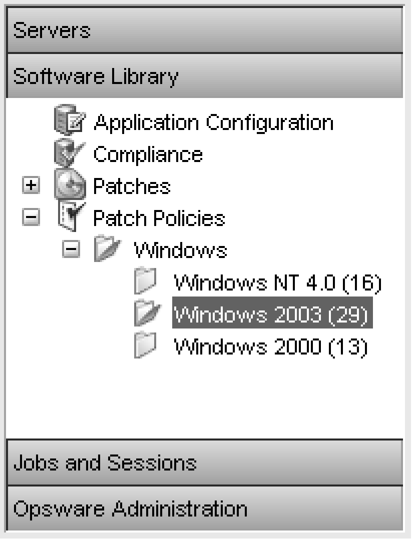

Content Pane

Depending on the selection in the Navigation pane, the Content pane lists the following information:

- All managed servers and server groups, including unmanaged servers
- Compliance snapshots and audits
- Patches and patch policies
- Application Configurations and Application Configuration Templates
- Agent deployment information
- Jobs that the user has run
- Access to the Global Shell sessions
- Patch configuration and patch compliance rules

You can perform actions on features in the Content area using the Action Menu or Action Pane, or you can right-click to perform various actions or double-click to open.

Figure 3-4: Feature Content Window Showing Managed Server

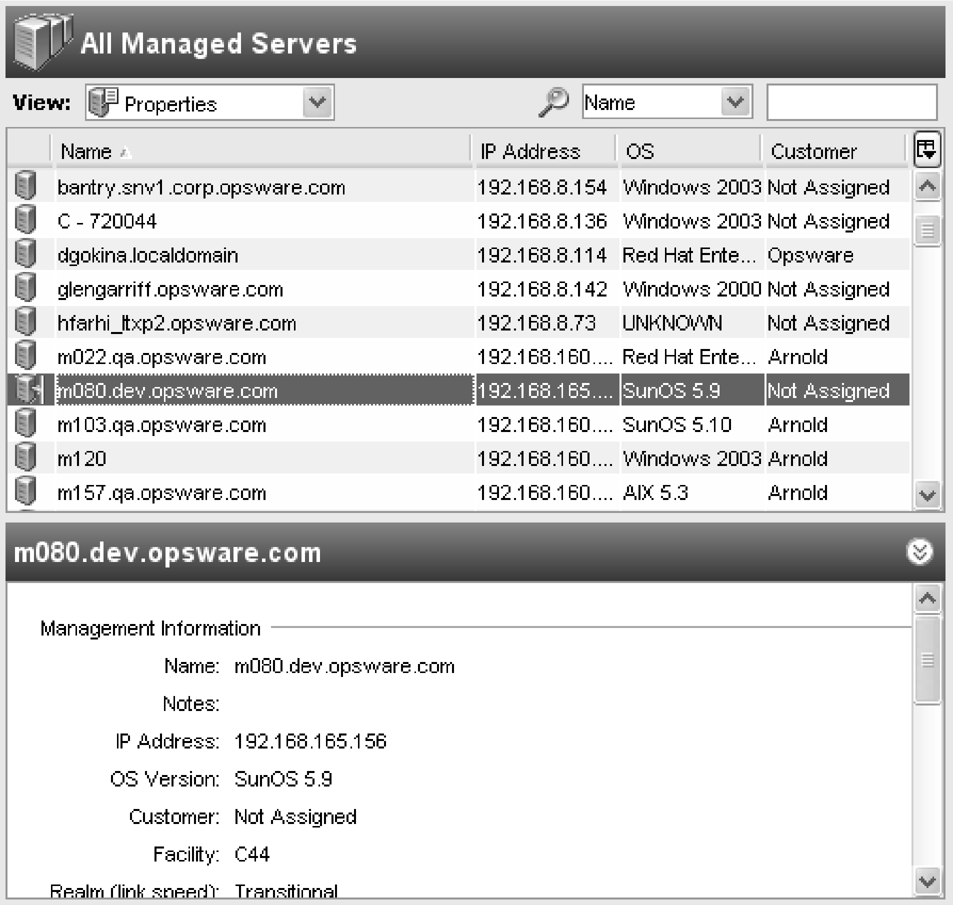

Name	IP Address	OS	Customer
bantry.snv1.corp.opsware.com	192.168.8.154	Windows 2003	Not Assigned
C - 720044	192.168.8.136	Windows 2003	Not Assigned
dgokina.localdomain	192.168.8.114	Red Hat Ente...	Opsware
glengarriff.opsware.com	192.168.8.142	Windows 2000	Not Assigned
hfarhi_ltxp2.opsware.com	192.168.8.73	UNKNOWN	Not Assigned
m022.qa.opsware.com	192.168.160....	Red Hat Ente...	Arnold
m080.dev.opsware.com	192.168.165....	SunOS 5.9	Not Assigned
m103.qa.opsware.com	192.168.160....	SunOS 5.10	Arnold
m120	192.168.160....	Windows 2003	Arnold
m157.qa.opsware.com	192.168.160....	AIX 5.3	Arnold

m080.dev.opsware.com

Management Information

Name: m080.dev.opsware.com

Notes:

IP Address: 192.168.165.156

OS Version: SunOS 5.9

Customer: Not Assigned

Facility: C44

Realm (link speed): Transitional

Content Pane Tools

From inside the Content pane, you can perform the following actions:

- With the View drop-down list, you can change the view of a selected feature. For example, you can select a patch from the Navigation pane, and then from the View drop-down list, choose Servers to see all the servers that the patch is attached to as shown in Figure 3-5.

Figure 3-5: View Drop-down List

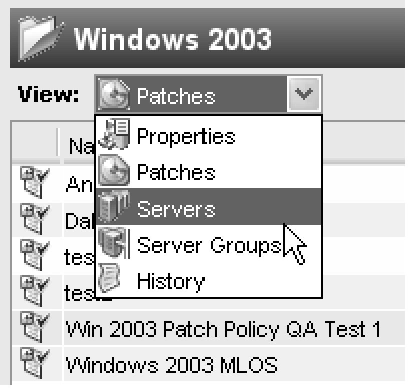

- In column headings of the Content pane, you can sort data about a feature. For example, for a managed servers, you can sort by Name, Patch Compliance, IP address, OS, and so on.
- With the search tool, you can search the Content pane by feature attribute, as shown in Figure 3-6.

Figure 3-6: Search Tool

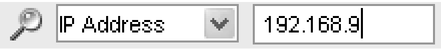

Preview Pane

The preview pane allows you to preview information about servers, server groups, patches, and patch policies selected in the Content pane without having to open a new window.

You can use the Preview pane to perform the following actions:

- To preview information about a server, server group, patch, or patch policy, select it in the Content pane.

- To filter the type of information you view in the Preview pane, from the top of the content area, choose a view from the View drop-down list.
- To deactivate the Preview pane, from the View menu, select **Preview Pane ➤ Minimize**.

For example, if you are viewing Windows 2003 patched from the Software Library, you can select a patch in the Contents pane and see information about the patch in the Preview pane. This is shown in Figure 3-7.

Figure 3-7: Main Application Windows Showing Patch Properties Information in the Preview Pane

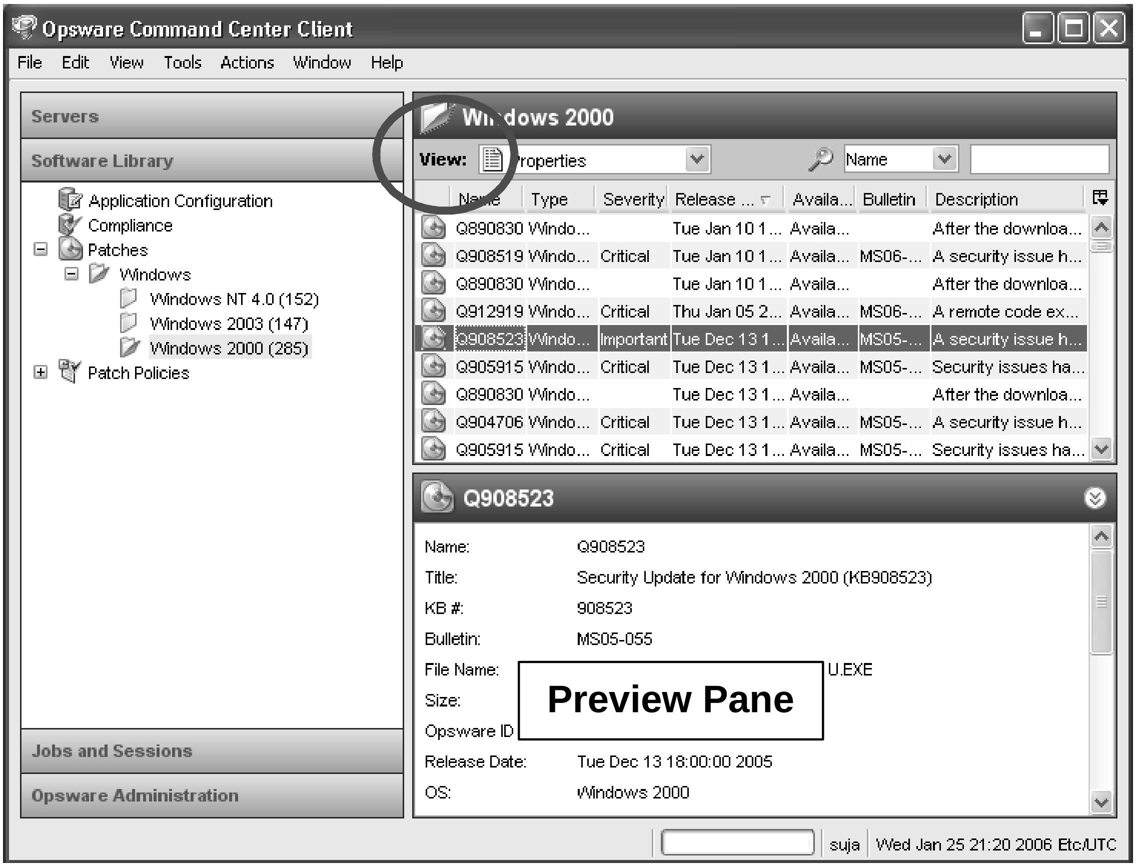

Na e	Type	Severity	Release ...	Availa...	Bulletin	Description
Q890830	Windo...		Tue Jan 10 1...	Availa...		After the downloa...
Q908519	Windo...	Critical	Tue Jan 10 1...	Availa...	MS06-...	A security issue h...
Q890830	Windo...		Tue Jan 10 1...	Availa...		After the downloa...
Q912919	Windo...	Critical	Thu Jan 05 2...	Availa...	MS06-...	A remote code ex...
Q908523	Windo...	Important	Tue Dec 13 1...	Availa...	MS05-...	A security issue h...
Q905915	Windo...	Critical	Tue Dec 13 1...	Availa...	MS05-...	Security issues ha...
Q890830	Windo...		Tue Dec 13 1...	Availa...		After the downloa...
Q904706	Windo...	Critical	Tue Dec 13 1...	Availa...	MS05-...	A security issue h...
Q905915	Windo...	Critical	Tue Dec 13 1...	Availa...	MS05-...	Security issues ha...

Name:	Q908523
Title:	Security Update for Windows 2000 (KB908523)
KB #:	908523
Bulletin:	MS05-055
File Name:	U.EXE
Size:	
Opsware ID	
Release Date:	Tue Dec 13 18:00:00 2005
OS:	Windows 2000

If you would like to view other types of information about the selected patch, from the View drop-down list, choose a view.

Preview Pane Show Filter

Some features displayed in the Preview panes allow you to further filter the feature. Using the Show drop-down list, you can choose different views of the feature.

For example, if you are viewing all of the servers that the patch policy is attached to, in the Preview pane you can select to filter either Servers with Policies Attached or Servers without Policies Attaches, as shown in Figure 3-8.

Figure 3-8: Show Drop-down List

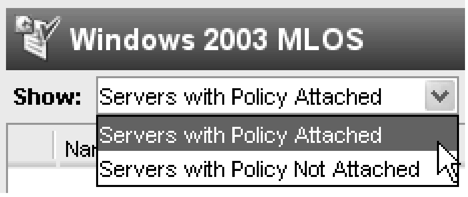

Windows 2003 MLOS
Show:
Servers with Policy Attached
Servers with Policy Attached
Servers with Policy Not Attached
Nar

Status Bar

At the bottom of the OCC Client application window, the status bar provides the following information:

- Informational text (left hand side) about the selected object
- A progress bar that shows progress on retrieving information from the core
- User ID
- Current Opsware SAS time

Figure 3-9: OCC Client Status Bar

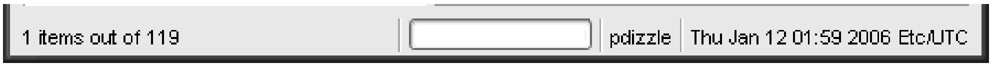

OCC Client Options

You can configure the following options for the OCC Client:

- **General Options**: This enables you to set options such as choosing the core you want to log into by default, how to handle caching, and so on.
- **Unmanaged Servers**: This enables you to set options for the Opsware Agent Discovery and Deployment (ODAD) feature.
- **Packaging Servers**: This enables you to view and specify packaging servers available on the core.
- **Terminal and Shell**: This enables you to configure your Terminal (UNIX) and RDP (Windows) client for the Global Shell and Remote Terminal connections.

- **Patch Policies**: This enables you to specify that a confirmation message will display when you try to remove a patch policy or a patch policy exception from a managed server.

Accessing OCC Client Options

To set OCC Client options, perform the following steps:

1. From the **Tools** menu, select **Options**.
2. From the left side of the Set Options dialog box, choose an option.

General Options

The following general options enable you to select your default core:

Core Server Defaults

This option allows you to configure the core that you log into from the OCC Client. Options include:

- **Host**: The name or IP address of the Opsware Command Center host that you log into by default.
- **Port**: The port number of the Opsware Command Center host that you log into by default; 443 is the default.

Cache

This option enables you to configure the caching of data displayed inside the OCC Client. You can configure the following cache settings:

- **Check for updates every X minute(s)**: Enter a value for how many minutes will pass before the cache is refreshed.
- **Check Now**: Click to check instantly for new information from the core.
- **Reload the Cache**: Click to immediately reload (refresh) the cache.

Progress Information

This option shows the progress of a job. When a job finishes, the progress window closes.

Unmanaged Servers

These options allow you to control the operation of the ODAD. You can set the following ODAD options:

- Installer Options
- Protocols
- Advanced

See Chapter 6, “Opsware Agent Management” on page 181 of this guide for more information.

Installer Options

From the Installer Options window, you can set the Installer options and control the installation of an Opsware Agent on a server. The Installer Options window enables the following actions:

- **Start the Opsware Agent after Installation**: Starts the Opsware Agent after installing it on the server. By default, the Opsware Agent Installer does not start the Opsware Agent.
- **Remove existing machine IDs and cryptographic material**: Removes any existing machine specific identifying materials, such as Machine ID file (MID), and all machine-specific cryptographic material from the server.
- **Ignore prerequisite check failures**: Ignores prerequisite check failures and forces Opsware Agent installation.
- **Set the server’s time from the Opsware core**: Synchronizes the time on the server in which the Opsware Agent is installed with the Opsware core.
- **Install Windows Installer (MSI) if required**: Installs MSI along with the Opsware Agent. If MSI is already installed, this option has no effect.
- **Install Windows Management Instrumentation (WMI) if required**: Installs WMI along with the Opsware Agent. If WMI is already installed, this option has no effect.
- **Reboot Windows servers after agent installation**: Reboots Windows servers after Opsware Agent installation is complete.
- **Install Red Hat Package Manager (RPM) on AIX and Solaris**: Installs the RPM handler with the Opsware Agent. Opsware recommends that you always include this option when you install Opsware Agents on Solaris and AIX servers.
- **Reset Opsware Agent configuration, if present**: Replaces the existing Opsware Agent configuration.
- **Delete gateway address list, if present**: Deletes the Opsware Gateway address list, if present and is no longer required.

- **Overwrite staged Opsware Agent installer**: Overwrites the existing Opsware Agent Installer.
- **Log Level**: Sets the log level for log messages. With this option, you can specify levels for error, warning, info, and trace.

Protocols

The Protocol window allows you to specify the standard port to connect to the servers for deployment. Protocols used are SSH, Rlogin, Telnet, Netbios, and WTS.

- **SSH**: The standard port to connect to the servers for deployment using the SSH protocol.
- **Rlogin**: The standard port to connect to the servers for deployment using the Rlogin protocol.
- **Telnet**: The standard port to connect to the servers for deployment using the Telnet protocol.
- **NetBIOS**: The standard port to connect to the servers for deployment using the NetBios protocol.
- **WTS**: The standard port to connect to the servers for deployment using the WTS protocol.

Advanced

The Advanced Installer Options window allows you to set the following:

- **Immediately do a full hardware registration**: This option forces the Opsware Agent Installer to report full hardware information to the core.
- **Immediately do software registration**: This option forces the Opsware Agent Installer to report full software information to the core.
- **Suppress Opsware Agent reachability check**: This option disables this check during installation. By default, the installer triggers the core to check if the server is reachable.
- **Disallow anonymous SSL connections if Opsware Agent is dormant**: This option configures the Opsware Agent so that browsers cannot connect without a valid certificate.
- **Force creation of new device record if conflict found**: This option suppresses this functionality. During registration, the Data Access Engine creates a new device record.

- **Fail if initial hardware registration fails (do not go dormant)**: This option does not allow the Opsware Agent to become dormant, if it fails to report hardware information.
- **Reconcile Type**: This option reconciles the server against any nodes assigned to the server. The reconcile type can be Full or Add only.
- **Attach to template ID**: This option assigns the nodes contained in the template to the server.
- **Extra Installer options**: This option allows you to specify any other installer options.
- **nmap parameters**: This option allows you to specify parameters used when scanning for unmanaged servers. If you find that the Opsware Discovery and Agent Deployment is unable to correctly locate and identify unmanaged servers due to the network firewall configuration, you can specify a different set of scan parameters. See the nmap documentation for more information.

Packaging Servers

In order to create packages with the OCC Client, you need to configure a managed server in your Opsware core as a packaging server.

The Set Options window lists all configured packaging servers on the core. To sort the list alphabetically by operating system, click the OS column heading. To sort this list of servers by existing configured packaging servers, click the Default Server column heading. For servers that are already configured as packaging servers, the IP address is displayed in this column. For servers that are not configured as packaging servers, Not Configured is displayed in this column.

To configure a packaging server, perform the following steps:

1. In the right pane, select a managed server and click **Edit**.
2. In the Select Server window, select the managed server that you want to configure as the packaging server and then click **Select**. In the Set Options window, the Default Server column will now display the IP address of the server. (Before you configured this as the packaging server, the Default Server column displayed Not Configured.)
3. In the Set Options window, click **Save** to save your changes or click **Cancel** to close this window without saving your changes.

Terminal and Shell

These settings define the command that the OCC Client invokes on your PC to open a Global Shell or Remote Terminal session. (For instructions on using an `ssh` client instead of the OCC Client, see "Opening a Global Shell Session" on page 387.)

- **Terminal Client**: Specifies the terminal client that the OCC Client uses for Remote Terminal sessions on Unix managed servers and for Global Shell sessions. The default value is:

```
cmd /c start /w cmd /c "telnet %h %p && echo %m && pause > nul"
```

  The `telnet` program emulates a command-line terminal session. The `%h` represents the host and the `%p` is for the port. (See Table 3-1.)

  If you change the Terminal Client setting from the default value, make sure that the command blocks until the terminal application terminates. The terminal application must not run in the background. If you specify `cmd /c start`, include the `/w` switch to make `cmd` block until the underlying command (such as `telnet`) completes.

  You are not required to use telnet as the terminal application. For example, to use a PuTTY client, specify the following command:

  ```
  "C:\\Program Files\\putty\\putty.exe" -telnet %h %p
  ```

- **RDP Client**: Specifies the remote desktop protocol (RDP) client that the OCC Client uses for Remote Terminal sessions on Windows managed servers. The default value is the Microsoft Terminal Services Client:

  ```
  mstsc "%r"
  ```

The OCC Client supports the Windows XP version of the Remote Desktop Connection Software, which can be downloaded from the following URL:
http://www.microsoft.com/windowsxp/downloads/tools/rdclientdl.mspx

The specified terminal client must be installed on your PC. To verify the existence of the terminal client, click **Test**.

The command can include variables such as `%h` and `%p`. When the terminal client is launched, these variables are replaced with the values shown in Table 3-1. To override a replacement value, specify a constant instead of a variable. For example, you might specify 435 for the port instead of `%p`.

Table 3-1: Variables for the Terminal and RDP Client Options

VARIABLE	DESCRIPTION	REPLACEMENT VALUE
`%e`	The character encoding.	For Remote Terminal sessions, the encoding of the managed server. For Global Shell sessions, the value of the Encoding field.
`%h`	The host name that the `telnet` client connects to.	The value of the `localhost` of the managed server.
`%m`	A locale-specific message on how to close the window.	For English locales, “Press the Enter key to close this window.”
`%p`	The port that the telnet client connects to.	A randomly chosen port.
`%r`	The name of the Remote Desktop (RDP) connection file. This variable is used only for the Microsoft Terminal Services Client (`mstsc`).	A temporary RDP file generated at runtime by the OCC Client.
`%t`	The title displayed in the terminal window.	For Remote Terminal sessions, the name of the managed server. For Global Shell sessions, the string “Global Shell.”

- **Encoding**: Sets the encoding for Global Shell and the Remote Terminal sessions. This option is the replacement value of the `%e` variable in the command specified by the Terminal Client field. The default value of the Encoding option is UTF-8.

Patch Policies

This option allows you to specify that a confirmation message will display when you try to remove a patch policy or a patch policy exception from a managed server.

Chapter 4: Server Tracking in the OCC

IN THIS CHAPTER

This section discusses the following topics:

- Server Tracking
- Server Lists
- My Servers
- Server Search
- Server Identification
- Server Histories and Reports
- Hardware Information for Managed Servers

Server Tracking

This section discusses the following topics:

- Ways to Locate, List, and Display Servers
- Tracked Server Properties
- Supported Operating Systems for Managed Servers.

Ways to Locate, List, and Display Servers

You can locate, list, and display servers in the Opsware Command Center in the following four ways:

- By searching when you know the name, host name, or IP address of the server you want to provision or manage.
- By viewing the Manage Servers list and Server Pool list when you want to see a complete list of all your servers. You can refine the lists by using filters.

- By browsing nodes in the Software Tree to determine which servers specific software should be installed on or find all servers that match certain rules; for example all servers with a specific application installed.
- By viewing servers sorted by hardware category. (Click Environment ➤ Hardware in the navigation panel.) The Servers tab in the Hardware pages shows each manufacturer and model that you have running in the operational environment. See "Hardware Information for Managed Servers" on page 155 in this chapter for more information.
- You can also browse managed servers and server groups using the OCC Client.

Tracked Server Properties

Every server that Opsware SAS manages has the following properties:

- IP addresses, host name, and the server ID
- Which nodes the server is assigned to in the Software Tree categories. Click a node link to display specific information about that node.
- What software should be installed on the server. Select the Install List tab from the Manage Servers: Server Properties page to display the list of software packages that should be installed on that server by virtue of that server's assigned nodes.

  Opsware SAS is able to determine what software should be installed on a server because of its model-based approach to server management. The software that should be installed is recorded in the Opsware Model Repository.
- All software that is installed on the server. Select the Installed Packages tab from the Manage Servers: Server Properties page to display the list of software that is reportedly installed on the server.

  Opsware SAS is able to determine what software is installed on a server because the Opsware Agent communicates with the Opsware core and reports the installed hardware and software for the server.

  In some cases, Solaris packages might only be partially installed. In these cases, the partially installed Solaris package does not show up in the installed list.

See "Server Information that the Opsware Agent Tracks" on page 185 in Chapter 6 for information about a complete list of all the hardware and software information that Opsware SAS tracks for managed servers.

Supported Operating Systems for Managed Servers

This section lists the supported operating systems for Opsware Agents, the Opsware Command Center, and the OCC Client.

The following table lists the supported operating systems for Opsware Agents, which run on the servers managed by Opsware SAS.

For the supported operating systems for Opsware Agents, Opsware SAS supports Red Hat Linux 3 AS/ES/WS and Red Hat Linux 4 AS/ES/WS on both 32 bit (also known as i386) and 64 bit (also known as AMD64 or EM64) `x86` architecture. All other versions of Red Hat Linux are supported on 32 bit architecture only.

Table 4-1: Opsware Agent Supported Operating Systems

SUPPORTED OPERATING SYSTEMS FOR OPSWARE AGENT	VERSIONS	ARCHITECTURE
AIX	AIX 4.3 AIX 5.1 AIX 5.2 AIX 5.3	POWER POWER POWER POWER
HP-UX	HP-UX 10.20 HP-UX 11.00 HP-UX 11.11 HP-UX 11i v2	PA-RISC PA-RISC PA-RISC PA-RISC and Itanium
Sun Solaris	Solaris 6 Solaris 7 Solaris 8 Solaris 9 Solaris 10	Sun SPARC Sun SPARC Sun SPARC Sun SPARC Sun SPARC, 64 bit x86 and Niagara
Fujitsu Solaris	Solaris 8 Solaris 9 Solaris 10	Fujitsu SPARC Fujitsu SPARC Fujitsu SPARC

Table 4-1: Opsware Agent Supported Operating Systems

SUPPORTED OPERATING SYSTEMS FOR OPSWARE AGENT	VERSIONS	ARCHITECTURE
Windows	Windows NT 4.0 Windows 2000 Server Family Windows Server 2003	32 bit x86 32 bit x86 32 bit x86
Red Hat Linux	Red Hat Linux 6.2 Red Hat Linux 7.1 Red Hat Linux 7.2 Red Hat Linux 7.3 Red Hat Linux 8.0 Red Hat Enterprise Linux 2.1 AS Red Hat Enterprise Linux 2.1 ES Red Hat Enterprise Linux 2.1 WS Red Hat Enterprise Linux 3 AS Red Hat Enterprise Linux 3 ES Red Hat Enterprise Linux 3 WS Red Hat Enterprise Linux 4 AS Red Hat Enterprise Linux 4 ES Red Hat Enterprise Linux 4WS	32 bit x86 32 bit x86 32 bit x86 32 bit x86 32 bit x86 32 bit x86 32 bit x86 32 bit x86 32 bit x86 and 64 bit x86 32 bit x86 and 64 bit x86 32 bit x86 and 64 bit x86 32 bit x86 and 64 bit x86 32 bit x86 and 64 bit x86 32 bit x86 and 64 bit x86
SUSE Linux	SUSE Linux Enterprise Server 8 SUSE Linux Standard Server 8 SUSE Linux Enterprise Server 9	32 bit x86 32 bit x86 32 bit x86 and 64 bit x86

The following table lists the operating systems supported for the OCC Client.

Table 4-2: OCC Client Supported Operating Systems

SUPPORTED OPERATING SYSTEMS FOR OCC CLIENT	VERSIONS	ARCHITECTURE
Windows	Windows XP Windows 2000 Windows 2003	32 bit x86 32 bit x86 32 bit x86

Java J2SE v 1.4.2 JRE must be installed on the system that runs on the OCC Clinet. To download this version of Java, go to http://java.sun.com/j2se/1.4.2/download.html

Server Lists

This section discusses the following topics:

- Types of Server Lists
- Server Pool
- Manage Servers List
- Filters on the Manage Servers List

Types of Server Lists

The Opsware Command Center displays lists for two types of servers, as Figure 4-1 shows.

Figure 4-1: Servers Section in the Navigation Panel

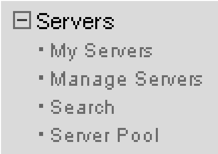

Server Pool: Servers in the Server Pool have registered their presence with Opsware SAS but do not have the target OS installed. An OS Build Agent is running on each server so that they can communicate with Opsware SAS.

See “Operating System Provisioning” on page 567 in Chapter 15 for information about how to use the Server Pool when you install the target OS on a server.

Manage Servers: The Manage Servers list contains servers on which Opsware SAS can perform management tasks, because Opsware Agents are installed on them. However, Opsware SAS might not have provisioned all software running on the servers.

You begin the OS provisioning process by reviewing the servers in the Server Pool list. From the Server Pool, you can install a target OS by selecting a server and clicking **Install OS**.

Server Pool

The Server Pool provides the following information about each server waiting to be provisioned with the target OS:

- The host name set by booting the server for the first time over the network or by using an Opsware Boot Floppy
- The MAC address
- The manufacturer and model
- The OS that the OS Build Agent is running – DOS (Windows operating systems), Linux, or Solaris

  You use this information to select the target OS for servers. If the server is in the process of installing an OS, this value might change.
- The last date and time that the Opsware Agent on the server communicated with Opsware SAS (by submitting the server's hardware and software information)

  If the server is in an unreachable state (that is, if the server icon has a red "x" on it), you can run a Communication Test to help you troubleshoot why that server is unreachable. See "Agent Reachability Communication Tests" on page 204 in Chapter 6 for more information.
- The life cycle value, such as whether the server is available to have a target OS installed on it
- The facility in which the server is located
- The customer association
- Additional hardware information (Click the server name to open a window that displays specific hardware information.)

Manage Servers List

The Manage Servers list contains servers on which Opsware SAS can perform management tasks because Opsware Agents are installed on them. When an existing operational server has an Opsware Agent installed successfully, it appears in the Manage Servers list and the server icon indicates that it is fully manageable, as Figure 4-2 shows.

Figure 4-2: The Manage Servers List in the Opsware Command Center

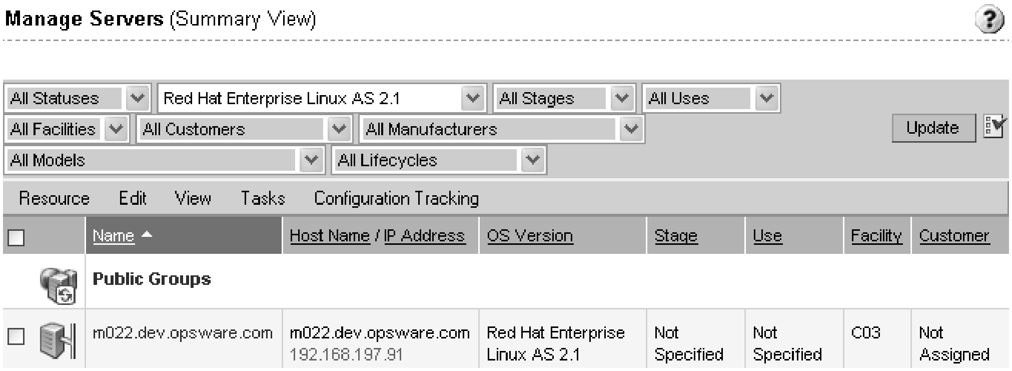

☐	Name ▲	Host Name / IP Address	OS Version	Stage	Use	Facility	Customer
	Public Groups						
☐	m022.dev.opsware.com	m022.dev.opsware.com 192.168.197.91	Red Hat Enterprise Linux AS 2.1	Not Specified	Not Specified	C03	Not Assigned

See “Opsware Agent on Managed Servers” on page 181 in Chapter 6 for more information. By default, servers in the Manage Servers list are sorted by the Name column. However, you can re-sort the list based on any of the column headings. For example, you can click the Hostname / IP Address column heading to re-sort the list by host name or IP address.

If you have many servers that Opsware SAS manages, the list of servers is grouped by pages. Click the page number links or the left arrow at the bottom of the list.

The Manage Servers list provides the following information about each server:

- The name of the server

  By default, the server’s host name appears in this field. However, you can edit it so that it is more descriptive or useful.

- The host name of the server determined by the Opsware Agent
- The IP address configured for the server, which users can edit by using the network configuration feature in the Opsware Command Center
- The reported OS, which is obtained by the Opsware Agent that is running on the server
- The stage of the server, which specifies the stages of deployment for servers
- The server’s use

- The facility in which the server is located
- The customer association
- Additional hardware information (Click the server name to open a window that displays specific hardware information.)

Filters on the Manage Servers List

The Manage Servers list displays the following filters that you can specify to qualify the servers that the Opsware Command Center displays, as Figure 4-3 shows.

Figure 4-3: Filters in the Manage Servers List

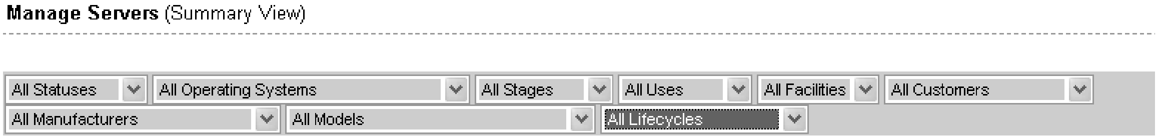

- **Status**: Specifies the ability of Opsware SAS to manage servers. Opsware SAS automatically detects the status of servers; a server's status is OK or Not Reachable.
- **OS**: Specifies the operating system on the server, which is obtained by the Opsware Agent that is running on the server.
- **Stages**: Specifies the stages of deployment for servers; for example, a server is live or offline. Users set this property for servers. The values in this list are customizable by the Opsware administrator.
- **Uses**: Specifies how an organization is utilizing servers; for example, a server is a staging server. Users set this property for servers. The Opsware administrator can customize the values in this list.
- **Facilities**: Specifies the location of servers. From an Opsware Command Center, users can manage servers located in any facility. For example, a user could log in to the Opsware Command Center running in facility A and manage the server located in facility B.
- **Customers**: Specifies the customer associated with each server. Your Opsware administrator defines the options for customer selections by using the Administration pages.
- **Manufacturers**: Specifies the manufacturer for the server as reported by the OS Build Agent running on the server.

- **Models**: Specifies the model of the server as reported by the OS Build Agent running on the server.
- **Life cycles**: Specifies the various Opsware server life cycle values which include Managed, Available, Build Failed, Installing OS, and Deactivated.

You can change the filters displayed in the Manage Servers page. To change the filters you want to be displayed on the Manage Servers page, click on the icon as shown in Figure 4-4 and specify the filters from the **Edit Filters** Menu.

Figure 4-4: Edit Icon

Figure 4-5 shows the filters that are in the Server Pool list.

Figure 4-5: Filters in the Server Pool List

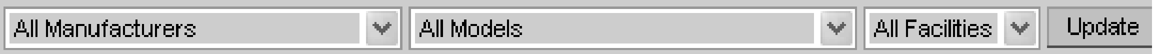

- **Manufacturers**: Specifies the manufacturer for the server as reported by the OS Build Agent running on the server.
- **Models**: Specifies the model of the server as reported by the OS Build Agent running on the server.
- **Facilities**: Specifies the location of the server. Users can manage servers in any facility from an Opsware Command Center in any facility.

My Servers

This section contains the following topics:

- Overview of My Servers
- Adding Servers to My Servers
- Removing Servers from My Servers

Overview of My Servers

The My Servers feature provides an efficient way to manage servers when your operational environment contains hundreds or thousands of servers.

When you search for servers or browse the server lists, you can add servers to My Servers (similar to a shopping cart on an e-commerce site). Using My Servers allows you to view and perform actions on selected servers.

When you use the same browser and login to the Opsware Command Center running in the same facility, servers stay in My Servers for one year or until you explicitly remove them. Each time that you login to the Opsware Command Center, you see the servers that were in My Servers the last time that you logged in.

The My Servers feature is available only on a per-user basis. You cannot log in as an Opsware administrator to see the servers in the My Servers area of other Opsware users.

Adding Servers to My Servers

Perform the following steps to add a server to My Servers:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the servers that you want to add to the My Servers.

   Or

   Search for the servers that you want to add to My Servers.

   See "Using the Search Feature" on page 125 in this chapter for more information. See "Server Searching by IP Address" on page 141 in this chapter for more information.

2. Select the servers that you want to add to My Servers.

3. Choose **Resource ➤ Add to My Servers** from the menu above the Manage Servers list. The Add To My Servers window appears, which indicates that you added the chosen number of servers to My Servers.
4. Click **Close** to close the window.
5. Next, select the My Servers link at the top of the page. You see the selected servers added to My Servers, as Figure 4-6 shows.

Figure 4-6: Servers in My Servers

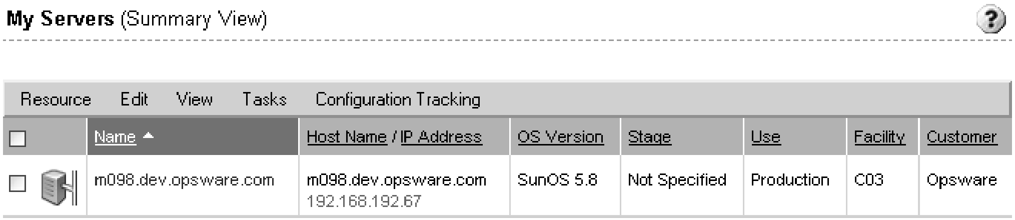

My Servers (Summary View)

□	Name ▲	Host Name / IP Address	OS Version	Stage	Use	Facility	Customer
□	m098.dev.opsware.com	m098.dev.opsware.com 192.168.192.67	SunOS 5.8	Not Specified	Production	C03	Opsware

You can perform the same server management tasks on servers in My Servers as on the servers in the Manage Servers list.

Removing Servers from My Servers

Perform the following steps to remove servers from My Servers:

1. Click the My Servers link in the navigation panel of the Opsware Command Center. The My Servers page appears that shows the servers currently added to it.
2. Select the servers that you want to remove from My Servers and choose **Edit ➤ Remove from My Servers** from the menu above the Server list.

The My Servers page refreshes and displays the remaining servers in My Servers.

Server Search

This section provides information about Server Search and contains the following topics:

- Searching for a Server By Using the Search Box
- Ways to Use Search
- Using the Search Feature
- Rules for Server Search and for Creating Dynamic Groups
- Line Break Workaround for Server Search
- Conditions for Searching with Multiple Rules
- Server Searching by IP Address
- Example: Server Search
- Searching for a Server Group

Searching for a Server By Using the Search Box

Perform the following steps to search for a server using the Search box:

1. On the Home page, click the down arrow in the top navigation panel to open the Search box, as Figure 4-7 shows.

Figure 4-7: Search Text Box on the Opsware Command Center Home Page

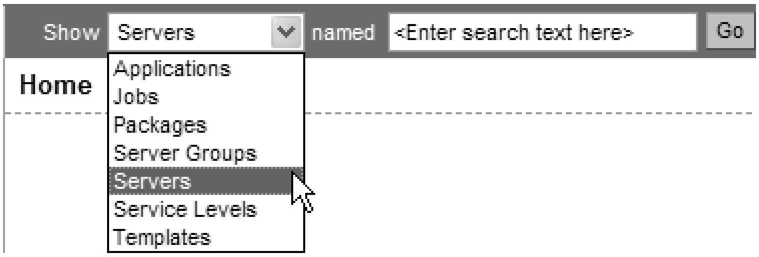

2. Verify that the Servers option is selected in the list.

3. Type the server's IP address, host name, or name in the Search box and then click **Go**.

   The search text that you enter can include an asterisk (*) wildcard character. However, the search feature automatically prepends and appends an asterisk to the text.

   For example, you can type any of the following search queries:

```
192.168.68.6
host02.coredev-va1.sample.com
192.168.*.19
host1*.xyz.samplecompany.com
```

The resulting page contains one or more servers, depending on the type of search query that you specified. If no servers are found, the Opsware Command Center displays a message that indicates that no servers were found that matched your query.

See “Using the Search Feature” on page 125 in this chapter for information about how to formulate complex, multiple rule search queries.

Ways to Use Search

By using the Opsware Command Center, you can perform searches in the following ways:

- From Opsware wizards

  While using the Opsware wizards from the Tasks panel, you are prompted to select (by browsing or searching) servers or server groups, operating systems, patches, applications, and templates at specific points in the process.

  What you can search for in the Opsware wizards is context sensitive to the type of operation that you are performing. For example, if you are using the Install Patch Wizard, you can select the Search tab to search for patches to install on servers.

- When adding operating systems, patches, or applications to templates

  Searching for an operating system, patch, or application to add to a template functions the same way as searching through the Opsware wizards.

- From the navigation panel, click Servers ➤ Search.

  The Search page allows you to search for managed servers that match specified rules.

- While adding or modifying rules for a dynamic group (By clicking **Search** on the Rules tab for a group).

Using the Search Feature

In the Opsware wizards, you can browse for servers, operating systems, patches, applications, templates, and service levels, or use the Search feature to search for these items.

Perform the following steps to search by using the Search feature:

1. In an Opsware wizard, select the Search tab. The following Search page appears. See Figure 4-8.

Figure 4-8: Search Tab in the Select Servers Step of an Opsware Wizard

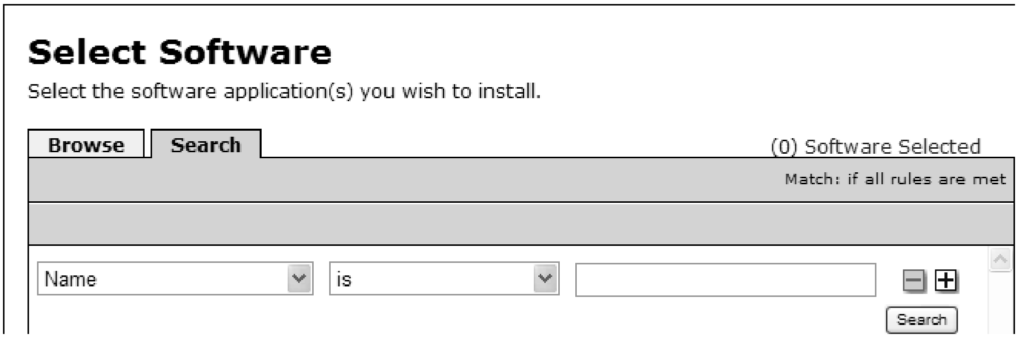

Select Software

Select the software application(s) you wish to install.

You can also use the Search tab at other steps in the wizards to search for operating systems, patches, applications, and templates.

Or

From the navigation panel, click Servers ➤ Search. The Search page appears, as Figure 4-9 shows.

Figure 4-9: Search Page

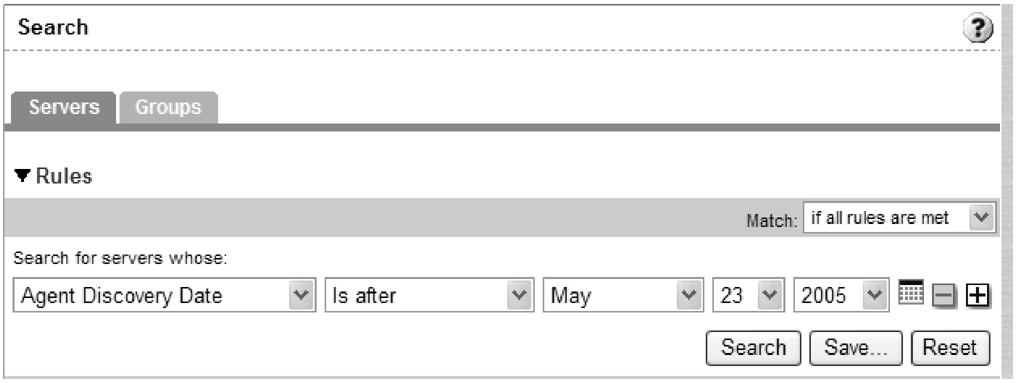

By default, one search rule is added to the search.

2 Specify the rule that you want to search for by selecting it from the first list, as Figure 4-10 shows.

Figure 4-10: Search Rules List in the Search Page

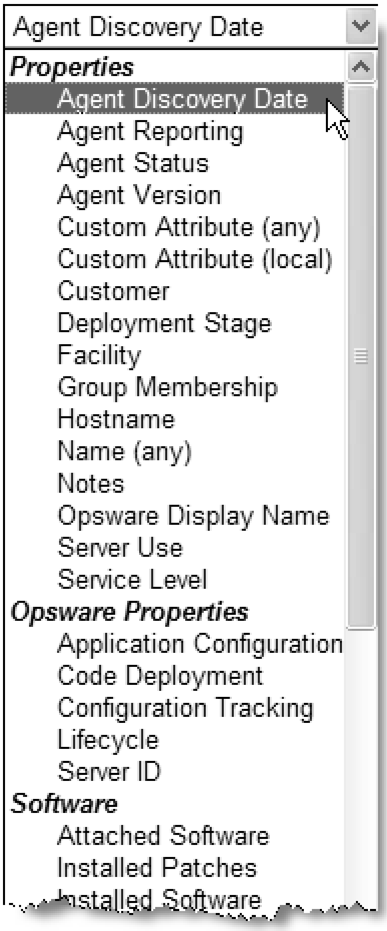

Agent Discovery Date

Properties

- Agent Discovery Date
- Agent Reporting
- Agent Status
- Agent Version
- Custom Attribute (any)
- Custom Attribute (local)
- Customer
- Deployment Stage
- Facility
- Group Membership
- Hostname
- Name (any)
- Notes
- Opsware Display Name
- Server Use
- Service Level

Opsware Properties

- Application Configuration
- Code Deployment
- Configuration Tracking
- Lifecycle
- Server ID

Software

- Attached Software
- Installed Patches
- Installed Software

Depending on the rule that you select, a popup window might appear in the page. For example, if you selected Deployment Stage, a popup window showing the stages appears in the page as Figure 4-11 shows:

Figure 4-11: Search Popup Window with Values

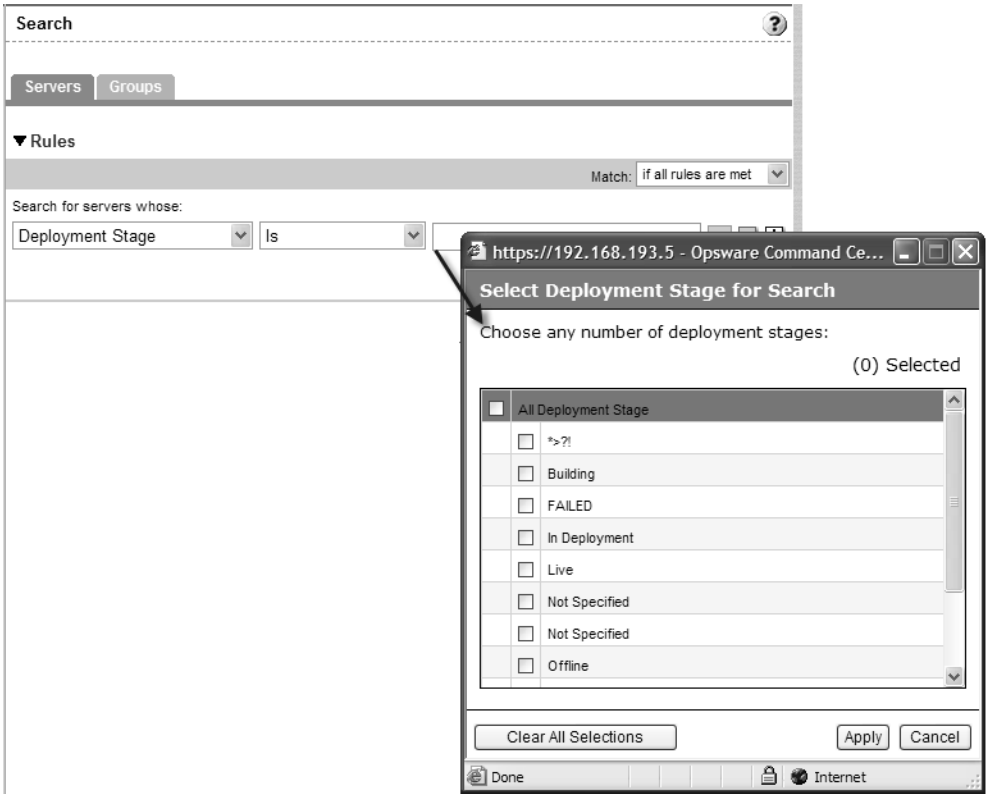

Match: if all rules are met

Search for servers whose:

Deployment Stage | Is

https://192.168.193.5 - Opsware Command Ce...

Select Deployment Stage for Search

Choose any number of deployment stages:

(0) Selected

- □ All Deployment Stage
- □ *>?!
- □ Building
- □ FAILED
- □ In Deployment
- □ Live
- □ Not Specified
- □ Not Specified
- □ Offline

You cannot search in Notes that contain line breaks. See "Line Break Workaround for Server Search" on page 139 in this chapter for more information and a workaround.

If you are searching while using an Opsware wizard, the first search rule list might not have all the options. The list only includes the options that are relevant for the Opsware wizard that is being used. For example, the Install OS Wizard does not include options to search for installed patches on the servers.

3 In the second list, specify how you want Opsware SAS to search by selecting a value. The operator selected defines how the search text is treated. Negative operators might not be available in all cases.

See "Rules for Server Search and for Creating Dynamic Groups" on page 130 in this chapter for more information about the operator for each search rule.

4 Enter the text that you want to search for in the text box or choose a value from the list or popup window. The search text that you enter can include an asterisk (*) wildcard character. The search text is case insensitive. You can also use the SHIFT or CTRL key to select multiple values from the list or popup window.

5 (Optional) To add additional rules, click the plus (+) button as Figure 4-12 shows and repeat Steps 2 through 4.

Figure 4-12: Multiple Rules in a Search

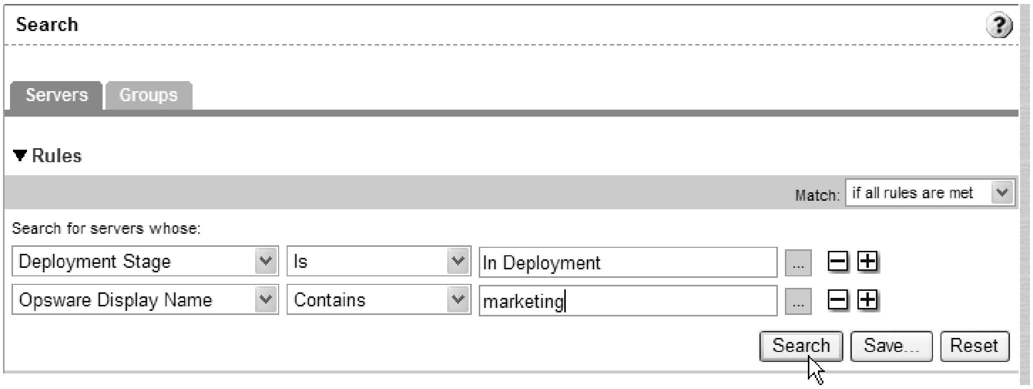

6 If you specified multiple rules for the search, select whether you want search results only if all rules are met or if any of the rules are met, as Figure 4-13 shows.

Figure 4-13: Operator Controlling Search Results

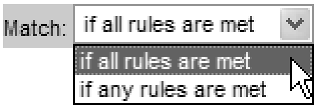

By default, search results appear for servers that match all the search rules. If you are searching from an Opsware wizard, this field is set to the value if all rules are met; you can change the value when searching for servers, but you cannot change it when searching for patches, software, operating systems, and so on.

7 Click **Search**. The list of servers that match your search rules appears in the page, as Figure 4-14 shows.

Figure 4-14: Displayed Search Results

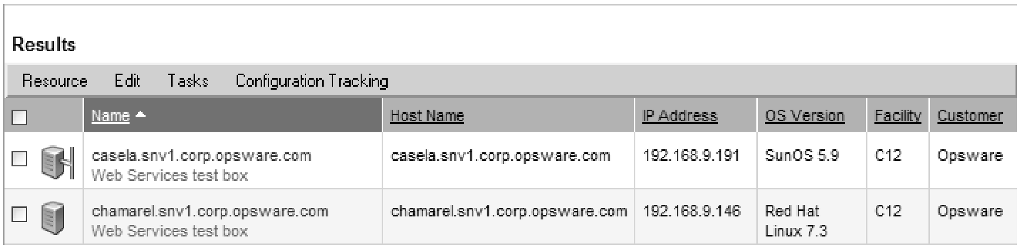

Results

□	Name ▲	Host Name	IP Address	OS Version	Facility	Customer
□	casela.snv1.corp.opsware.com Web Services test box	casela.snv1.corp.opsware.com	192.168.9.191	SunOS 5.9	C12	Opsware
□	chamarel.snv1.corp.opsware.com Web Services test box	chamarel.snv1.corp.opsware.com	192.168.9.146	Red Hat Linux 7.3	C12	Opsware

The search results include columns for Name, IP Address, OS Version, Facility, and Customer.

When you search for installed software or patches and include an asterisk in the search text, Opsware SAS might take several minutes to display the search results.

Rules for Server Search and for Creating Dynamic Groups

The following table describes the rules that you can use to search for servers or to create dynamic server groups.

Note that anywhere you can enter text, you can enter a wildcard (*) character to broaden your results.

Table 4-3: Rules for Server Search and for Creating Dynamic Groups

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
PROPERTIES		
Agent Discovery Date: The date that the Opsware Agent was installed.	• Is after • Is before	Drop-down lists with the day, month, and year
	• Is within the last	User-entered text
	• Is today	N/A

Table 4-3: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Agent Reporting: Whether the Opsware Agent is reporting to Opsware SAS.	• Is • Is not	• Has not reported • OK • Registration in progress • Reporting error
Agent Status: Whether the Opsware Agent is reachable by Opsware SAS.	• Is • Is not	• Not reachable • OK
Agent Version: The version of the Opsware Agent – such as 14.2.3b.	• Contains • Does not contain • Is • Is not	User-entered text
Custom Attribute (any): The name of a custom attribute that is associated with the server through attachment or inheritance.	• Contains • Is	User-entered text
Custom Attribute (local): The name of a custom attribute that is locally attached to the server,	• Contains • Is	User-entered text
Customer: The customer or account that the server is associated with.	• Is • Is not	Popup window of customers

Table 4-3: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Deployment Stage: The stage the server performs within a lifecycle environment.	• Is • Is not	• In Deployment • Live • Not Specified • Offline The values that appear in this list are customizable; in addition to the values above, values specific to your environment might appear.
Facility: The collection of servers managed by an Opsware SAS installation.	• Is • Is not	Popup window of facilities When Opsware SAS is running multimaster mode, the list can contain many facilities.
Group Membership: Whether the server belongs to a group.	• Is	Popup window of groups
Host name: The host name of the server – such as m004.company.com.	• Contains • Does not contain • Is • Is not	User-entered text
Name (any): This enables searching for any name or IP address associated with a server.	• Contains • Does not contain • Is • Is not	User-entered text

Table 4-3: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Notes: The contents of the Notes field from the Properties tab for a server.	• Contains • Does not contain • Is • Is not	User-entered text
Opsware Display Name: The user-configurable name for the server.	• Contains • Does not contain • Is • Is not	User-entered text By default, Opsware SAS uses the configured host name of the server until a user edits it.
Server Use: How the server is being used – such as Development, Staging, Production.	• Is • Is not	• Development • Not Specified • Production • Staging You can customize values that appear in this list. In addition to the values above, values specific to your environment might appear.

Table 4-3: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Service Level: The user-defined category that can be used as an organizational tool.	• Is • Is attached here or below	Popup window of service levels Servers can be associated with multiple service levels. See the *Opsware® SAS User's Guide* for more information about working with service levels.
OPSWARE PROPERTIES		
Application Configuration: Whether the server uses the Application Configuration feature.	• Is not used • Is used	N/A
Code Deployment: Whether the server uses the Code Deployment feature.	• Is not used • Is used	N/A
Configuration Tracking: Whether the Configuration Tracking feature is monitoring or backing up specific files or configurations on a server.	• Is off • Is on	N/A
Lifecycle: The server states that are part of bringing a server into Opsware SAS.	• Is • Is not	• Available • Build Failed • Deactivated • Installing OS • Managed

Table 4-3: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Server ID: The internal ID Opsware SAS uses to identify the server.	• Is • Is not	User-entered text In most cases, the Server ID is the same as the MID.
SOFTWARE		
Attached Software: The software that is assigned or modeled through the Opsware reconcile operation – the installation process.	• Is • Is attached here or below	Popup window of software
Installed Patches: Whether a patch is installed on the server.	• Contains • Does not contain • Is • Is not	User-entered text
Installed Software: The package reported installed on the server.	• Contains • Does not contain • Is • Is not	User-entered text A package does not have to be installed by Opsware SAS to be reported as installed on a server.
OS Version: The OS version defined by OS definitions in the OS Provisioning feature.	• Is • Is not	Popup window of operating systems
Reported OS: For Windows – the version reported by the OS, for Unix – the version returned by the uname command.	• Contains • Does not contain • Is • Is not	User-entered text

Table 4-3: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Windows Service: The names of the Windows services that are reported to Opsware SAS.	• Contains • Does not contain • Is • Is not	User-entered text
NETWORK		
DNS Search Domains: The domains configured to be searched in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
DNS Servers: The IP addresses of the DNS servers configured in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
Default Gateway: The IP address of the default router.	• Contains • Does not contain • Is • Is not	User-entered text
IP Address: Any Internet Protocol address for the server.	• Contains • Does not contain • Is • Is not	User-entered text
MAC Address: Any Media Access Control address, which is the network interface card's unique hardware number.	• Contains • Does not contain • Is • Is not	User-entered text

Table 4-3: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
WINS Servers: The Windows Internet Naming Servers configured in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
HARDWARE		
CPU Make and Model: The vendor name and CPU model for the server – such as GENUINEINTEL Intel(R) Pentium(R) 4 CPU 2.60GHz.	• Is • Is not	Popup window of CPU makes and models
CPU Speed: The Central Processing Unit speed in gigahertz [GHz].	• Does not equal • Equals • Is greater than • Is less than	User-entered text A 600 Mhz machine should be entered as `0.6`.
Make and Model: The vendor name and server model for the server – such as Compaq - DL360.	• Is • Is not	Popup window of server makes and models
Number of CPUs: The number of CPUs on the server.	• Does not equal • Equals • Is greater than • Is less than	User-entered text
RAM: The amount of RAM on the server in megabytes [MB].	• Does not equal • Equals • Is greater than • Is less than	User-entered text To enter 1 Gigabyte, type `1024`.

Table 4-3: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Serial Number: The serial number of the server.	• Contains • Does not contain • Is • Is not	User-entered text
Storage Make and Model: The vendor name and storage model for the server – such as WDC - WD800BB-75DKA0.	• Is • Is not	Popup window of storage makes and models
CUSTOM FIELDS		
A **Numeric** field	• Does not equal • Equals • Is greater than • Is less than	User-entered text
A **String** field	• Contains • Does not contain • Is • Is not	User-entered text
A **URI** field	• Contains • Does not contain • Is • Is not	User-entered text
A **File** field	• Contains • Does not contain • Is • Is not	User-entered text

Table 4-3: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
A **Date** field	• Is after • Is before	Drop-down lists with the day, month, and year
	• Is within the last	User-entered text
	• Is today	N/A

Line Break Workaround for Server Search

You cannot search in notes that contain line breaks. For example, you cannot search for text in a note when it is this type of text:

```
line1 <line break>
```

```
line2
```

For example, the following query does not return any results:

```
Any/all notes contain line1 <line break> line2
```

However, the following query does return all servers that have notes:

```
Any/all notes is not line1 <line break> line2
```

To work around this limitation, include an asterisk (*) where the line break occurs, as Figure 4-15 shows. For example:

```
Any/all notes is line1*line2
```

Figure 4-15: Line Break Workaround in the Server Search Feature

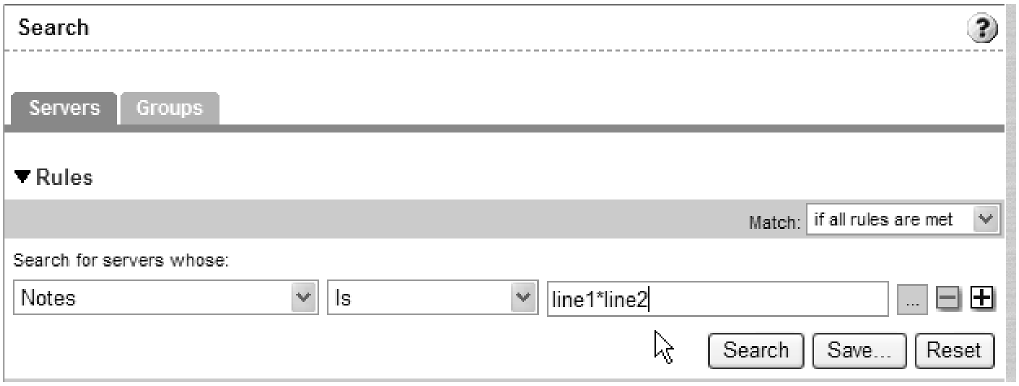

Match: if all rules are met

Search for servers whose:

Notes | Is | line1*line2

Conditions for Searching with Multiple Rules

When you perform a search with multiple rules, the following conditions apply to the way Opsware SAS provides search results:

- When it evaluates rules, the Search feature considers each rule individually, finding servers (or operating systems, patches, and so forth) that match the individual rule, and then the results of each rule are combined.
- You must select at least one rule, and it must have a value. Default filter rules count as rules with a value. For example, in the MS Updates Wizard, when searching for a server, the Search feature automatically enters a default rule to search only for servers running a Windows OS. Therefore, users are not required to specify any more rules.
- Empty rules are ignored in searches. Users do not have to manually remove them for the search to proceed.
- You cannot search for empty values. For example, you cannot search for all servers where the Notes field is empty.

At certain steps in the Opsware wizards, the Search feature provides default values based on previous selections in the wizard.

For example, in the Select Servers step of the Install Patch Wizard, the search query automatically includes the value for OS version that you specified in earlier steps, as Figure 4-16 shows.

Figure 4-16: Default Values Entered in Searches

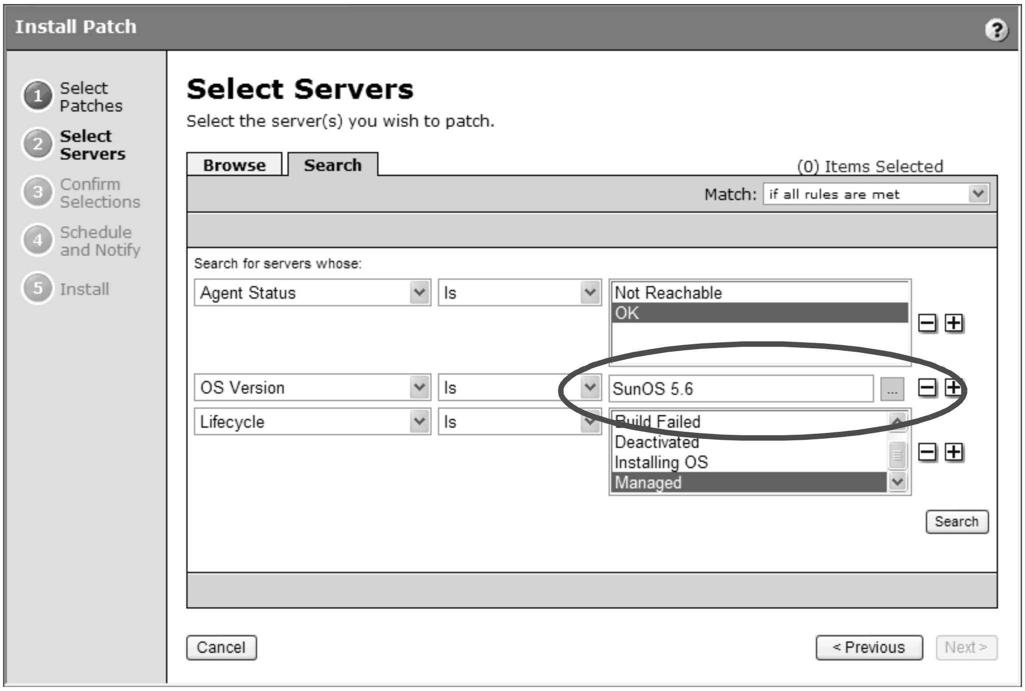

Install Patch

1 Select Patches
2 Select Servers
3 Confirm Selections
4 Schedule and Notify
5 Install

Select Servers

Select the server(s) you wish to patch.

(0) Items Selected

Match: if all rules are met

Search for servers whose:

Agent Status | Is | Not Reachable, OK
OS Version | Is | SunOS 5.6
Lifecycle | Is | Build Failed, Deactivated, Installing OS, Managed

The wizards are flexible; for example, when you select multiple servers and applications, Opsware SAS will install the correct applications on the servers, even if you selected different OS versions for the servers and applications in the respective steps. Opsware SAS also matches the customer association for applications, operating systems, and templates with the customer association of servers.

If Opsware SAS cannot find a match, an error message appears at the end of the wizard; therefore, use caution when modifying these default values.

Server Searching by IP Address

Users can search for servers by entering a specific IP address in the Search box in the top navigation or in the Search feature.

Opsware SAS includes support for static Network Address Translation (NAT). This feature introduces the concept of a management IP, which might be different from any of the local IP addresses that the Opsware Agent reported for a server.

When searching for a server based on its IP address, the Search feature searches based on the server's primary IP address and based on the IP address for any interface that server has, including its management IP address. In the search results, an extra column is shown that lists all matching IP addresses for all interfaces. The management IP address is included if the server's networking is configured for static NAT.

See "Communication Between Managed Servers and Opsware SAS" on page 221 in Chapter 7 for information about how Opsware SAS handles servers that are affected by static NAT.

Example: Server Search

Using the Search feature, a user creates a query with the following conditions:

- Installed Software contains qa
- Installed Software contains man
- If all rules are met is selected

The results of this search will be all servers that have at least one installed package with qa somewhere in its name and at least one installed package (not necessarily the same one) with man in its name.

The search results are not limited to packages that contain both qa and man in the package name.

Find all servers that have some version of Apache or some version of Java installed:

- Installed Package contains apache
- Installed Package contains java
- If any rules are met is selected

Searching for a Server Group

1. From the navigation panel, click Servers ➤ Search. The Search page appears. By default, the Servers tab is selected.

   Select the Groups tab. The rules for server group search appear as Figure 4-17 shows. By default, one search rule is added to the search.

Figure 4-17: Rules for Server Group Search

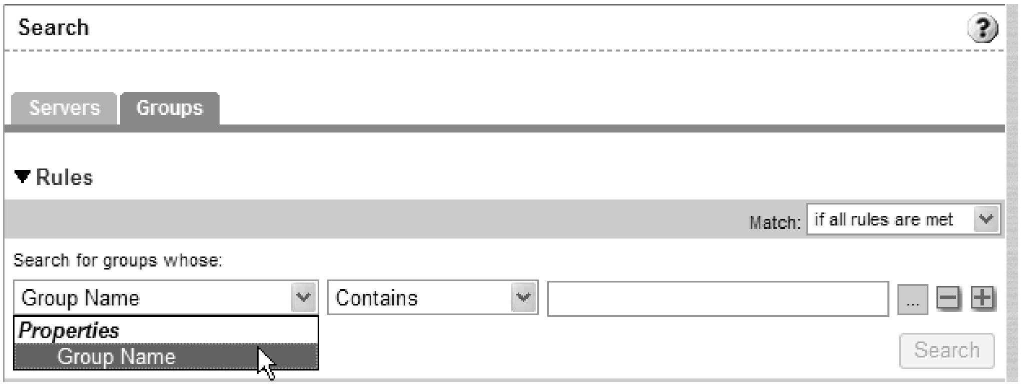

2. In the second list, specify how you want Opsware SAS to search by selecting either “Contains” or “Is”. The operator selected defines how the search text is treated.

3. Enter the text that you want to search for in the text box. The search text that you enter can include an asterisk (*) wildcard character. (The search text is case sensitive. **Search** is not enabled until you enter text in the text box.)

When searching for a server group, you can only specify one rule for Group Name. The plus (+) button is disabled.

4 Click **Search**. The search results appear as Figure 4-18 shows.

Figure 4-18: Server Group Search Results

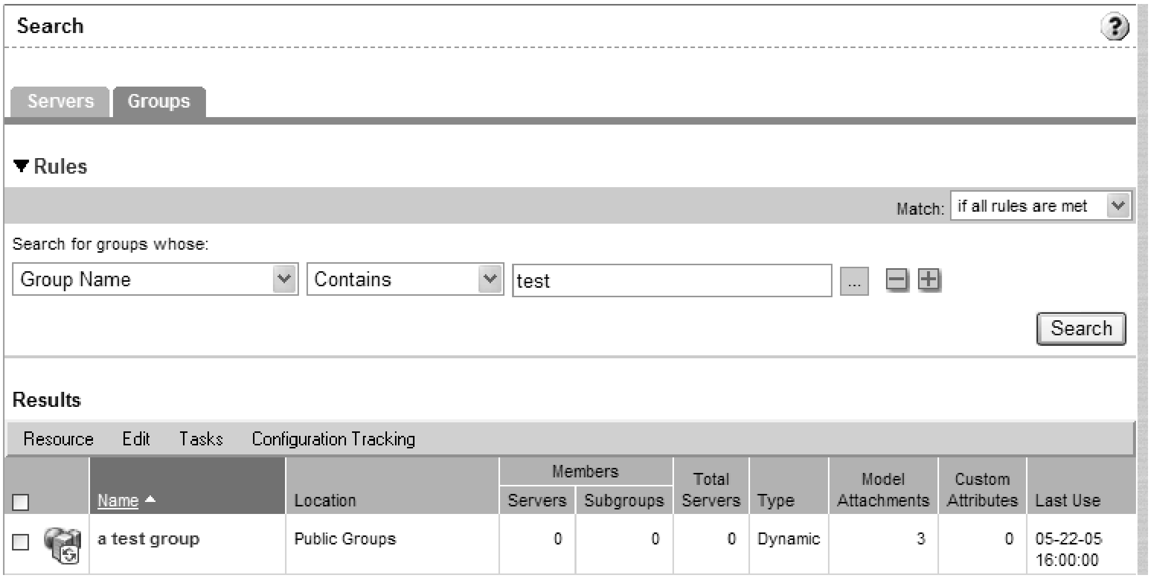

Match: if all rules are met

Search for groups whose:

Group Name | Contains | test

Results

□		Name ▲	Location	Members		Total Servers	Type	Model Attachments	Custom Attributes	Last Use
				Servers	Subgroups					
□		a test group	Public Groups	0	0	0	Dynamic	3	0	05-22-05 16:00:00

Server Identification

This section provides information on server identification within Opsware SAS and contains the following topics:

- Overview of Server Identification
- Ways that Servers are Identified by Opsware SAS
- Customer Accounts in Opsware SAS
- Associating Servers with Customers

Overview of Server Identification

Opsware SAS uses the following IDs to track managed servers:

- **MID**: Machine ID. The unique identifier that Opsware SAS uses to identify the server. The MID is usually equal to the server ID.

  The MID is stored in a file on a server's disk so that the MID can persist and be read by the Opsware Agent.

The MID follows the hard disk, not the chassis, so system administrators can swap chassis for servers without affecting how Opsware SAS tracks those servers.

See "Example: Opsware SAS Swaps a Server's Hard Disk" on page 146 in this chapter for information about swapping hard disks.

- **Server ID**: The primary key in the Opsware Model Repository (database) that represents a given server. The Server ID is used internally in Opsware SAS. Generally, users do not need this value for servers to manage them in Opsware SAS.
- **MAC Address**: Media Access Control address, which is the network interface card's unique hardware number. The MAC is used as the server's physical address on the network.
- **Chassis ID**: A unique hardware-based identifier that the Opsware Agent discovers, typically derived from some property of the server's chassis. As a common source for this ID, Opsware SAS uses an interface's MAC address or the host ID on Solaris servers, or the serial number for one of the interfaces.

Ways that Servers are Identified by Opsware SAS

Servers in the Manage Servers list are identified in the following ways when they register their hardware and software with Opsware SAS:

- Opsware SAS identifies each server by using the MID first.
- If the MID cannot be determined, the chassis ID is used to identify the server.
- If the server cannot be identified with the chassis ID, the MAC addresses are used to identify the server.

In the Server Pool, the MAC Address column displays values by which Opsware SAS tracks the servers. The value used varies by platform:

- Intel x86 processor-based servers are identified by the MAC address of the server.
- Sun SPARC processor servers are identified by the host ID of the server.

  The host ID for Sun SPARC processor servers appears in the MAC Address column in the Server Pool list.

  To determine the value in the MAC Address column, Opsware SAS uses the hardware address by which the server contacted the Opsware Build Manager (a component of the OS Provisioning feature).

Example: Opsware SAS Swaps a Server's Hard Disk

The following steps show how Opsware SAS handles swapping a hard disk for a server:

1. A system administrator swaps the hard disk of Server A (MID `1230001`, chassis ID `AB:08`) with the hard disk of Server B (MID `98730001`, chassis ID `XY:96`).
2. The Opsware Agent on Server A registers its hardware with Opsware SAS. The MID for Server A equals `1230001` and the chassis ID equals `XY:96`.
3. Opsware SAS locates Server A by using the MID.
4. Opsware SAS updates the data it has for Server A in the Model Repository. It sets the chassis ID equal to `XY:96`.
5. The Opsware Agent on Server B registers its hardware with Opsware SAS. The MID for Server B equals `98730001` and the chassis ID equals `AB:08`.
6. Opsware SAS locates Server B by using the MID.
7. Opsware SAS updates the data it has for Server B in the Model Repository. It sets the chassis ID equal to `AB:08`.

Customer Accounts in Opsware SAS

Many enterprise customers have consolidated disparate IT operations into a single operation, yet they still need separate reporting, billing, and management for different business units or groups (for example, West Coast Office, East Coast Office, and London Office).

Opsware SAS accommodates these requirements with customer accounts created by your Opsware administrator.

When your Opsware administrator creates a customer in Opsware SAS, a value for that customer is automatically added to the customer filter in the Manage Servers list, as Figure 4-19 shows.

Figure 4-19: Customer Filter in the Manage Servers List

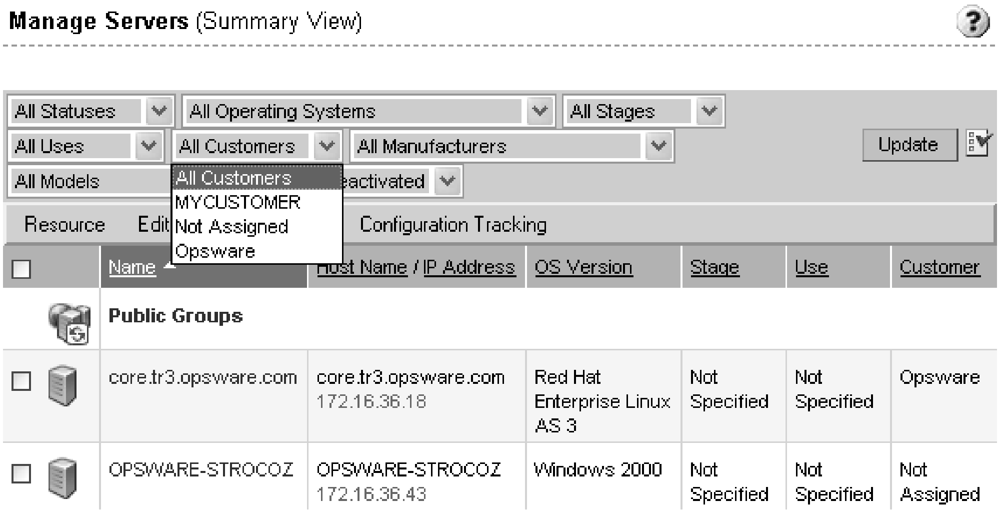

□	Name	Host Name / IP Address	OS Version	Stage	Use	Customer
	Public Groups					
□	core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	Red Hat Enterprise Linux AS 3	Not Specified	Not Specified	Opsware
□	OPSWARE-STROCOZ	OPSWARE-STROCOZ 172.16.36.43	Windows 2000	Not Specified	Not Specified	Not Assigned

By using customer accounts in the Opsware Command Center, you can segregate servers that belong to different business units. By segregating servers, you can have separate accounting for each customer or different levels of security for different customers. You might want to segregate the servers based on the department or business unit.

By default, Opsware SAS is shipped with the following two customers:

- **Customer Independent**: A global customer in Opsware SAS. Resources (applications, patches, and templates) that are associated with “Customer Independent” can be installed on any managed server, no matter what customer it is associated with.
- **Not Assigned**: The servers are not associated with a customer. You can install applications, patches, or templates that are Customer Independent on Not Assigned servers. However, you cannot install or use any resources associated with a customer on a server that is not assigned to a customer.

When you install an Opsware Agent in a server, the server is associated with the Not Assigned customer if IP ranges were not created to automatically associate Opsware SAS managed servers with customers. See Figure 4-20.

Opsware, Inc. recommends that you associate servers with customers, if necessary, by using the Server Properties pages.

Figure 4-20: Customers List Under Environment in the Opsware Command Center

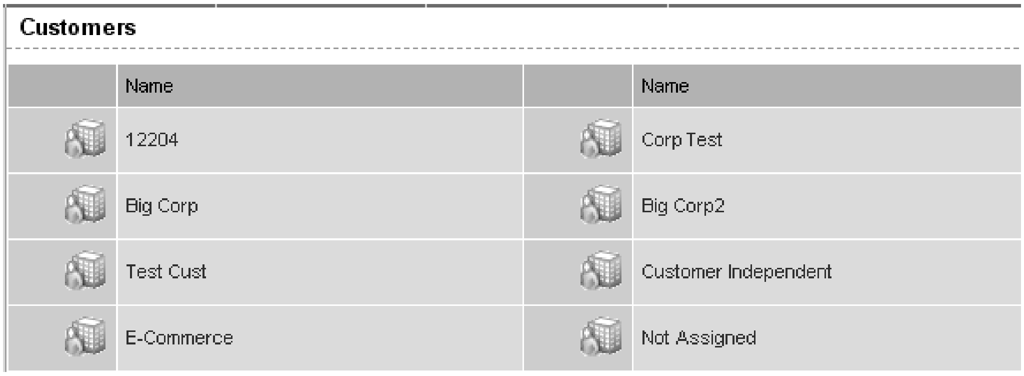

Customers

	Name		Name
	12204		Corp Test
	Big Corp		Big Corp2
	Test Cust		Customer Independent
	E-Commerce		Not Assigned

Associating Servers with Customers

An Opsware user or an Opsware administrator can set up an IP range group so that servers are automatically associated with customers when users perform the following server management tasks:

- Install Opsware Agents on the servers.

  See "Agent Reachability Communication Tests" on page 204 in Chapter 6 for more information.

- Use the OS Provisioning feature to install operating systems on bare-metal servers.

  See "Operating System Provisioning" on page 567 in Chapter 15 for more information.

To set up this automatic customer association, you must create IP range groups for customers and specify the ranges of IP addresses that the groups contain.

In the Opsware Command Center, an IP range group is both a physical and logical list – an accounting way to group ranges of IP address and assign them to a particular customer. An IP range identifies a range of IP addresses within an IP range group.

When you set this up, IP addresses get their customer association through the IP range, which, in turn, gets its customer association from the IP range group.

IP Address > IP Range > IP Range Group ↗ Customer
↘ Facility (data center
or server room)

See the *Opsware® SAS Configuration Guide* for more information.

The loose relationship between server and IP address means that you can associate a server with a different customer from its IP address.

Even when IP range groups are set up for a customer, a server's IP address does not necessarily determine the customer to which the server is associated because a user can change the customer association in the Server Properties page.

See “Editing the Properties of a Server” on page 265 in Chapter 7 for information about how to change the customer association for a server.

The customer association for a server is based on the management IP address of the server and not the primary IP address.

See “Communication Between Managed Servers and Opsware SAS” on page 221 in Chapter 7 for information about how Opsware SAS uses management IP addresses for servers.

However, a server always belongs to the same facility (data center or server room) as its primary IP address. Opsware SAS enforces the relationship between server and facility at hardware registration. See Figure 4-21.

Figure 4-21: Primary IP Addresses in Opsware SAS

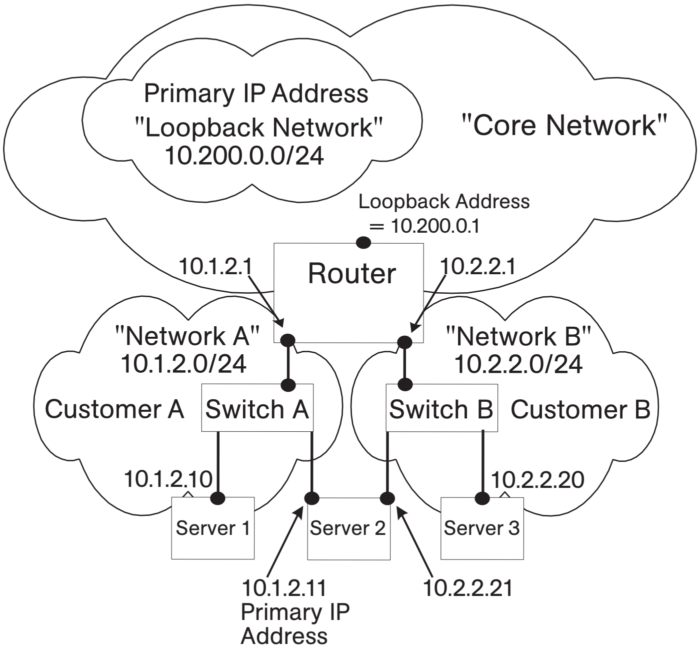

In this illustration, the following conditions apply:

- Server 1 belongs to Customer A.
- Server 2 belongs to Customer A but has IP addresses in Network A and Network B.
- Server 3 belongs to Customer B.
- The Router belongs to the Core Network but has IP addresses in Network A and Network B.

Server Histories and Reports

This section provides information on server histories and reporting with Opsware SAS and contains the following topics:

- Overview of Server History
- Viewing Server History
- Time Stamp for Server Operations
- DCI Reporting

Overview of Server History

By using the Opsware Command Center, you can view the history of changes made to a server. Each action performed on a managed server is logged in the history with the associated user who performed the action and the time of the action. You can view the history at any time, but you cannot change it. History is read-only. See Figure 4-22 for examples of recorded information.

Figure 4-22: Server History for a Server in the Manage Servers List

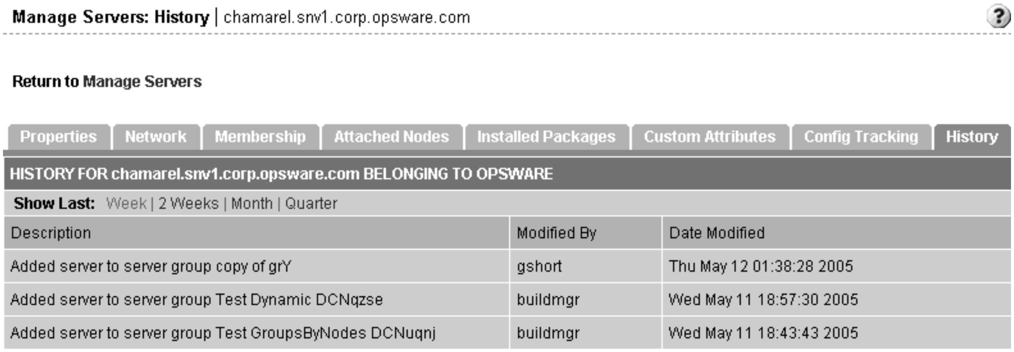

HISTORY FOR chamarel.snv1.corp.opsware.com BELONGING TO OPSWARE

Description	Modified By	Date Modified
Added server to server group copy of grY	gshort	Thu May 12 01:38:28 2005
Added server to server group Test Dynamic DCNqzse	buildmgr	Wed May 11 18:57:30 2005
Added server to server group Test GroupsByNodes DCNuqnj	buildmgr	Wed May 11 18:43:43 2005

Each History entry contains three pieces of information, as Table 4-4 shows.

Table 4-4: 'Description of the Entries in the Server History Tab

HISTORY ENTRY	DESCRIPTION
Event Description	Description of the operation performed, for example: `Install Template (Job ID: 21870101L) completed successfully.`
Modified By	The name of the Opsware user who made the change.
Date Modified	The date and time the change was made, for example: `Mon Aug 06 18:14:41 GMT+00:00 2001)`

Actions Logged in History

Opsware SAS also logs the following actions in the history for each managed server:

- Addition of the server to a node
- Removal of the server from a node
- Reassignment of the server from one node to another
- Addition of a cloned server to a node (when cloning a server, node assignments occur in Opsware SAS and are logged)
- Installation of a template on the server
- Preview reconcile failure or success
- Reconcile failure or success

Time Stamps used in History

Data is maintained for servers in Opsware SAS for the following periods of time:

- The Opsware Command Center maintains the history of changes for the last three-month interval.
- Command Engine session logs are retained for 30 days, except for the last reconcile session for a server, which is retained indefinitely.

  The Command Engine is the Opsware SAS component that enables distributed programs to run across many servers.
- Server node history is retained for 6 months.

If longer periods of time are required, Opsware, Inc. recommends regular backups to enable offline storage of Opsware SAS data.

Opsware SAS deletes old data from the Opsware Model Repository, and does not copy the data before it removes it, but you can retain information for longer periods of time, by using Oracle commands to manipulate scheduled jobs. Contact your Opsware, Inc. support representative for assistance in changing these retention periods.

Viewing Server History

Perform the following steps to view the server history:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server whose history you want to view.

   Or

   Search for the server whose history you want to view.

   See “Using the Search Feature” on page 125 in this chapter for more information. See “Server Searching by IP Address” on page 141 in this chapter for more information.

2. Click the name of the server whose information you want to see. The Manage Servers: Properties page appears for that server.

3. Select the History tab.

   By default, the view shows changes made within the past week.

Time Stamp for Server Operations

Opsware SAS maintains a comprehensive audit trail of the software that the Opsware users install, configure, and remove from a server. By using the Opsware Command Center, you can view the history of the changes made to a server. Entries are generated when actions are performed for managed servers in the Opsware Command Center. The history is read-only. See “Server Histories and Reports” on page 151 in this chapter for more information.

The time stamps for system events are determined based on the system clocks of the servers running the Opsware core components. To obtain accurate time stamps in server histories (displayed in the Opsware Command Center) and Opsware component logs, you must:

- Synchronize the system clocks on all servers running Opsware components so that all the servers are running with a common time.
- Set the time zone for all servers running Opsware components to Coordinated Universal Time (UTC).

See the *Opsware® SAS Deployment and Installation Guide* for more information on facility time requirements.

Additionally, Opsware, Inc. recommends that after installing an Opsware Agent on the server (so that it becomes managed by Opsware SAS), you should synchronize the system clock on the server with the system clocks of the servers running the Opsware core components.

The time stamps appearing in the OCC server History tab and the Opsware Agent logs are obtained from the Opsware core; however, you might need to review the server's logs (such as, stdout). Having a consistent time stamp in the server's logs and Opsware SAS is essential for effective troubleshooting.

DCI Reporting

The Model Repository maintains precise information about the state and configuration of every server under management. You can access this information through Opsware Data Center Intelligence (DCI) Reporting. DCI Reporting provides dynamic and detailed information about the operational environment, and it includes the following features:

- Exact information about the latest system state and configurations
- Visibility across the entire operational environment
- Accurate and detailed change history information
- A comprehensive set of patch reports
- The ability to extend the DCI reports

See the *Opsware System DCI 1.8 Administrator's Guide* for information about how to set up the DCI Reporting component. See the online DCI Reporting documentation for information about how to use and run the reports. Note that the DCI Reporting feature must be installed and running in the facility to view the online documentation.

The Opsware Data Center Intelligence Reporting feature is an optional component. By default, it is not installed with Opsware SAS. If this reporting feature is not available for your organization, contact your Opsware, Inc. support representative for information about how to obtain it so that you can generate reports for your managed servers.

Hardware Information for Managed Servers

The Hardware link of the Opsware Command Center provides a read-only view of all servers in your managed environment categorized by hardware manufacturer and model. The Hardware link provides hardware related information for each server, such as:

- Manufacturer
- Model number
- MAC ID
- Serial number
- CPUs used on the server
- Memory
- Storage capacity

See "Server Information that the Opsware Agent Tracks" on page 185 in Chapter 6 for more information.

Viewing Managed Server Hardware Information

Viewing hardware information allows you to see all the servers in your managed environment by hardware vendor, and view such information as MAC ID, CPUs, memory, and so on.

To view hardware of managed servers:

1. From the navigation panel, click Environment ➤ Hardware. You see the top level of the Hardware category in the managed environment, as shown in Figure 4-23.

Figure 4-23: Top level hardware node

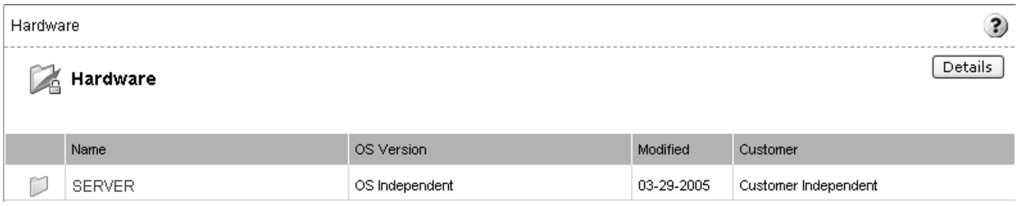

	Name	OS Version	Modified	Customer
	SERVER	OS Independent	03-29-2005	Customer Independent

2 To view the servers in your managed environment, click the Servers link.

3 Drill down to the type of server you want to look at. For example, you might want to look at all Dell POWEREDGE 650s, as shown in Figure 4-24.

Figure 4-24: Hardware home page for Dell POWEREDGE 650s

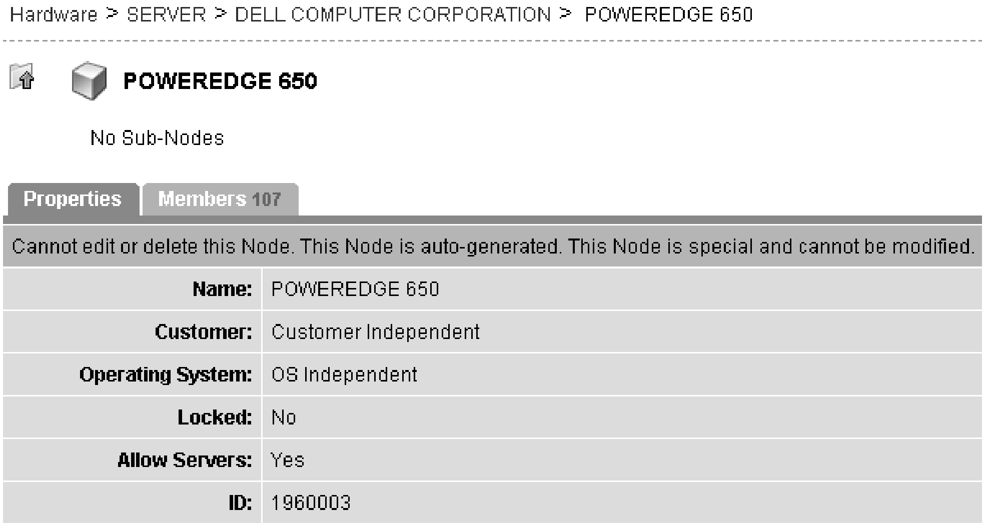

Hardware > SERVER > DELL COMPUTER CORPORATION > POWEREDGE 650

POWEREDGE 650

No Sub-Nodes

Cannot edit or delete this Node. This Node is auto-generated. This Node is special and cannot be modified.

Name:	POWEREDGE 650
Customer:	Customer Independent
Operating System:	OS Independent
Locked:	No
Allow Servers:	Yes
ID:	1960003

4 Next, select the Members tab. You will see a list of all the Dell POWEREDGE 650s in your managed environment.

5 To view specific hardware information, from the **View** menu, choose **Hardware**. You now see more detailed information about all the Dell POWEREDGE 650s computers being managed by Opsware SAS.

Chapter 5: Exploring Servers and Jobs in OCC Client

IN THIS CHAPTER

This section discusses the following topics:

- Server Explorer and Groups Browser
- Accessing the Server Explorer
- Server Information
- Server Groups Browser
- Browsing Job Logs

Server Explorer and Groups Browser

The Servers feature of the OCC Client allows you to browse and manage servers and server groups in your facility. Using the Server Explorer enables you to perform the following tasks on individual servers:

- Perform an audit.
- Take a snapshot. lion
- Configure applications.
- View summary and history information.
- Browse the file system, registry, hardware inventory, software and patch lists, and services.
- Browse Opsware information such as properties, configurable applications, and even server history.
- Drag-and-drop files between your desktop and servers.

You can also browse server groups and perform the following operations for groups of servers:

- Browse server groups and access servers inside of groups.
- Perform an audit.
- Take a snapshot.

When you perform an action on a server group, it applies to all members of the group, including other groups and their members. However, changes applied to an Application Configuration at the group level only apply to members of the group that are servers, but do not apply to sub groups of the parent group.

Accessing the Server Explorer

You can open the Server Explorer to view server information or perform operations on a server. To access the Server Explorer, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and then select All Managed Servers.

2. A list of servers will display in the Content pane.

   If the list of servers is long, You can use the search tool 🔍 to locate a server (upper right corner) by name, IP address, OS, customer, facility, or description. If you search by user name, the text entry is case insensitive.

   You can also sort the list by clicking a column heading, such as name, IP address, OS, customer, and so on. To reverse sort, click the column heading a second time.

3. Select a server and open it. This opens the Server Explorer. From the Server Explorer, you can perform the following operations from the **Actions** menu:

   - View more detail information about the selected server in the OCC.
   - Configure applications.
   - Audit application configurations.
   - Perform an audit.
   - Create a snapshot.
   - Create a package.
   - Open a remote terminal.

Server Explorer Interface

The Server Explorer consists of two main components, the server object tree in the left pane and the content area in the center pane. The server object tree lists objects from the managed server, and the content area displays content for each of the server's objects. If you select a server object, the corresponding content appears in the content area. See Figure 5-1.

Figure 5-1: Server Explorer Interface

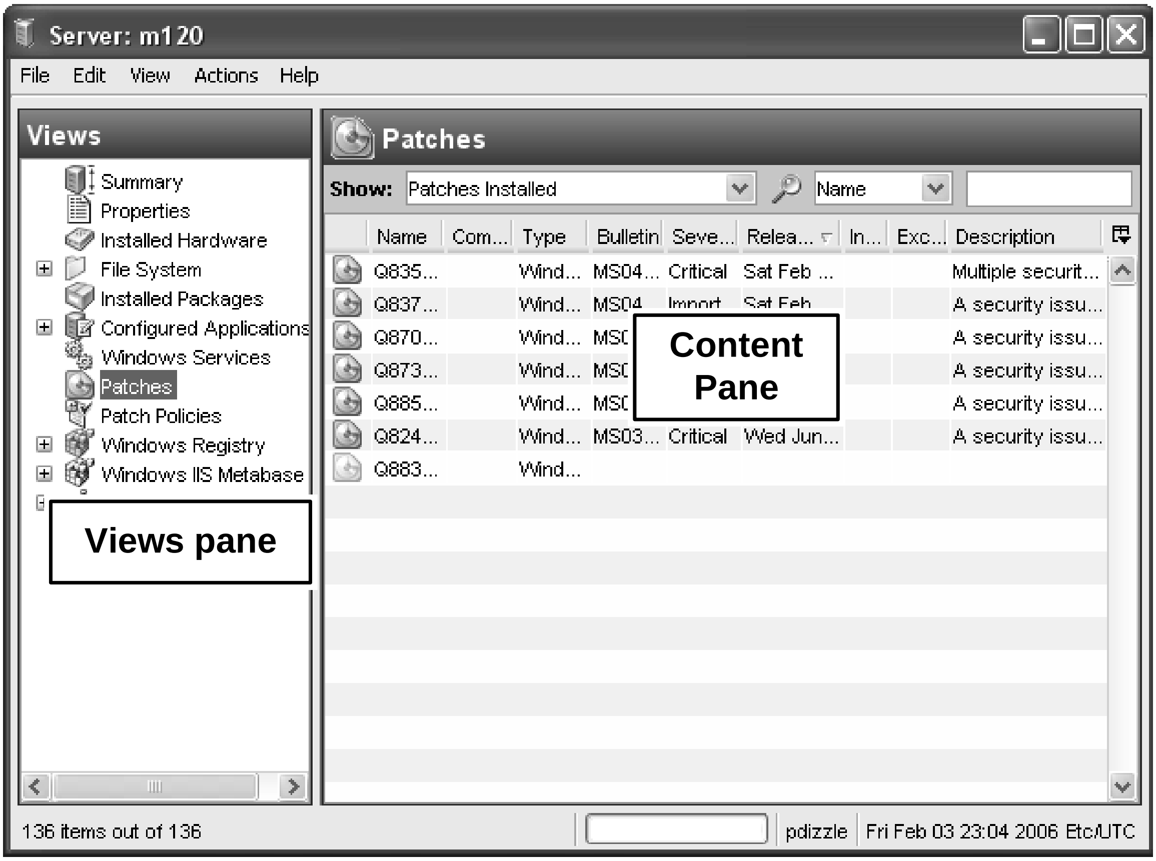

Show: Patches Installed   Name

Name	Com...	Type	Bulletin	Seve...	Relea...	In...	Exc...	Description
Q835...		Wind...	MS04...	Critical	Sat Feb ...			Multiple securit...
Q837...		Wind...	MS04	Import	Sat Feb			A security issu...
Q870...		Wind...	MS0					A security issu...
Q873...		Wind...	MS0					A security issu...
Q885...		Wind...	MS0					A security issu...
Q824...		Wind...	MS03...	Critical	Wed Jun...			A security issu...
Q883...		Wind...						

136 items out of 136   pdizzle   Fri Feb 03 23:04 2006 Etc/UTC

Actions Menu in the Server Explorer

To access features from the Server Explorer, select a server object, and select an item from the **Actions** menu. Action menu items change according to the server object selected.

For example, if you select the Configured Applications object from the server object tree, then from the **Actions** menu, you can add a configuration, remove a configuration, create a package, and so on.

Server Information

The Server Explorer allows you to review the following server information:

- Server Summary
- Server Properties
- Server Explorer File System
- Creating a Configuration Template from a File
- Creating a Package from a File
- Opening a Remote Terminal
- Viewing Installed Packages
- Configured Applications
- Configuration History
- Services
- Patches
- Patch Policies
- Windows Registry
- Windows IIS Metabase
- Windows COM+ Objects
- Compliance
- Server History Window

Server Summary

Inside the Server Explorer, the Summary window lists the following information:

- **System**: This displays operating system information.
- **Manufacturer**: This displays server manufacturer, hardware, and system details.
- **Opsware Registration**: This displays communication status, when the server was last registered, how many applications and patches are registered with Opsware SAS, and so on.

Server Properties

The Server Explorer: Properties Window lists the following property information for the server that you are browsing:

Management Information

- **Name**: This displays the name of the managed server.
- **Notes**: This displays any notes listed.
- **IP Address**: This displays the IP address of the managed server.
- **OS Version**: This displays the operating system (platform) that the managed server is running on.
- **Customer**: This displays an account within Opsware SAS that has access to designated resources, such as servers and software.
- **Facility**: This displays the location of the server. Users can manage servers in any facility from an Opsware Command Center.
- **Realm (link speed)**: This displays the minimum bandwidth limit between the Opsware Agent and the core (if the agent is going through gateways).
- **Server Use**: This displays how an organization is using the managed server; for example, a server could be a staging server, a production server, a development server, and so on.
- **Deployment Stage**: This displays the stages of deployment for a server; for example, a server could be live or offline.
- **Opsware Lifecycle**: This displays the server's stage in the Opsware Lifecycle; for example, unprovisioned, available, managed, or deactivated.
- **Server ID**: This displays the internal ID that Opsware SAS uses to identify the server.
- **Status**: This displays whether or not the server is reachable and thus managed by Opsware SAS. "OK" means that the server (its Opsware agent) is reachable; unreachable means there is a communication problem and Opsware SAS cannot communicate with the server.
- **Encoding**: This displays the character encoding of the managed server, such as Shift_JIS (Japanese) or Windows 1252 (Western).

Reported Information

- **Reporting**: This displays information about the ability of the server's agent to communicate with the core. Statuses include Has not reported, OK, Registration in progress, and Reporting error.
- **Agent Version**: This displays the version number of the agent.
- **Name**: This displays the name of the managed server.
- **Reported OS**: This displays the operating system (platform) that the managed server is running on.
- **MAC Address**: This displays the Media Access Control (MAC) address. This is the network interface card's unique hardware number. The MAC address is used as the server's physical address on the network.
- **Serial Number**: This displays the serial number of the system. Opsware SAS attempts to report a chassis ID if possible.
- **Chassis ID**: This displays a unique hardware-based identifier that the Opsware Agent discovers, typically derived from some property of the server's chassis. As a common source for this ID, Opsware SAS uses an interface's MAC address or the host ID on Solaris servers, or the serial number for one of the interfaces.

From this window, you can also open a remote terminal on the selected server.

Installed Hardware

The Installed Hardware window lists all the reported hardware on the selected managed server. This includes the following information:

- **CPUs**: This lists CPU information for all CPUs on the managed server.
- **Memory**: This lists the total amount and the types of memory on the managed server.
- **Storage**: This lists all storage devices on the managed server.

Server Explorer File System

The Server Explorer's File System element enables you to browse the file system of a managed server and perform the following operations:

- Viewing File Contents
- Copying Files Between Managed Servers
- Copying Files from Your Computer to a Managed Server
- Deleting Files
- Renaming Files

The File System has two main sections (similar to the Windows file system explorer): the server's directories and the contents of the selected directory.

The left side navigation panel of the Server Explorer shows all the directories of the selected server, and the right side of the Server Explorer lists the contents of the selected directory.

For each file, the OCC Client lists the file's name, size, type, and date modified. To sort the files by any of these categories, click on the top of the column.

Depending upon your user permissions, you might not have access to a particular server's file system. In such a case, you cannot select and view the server's file system in the Server Explorer. If you have access to a server's file system, then you will see user names, such as Administrator, root, and Local System. These are user names used to access that server's file system.

Viewing File Contents

To view file contents, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and then select All Managed Servers.
2. A list of servers will display in the Content pane. Select a server and open it. This opens the Server Explorer.
3. From the left side of the Server Explorer, select a File System object.
4. You are prompted to select a user name to log into the computer, such as Administrator, LocalSystem, or root. Select a user.

5. To view the contents of disk drives or folders, expand the icon. Select a directory.
6. From the **Actions** Menu, select **View Contents**. The file content view pane appears at the bottom of the window.
7. To change the character encoding, select an item from the Encoding drop-down list.

Ways to Copy Files

You can copy files from a server to another directory on the same server, to a directory on another Opsware-managed server, or to your local computer (where the OCC Client is running). You can also copy a file from your local computer to a directory on the managed server. A few restrictions apply when copying files on a managed server's file system using the Server Explorer:

- You cannot copy folders/directories.
- You can only copy to servers that you have permissions to write to and to view.
- You can only copy one file at a time.
- You cannot undo a deletion – once you delete a file, it's gone.

Copying Files Between Managed Servers

To copy files between managed servers, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select the Servers and then Managed Servers.
2. To launch the Server Explorer, open a server from the server list.
3. From the left side of the Server Explorer, select a File System object. To view the contents of disk drives or folders, expand the icon.
4. Navigate to the directory that contains the file that you want to copy and select it.
5. From the **Actions** Menu, select **Copy To**.
6. In the Copy To dialog box, select from the following locations in the drop-down list:
   - Managed Servers (other managed servers in the core)
   - This Server
   - Local File System
7. Navigate to the desired directory.
8. Click **Select**.

Copying Files from Your Computer to a Managed Server

To copy files from your computer to a managed server, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select the Servers and then Managed Servers.
2. To launch the Server Explorer, open a server from the server list.
3. From the left side of the Server Explorer, select a File System object.
4. Navigate to the target directory where you want to copy the file.
5. Use your system's file system explorer to select the file that you want to copy, then drag the file to the desired location in the Server Explorer.

Deleting Files

Once you delete a file, it cannot be recovered. (However, before you delete, you are prompted with a confirmation dialog box.)

To delete a file, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select the Servers and then Managed Servers.
2. To launch the Server Explorer, open a server from the server list.
3. From the left side of the Server Explorer, select a File System object.
4. Select a file to delete from the Content pane.
5. From the **Actions** menu, select **Delete**.
6. Click **Yes** in the confirmation dialog box.

Renaming Files

To rename a file, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select the Servers and then Managed Servers.
2. To launch the Server Explorer, open a server from the server list.
3. From the left side of the Server Explorer, select the File System object. To view the contents of a folder, expand the folder.
4. Select the file that you want to rename, and from the **Actions** menu, select **Rename**.

5. Enter a new name for the file, then press ENTER. Pressing the ESC key on your keyboard will cancel the rename operation.

Creating a Configuration Template from a File

For any file on a managed server, you can create an Application Configuration Template.

To create an Application Configuration Template from a file, select the file. From the **Actions** Menu, select **Create Configuration Template**. See "Creating a Configuration Template" on page 618 in Chapter 16 for more information.

Creating a Package from a File

For any file on a managed server, you can create an installable software package. For each package, you can specify the customer assignment, the reboot requirements, and the pre/post install and pre/post uninstall scripts.

To create a package from a file on the managed server file system, select the file. From the **Actions** Menu, select **Create Package**. See "Creating a Package" on page 373 in Chapter 9 for more information on how to create a package.

Opening a Remote Terminal

You can open a remote terminal for any managed server, but not a server group. To do so, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and then Managed Servers.
2. Select a managed server and open it.
1. In the Server Explorer window, from the **Actions** menu, select **Launch Remote Terminal**.
2. Log into the remote terminal.

See "Opening a Remote Terminal" on page 386 in Chapter 10 for more information.

Viewing Installed Packages

To view the installed packages on a server in the Server Explorer, select the Installed Packages object.

The Installed Packages window enables you to view any installed packages on the selected managed server. For each package, you can view name, type, size, last modified, and description. To sort the list by these categories, click the title of each column. See "Creating a Package" on page 373 in Chapter 9 for more information on how to create a package.

Configured Applications

To view Application Configurations on a server in the Server Explorer, select Configured Applications.

In the Content pane, the Installed Configurations tab of the Configured Applications window allows you to browse and edit all Application Configurations attached to the managed server.

In the Views pane, expand the Configured Application icon to display all applications being managed by the Application Configuration Management System. (If you cannot expand the Application Configurations folder, then no Application Configurations have been attached to the server.)

Each application can contain one or more instances of a managed application on the server. You can expand each application to view and edit values sets for each managed application. For more information about setting values for an Application Configuration, see "Setting Application Configuration Values on a Server or Group" on page 630.

Configuration History

The Backup Configurations tab of the Configured Applications window provides a history of all changes made to the selected application configuration template, and allows you to revert to a previous version of the configuration. Thus, you can rollback the current state of an application configuration to any previous state in this list.

Reverting to a Previous Configuration State

To rollback the configuration to a previous state, select an item in the list and click **Revert**.

Services

The Services window shows you a list of all running services on the selected managed server. Depending on the installed operating system, you can perform different operations on the services:

- For Windows services, you can start, stop, pause, resume, and restart a service. You can also set the service to start manually, to start automatically when the system is rebooted, or to be disabled altogether.
- For Linux servers (supported by Red Hat and SuSE versions), you can perform any action that a particular service supports. Supported actions may vary from service to service, for example, start, stop, restart, condrestart, or status. You can also specify the run levels that you want a service to run under.

To perform an operation on a service, select the service and right-click.

Depending upon your user's permissions, you might not have access to a server's services. In such a case, you cannot select and view the server's services in the Server Explorer. If you do have access to a server's services, then user names appear such as Administrator, root, and LocalSystem. These are user names that allow access to a server's services.

Patches

This window displays all patches associated with the selected managed server. You can use the Show drop-down list to filter what type of patch information to display in the Server Explorer.

Show Options

- **Patches Installed**: This option displays all patches that have been installed on the server.
- **Patches Recommended By Vendor**: This option displays all application and operating system patches that have been recommended by Microsoft (MBSA 2.0) for the selected server. If multiple patches have the same QNumber, Patch Management detects which application files are already installed on a managed server and, subsequently, recommends the correct patch to install.

- **Patches with Policies or Exceptions**: This option displays patches in policies attached to the selected server or patches that have always install exceptions *and* if one of the following conditions exist:
  - The patches are not currently installed and are recommended by the vendor.
  - The patches are currently installed.
- **Patches Needed**: This option displays all patches that should be installed on the selected server but are not. These include patches that are in policies attached to that server or patches that have always install exceptions *and* are recommended by the vendor.
- **Patches with Exceptions**: This option displays all patches that have exceptions (such as always install or never install) a*nd* if one of the following conditions exist:
  - The patches are not currently installed and are recommended by the vendor.
  - The patches are currently installed.
- **All Patches**: This option displays all patches that are associated with the operating system of the server.

Patch Contents

- **Icon**: A dimmed icon means that the patch has not yet been uploaded to the Software Library.
- **Name**: This is the QNumber of a patch that is a hotfix or an update rollup. Service pack patches do not have a QNumber.
- **Compliance**: This shows one of the following three levels of patch compliance, as defined by a patch administrator:
  - **Non-Compliant** (red): This indicates that the patch is installed on the server, but is not in the policy, or the patch is not installed on the server but is in the policy.
  - **Partial** (yellow): This indicates that the policy and exception do not agree, and the exception does not have data in the Reason field.
  - **Compliant** (green): This indicates one of the following conditions:
    - A patch is installed on the server and is in a policy, or a patch is not installed on the server and is not in a policy.
    - A patch is installed on the server and there are additional patches with the same QNumber in a patch policy or exception. In this case, all patches with the same

QNumber are considered installed when Patch Management calculates patch compliance.

- A patch is not installed on the server and is in a patch policy or has an always install exception, and is not recommended by the vendor. In this case, the patch is considered as if it has a never install exception because it is not recommended by the vendor.

In the Preview pane, move the cursor over the icon or text in the Compliance column to view patch compliance information about a server.

- **Type**: This shows the type of patch, such as Windows Hotfix or Windows Update Rollup.
- **Bulletin** (Optional): This shows the Microsoft Security Bulletin ID number for this patch.
- **Severity** (Optional): This shows one of the following three Microsoft severity ratings for this patch:
  - **Critical**: This indicates a patch whose exploitation could allow the propagation of an internet worm, without user action.
  - **Important**: This indicates a patch whose exploitation could result in a compromise of the confidentiality, integrity, or availability of user data, or of the integrity or availability of processing resources.
  - **Moderate**: This indicates a patch whose exploitability is mitigated to a significant degree by certain factors, such as default configuration, auditing, or difficulty of exploitation.
  - **Low**: This indicates a patch whose exploitation is extremely difficult, or whose impact is minimal.
- **Release Date**: **Release Date**: This displays the date that Microsoft released this patch.
- **Exception**: This displays the type of patch policy exception set for the selected server.
- **Installed**: This shows if the patch is installed on the selected server.
- **Recommended**: A check mark indicates that this patch was recommended by the vendor (MBSA 2.0) during the last software registration.
- **Description**: This displays a description of the server.

Patch Policies

This window displays all patch policies associated with the selected managed server (or server group). You can use the Show drop-down list to filter the type of patch policies to display in the Server Explorer.

Show Options

- **Policies Attached to the Server**: This displays all policies attached to the server, or policies attached to a server group to which the selected managed server belongs.
- **Policies Not Attached to the Server**: This displays a list of all patch policies relevant to the selected server that are not attached to the server.

Patch Contents

- **Name**: This displays the name of the patch policy.
- **OS**: This displays the operating system associated with the patch policy.
- **Description**: This shows a description of the patch policy.

Windows Registry

This window displays a read-only view of the Windows registry on the selected Windows managed server. You can navigate to this registry much like the regedit tool on the Windows operating system.

Folders on the left side of the window represent keys in the registry. Clicking a folder on the left displays entries in a key in the right window.

To view Windows Registry items, select the top-level Windows Registry icon in the Server Explorer and select a user from the menu. Your user must have proper permissions to view Windows Registry keys. If your user is unable to access the Windows Registry for the selected managed server, contact your Opsware administrator.

The HKEY_CLASSES_ROOT might have thousands of entries and can take time to load.

Windows IIS Metabase

This window displays a read-only view of the IIS Metabase on the selected Windows managed server. You can browse the IIS Metabase much like one of the metabase browsing tools such as metaedit or the IIS Metabase Browser.

The left side of the Metabase window displays the hierarchical layout of the metabase tree. Selecting an item in the tree on the left shows the data items associated with the selected key in the right hand view. Clicking the (+) symbol to the left of a key item will expand the item's child keys.

To view Windows IIS Metabase items, select the top-level Windows Metabase icon in the Server Explorer, right-click, and select a user. Your user must have permissions to view Windows Metabase items. If your user is unable to access the Windows Registry, contact your Opsware administrator.

Windows COM+ Objects

This window displays a read-only view of all the COM+ objects on the selected managed server. In the Server Explorer window, the Views pane displays two main folders for browsing COM+ objects:

- **All Objects**: This is a flat list of all the COM+ objects on the managed server.
- **Component Categories**: This contains an alphabetical list of all COM component categories.

To view the contents of a COM+ object:

1. Select the All Objects or the Component Categories folder. In the Content pane, expand the folder until you reach an object.
2. To view the contents of a COM+ object, from the Actions **Menu**, select **View Contents**. The contents will then display.
3. If the content of the COM+ object uses a different encoding, choose the appropriate encoding type from the Encoding drop-down list.

You must have specific user permissions to view Windows COM+ objects. If you are unable to access the Windows Registry, contact your Opsware administrator.

Compliance

The Compliance feature displays the following tabs:

- Server: Audit Templates Tab

- Server: Audit Results Tab
- Server: Snapshot Templates Tab
- Server: Snapshots Tab

Server: Audit Templates Tab

The Audit Templates tab displays a list of all compliance audit templates on the selected server and provides the following information:

- **Name**: The name of the compliance audit template.
- **Last Modified**: The date when the audit template was last modified.
- **Modified By**: The user who modified the audit template.

You can select an audit template from the **Actions** menu, choose to create a new audit template (this chooses the selected server as the source), open the template, delete the audit template, or perform an audit.

Server: Audit Results Tab

The Audit Results tab displays a list of all compliance audit results that used the selected server as a source or target for the audit. This window provides the following information:

- **Name**: The name of the compliance audit results.
- **Last Modified**: The date when the compliance audit results were created.
- **Modified By**: The user who ran the compliance audit results.

You can select an audit result from the **Actions** menu, choose to view the results, rerun an audit, delete the audit results, or create a package from the audit.

Server: Snapshot Templates Tab

The Snapshot Templates tab displays a list of all your snapshot templates on the selected server and provides the following information:

- **Name**: The name of the compliance snapshot template.
- **Last Modified**: The date when the compliance snapshot template was last modified.
- **Modified By**: The user who modified the compliance snapshot template.

You can select a snapshot template from the **Actions** menu, choose to create a new snapshot template, open the template, create a snapshot from the selected template, delete the snapshot template, or create a package.

Server: Snapshots Tab

The Snapshots tab displays a list of all snapshots performed on the selected server and provides the following information:

- **Name**: The name of the compliance snapshot.
- **Modified**: The date when the compliance snapshot was last modified.
- **By**: The user who ran or modified the compliance snapshot.

You can select a snapshot from the **Actions** menu, choose to open the snapshot, delete the snapshot, perform an audit, or create a package of the snapshot.

Server History Window

The Server History window shows changes made to a managed server. For example, it displays who modified a server, what change was made, when it was modified, and so on. Server history specifically shows when a user has performed one of the following actions:

- Added the server to a node
- Removed the server from a node
- Reassigned the server from one node to another

Entries are generated when actions are performed for managed servers in the Opsware Command Center. The History is read-only. Each entry shows the following information:

- **Description**: A description of the change.
- **Last Modified**: The date when the last change occurred.
- **Modified By**: The user who made the change.

Use the View drop-down list to sort the server history list according to a range of time, such as last week, the last two months, and so on.

Server Groups Browser

This section discusses the following topics:

- Accessing the Server Groups Browser
- Server Group Members
- Configured Applications for Server Groups
- Server Group History Properties

The Server Groups Browser provides access to all server groups in the core. From the Server Groups Browser, you can perform the following actions:

- Browse server groups and access servers inside of groups.
- Perform an audit.
- Take a snapshot.

Accessing the Server Groups Browser

To access the Server Groups Browser, perform the following steps:

1. Launch the OCC Client. Select Servers and then select Server Groups.
2. Select a server group and open it. (You can expand a group to find sub groups.)
3. If the list of server groups is long, sort the list by clicking a column name, such as name, IP address, OS, customer, or facility.
4. For each server group, you can also perform an audit, take a snapshot, or configure applications.

Server Group Members

From inside each server group, you can view all members that belong to the group. This can include servers as well as other groups. For each server that belongs to the group, the system displays its name, IP address, OS, customer, facility, and any description.

Configured Applications for Server Groups

If the group is public, then you can add an Application Configuration to the group. The Application Configuration applies to all servers and groups in this group.

- The Installed Configurations tab allows you to browse and edit all Application Configurations attached to the server group.
- The Backup Configurations tab provides a history of all changes made to the selected application configuration template, and allows you to revert to a previous version of the configuration.

See Chapter 16, “Application Configuration Management” on page 599 of this guide for more information.

Server Group Patches

This window displays all patches associated with the selected server group. You can use the Show drop-down list to filter the type of patch information displayed in the Server Groups Browser.

Show Options

- **Patches with Exceptions**: This option displays all patches that have exceptions, such as always install or never install a*nd* if one of the following conditions exist:
  - The patches are not currently installed and are recommended by the vendor.
  - The patches are currently installed.
- **All Patches**: This displays all patches that are associated with the operating system of a server.

Patch Contents

- **Icon**: This displays a dimmed patch icon when the patch has not yet been uploaded to the Software Library.
- **Name**: This indicates the QNumber of a patch that is a hotfix or an update rollup. Service pack patches do not have a QNumber.
- **Compliance**: This shows one of the following three levels of patch compliance, as defined by a patch administrator:
  - **Non-compliant** (red): This indicates that the patch is installed on the server, but that it is not in the policy, or the patch is not installed on the server but is in the policy.
  - **Partial** (yellow): This indicates that the policy and exception do not agree, and that the exception does not have data in the Reason field.
  - **Compliant** (green): This indicates any of the following conditions:
    - – A patch is installed on the server and is in a policy, or a patch is not installed on the server and is not in a policy.
    - – A patch is installed on the server and there are additional patches with the same QNumber in a patch policy or exception. In this case, all patches with the same QNumber are considered installed when Patch Management calculates patch compliance.
    - – A patch is not installed on the server and is in a patch policy or has an always install exception, and is not recommended by the vendor. In this case, the patch is considered as if it has a never install exception because it is not recommended by

the vendor.

In the Preview pane, move the cursor over the icon or text in the Compliance column to view patch compliance information about a server.

- **Type**: This indicates the type of patch, such as Windows Hotfix or Windows Update Rollup.
- **Bulletin** (Optional): This indicates the Microsoft Security Bulletin ID number for this patch.
- **Severity** (Optional): This displays one of following Microsoft severity ratings for this patch:
  - **Critical**: This indicates a patch whose exploitation could allow the propagation of an internet worm, without user action.
  - **Important**: This indicates a patch whose exploitation could result in a compromise of the confidentiality, integrity, or availability of user data, or of the integrity or availability of processing resources.
  - **Moderate**: This indicates a patch whose exploitability is mitigated to a significant degree by factors, such as default configuration, auditing, or difficulty of exploitation.
  - **Low**: This indicates a patch whose exploitation is extremely difficult, or whose impact is minimal.
- **Release Date**: This shows the date that Microsoft released this patch.
- **Exception**: This indicates the type of patch policy exception set for the selected server.
- **Installed**: This indicates if the patch is installed on the selected server.
- **Recommended**: A check mark indicates that this patch was recommended by the vendor (MBSA 2.0) during the last software registration.
- **Description**: This shows a description of the server.

Server Group Patch Policies

This window displays all patch policies associated with the selected server group. You can use the Show drop-down list to filter the type of patch policies to display in the Server Explorer.

Show Options

- **Policies Attached to Server Group**: This displays all policies attached to the server group, or policies attached to a server group to which the selected managed server belongs.
- **Policies Not Attached to the Server:** This displays a list of all patch policies relevant to the selected server group that are not attached to the server group.

Patch Contents

- **Name**: This displays the name of the patch policy.
- **OS**: This displays the operating system associated with the patch policy.
- **Description**: This shows a description of the patch policy.

Server Group History Properties

The Server Groups Explorer shows the following information for each server group:

- **Type**: This indicates if the group is dynamic (rule-based) or static.
- **Rules**: If the group is dynamic, this displays all of the group's rules.
- **Status**: This displays the group's status. Active means that the group is accessible. "In use by access control" means that an administrator has set access control boundaries (security permissions) for the group.
- **Accessibility**: This indicates if the group is public or private.
- **Servers (this level)**: This indicates the number of servers that belong to this group at this level in the group hierarchy.
- **Groups (this level)**: This indicates the number of groups that belong to this group at this level in the group hierarchy.
- **Unique Servers (all levels)**: This shows the number of servers for the current group and all the servers in its subgroups, recursively. For example, if the group has one server, ServerA, and a sub-group, Group, containing two servers, then this group will contain a total of three unique servers.
- **Date last used (by a job)**: This displays the most recent job run on the group.

Browsing Job Logs

A job is any major process run by the Opsware Command Center or the Opsware Command Center Client, such as a communication test, a software installation, an Application Configuration push, a server compliance audit, and so on.

The Job Logs window shows the details of all jobs run under the currently logged in user name and jobs that are currently in progress. It also displays jobs scheduled to run, including the name of the job, the start time, the number of servers affected by the job, and the status of the job.

The display of a job's start time and finish time is determined by original user preferences set in the Opsware Command Center, which may be different than those of the current user. If a user is working in multiple times zones, it is a good idea to make sure that these preferences are set to display the time zone in the date.

To see details of a finished job, open a job. If the job is recurring (scheduled to be run), then opening the job will cause the Schedule Job dialog box to display. You will only be able to modify the scheduled job if you created the job, or have the Edit All Jobs permissions. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

You can search the job list by selecting the following criteria (located at the top of the Job Logs window):

- **Job ID**: Enter the job ID.
- **Job Type**: Choose a type of job from this list.
- **Time Restrictions**: Limit the search for a job by a time restriction such as the last week, last two weeks, or last month.
- **Job Status**: Choose a job status, such as job completed, completed with errors, cancelled, and so on.
- **My Jobs/All Jobs**: Choose to show all jobs or just your jobs. Click on the column header to sort the list. Only users with the View All Jobs or Edit All Jobs permissions will be able to view all jobs in the core. If you do not have these permissions, you will not see these options.

Chapter 6: Opsware Agent Management

IN THIS CHAPTER

This section discusses the following topics:

- Opsware Agent on Managed Servers
- Opsware Discovery and Agent Deployment
- Opsware Agent Installation Using ODAD
- Viewing Unmanaged Server Information
- Agent Reachability Communication Tests

Opsware Agent on Managed Servers

This section provides information about the Opsware Agent on managed servers and contains the following topics:

- Overview of the Opsware Agent on Managed Servers
- Security for Opsware Agents Running on Managed Servers
- Opsware Agent Functionality on Managed Servers
- Server Information that the Opsware Agent Tracks

Overview of the Opsware Agent on Managed Servers

The Opsware Agent regularly performs management tasks on each server autonomously:

On a regular interval, the Opsware Agent gathers a hardware and software inventory of each managed server. It opens a secure communication channel to the Opsware core, presenting its IP address and public-key certificate for authentication purposes. If properly authenticated, the Opsware Agent is permitted to write its updates about the server to the Opsware Model Repository.

Every 12 hours, the Opsware Agent submits hardware information for the managed server on which it is running. Hardware registration also occurs during Opsware Agent installation or software installation. Every 24 hours, the Opsware Agent submits software information for the managed server to the Opsware core. See Figure 6-1.

Figure 6-1: Server Management Tasks Performed by Opsware Agent

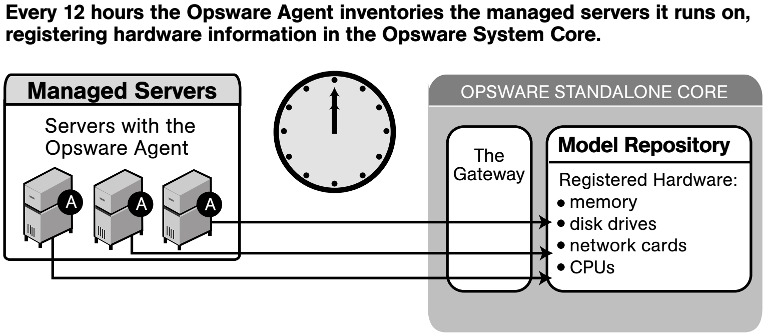
Every 12 hours the Opsware Agent inventories the managed servers it runs on, registering hardware information in the Opsware System Core.
Managed Servers
Servers with the Opsware Agent
A
A
A
OPSWARE STANDALONE CORE
The Gateway
Model Repository
Registered Hardware:
• memory
• disk drives
• network cards
• CPUs

The Reporting field indicates the status of the Opsware Agent's reporting capability and tells you whether or not the Opsware Agent is reporting regularly and successfully. The four possible reporting states for the Opsware Agent are as follows:

- **OK**: The Opsware Agent is reporting properly.
- **Registration in progress**: The Opsware Agent is currently registering server hardware information.
- **Reporting error**: The Opsware Agent encountered an error while trying to report hardware information.
- **Last reported days ago**: This indicates when the Opsware Agent last reported.

You can access Opsware Agent reporting information in Server Properties by using advanced search, and by viewing managed servers by Communication status. When viewing by Communication status, the Opsware Command Center user interface displays this information in the registration column.

If the Opsware Agent experiences an error in reporting, or has not reported within 24 hours, you can run a Communication Test to troubleshoot the problem.

If you modify the server hardware, it could take up to 12 hours for the change to appear in the Opsware Command Center user interface, depending on the time that the Opsware Agent for that server contacted the Opsware core.

If you install or uninstall software on a managed server outside of Opsware SAS, it could take up to 24 hours for the change to appear in the Opsware Command Center user interface. For example, if you update the Microsoft Patch Database, it could take up to 24 hours for all managed servers to display whether they need new patches based on the updated Microsoft Patch Database.

In some cases, not all of a server's hardware information is reported. For example, if the Opsware Agent was installed with its default settings, not all hardware information is reported to the Opsware Command Center until an hour after the agent is installed. There also might be a problem retrieving certain hardware information, such as a disk failure, that could prevent some hardware information from being reported. In these cases, the server's property page lists unreported information as not set.

If configuration tracking is enabled for a server, the Opsware Agent sweeps through the managed server on a regular interval to see if any of the configurations being tracked are changed. If a tracked configuration is changed, the Opsware Agent performs the action specified by the tracking policy, namely, writing the information to a log file, generating a backup, or sending an email message through SMTP to the email address specified in the tracking policy.

Security for Opsware Agents Running on Managed Servers

Opsware Agents act as both clients and servers when they communicate with Opsware SAS. All communication is encrypted, integrity-checked, and authenticated using X.509v3 client certificates using SSL/TLS.

A small number of Opsware core components can issue commands to the Opsware Agent over a well-defined TCP/IP port. The Opsware Agent can also call back to Opsware core components, each with its own well-defined port.

The Code Deployment & Rollback (CDR) feature uses agent-to-agent communication for performance reasons. In particular, CDR synchronizations (the process of copying changed files and directories from one server to another) happen when an Opsware Agent connects to another Opsware Agent and sends across the network files that have changed since the last time the two Opsware Agents connected.

To further safeguard the SSL/TLS-based communication channel, the two Opsware Agents participating in the code deployment also need to have a common shared secret provided by the Command Engine. Before one Opsware Agent can begin a file transfer to another Opsware Agent, the two agents must verify that they have a common shared secret provided on a per-session basis by the Command Engine. This safeguard prevents unauthorized users from copying files from one managed server to another.

Opsware Agent Functionality on Managed Servers

The Opsware Agent is designed such that:

- It can only discover information about its own managed server (and no others).
- It cannot make changes on a server unless explicitly instructed to do so by an Opsware core component.

Opsware SAS runs with administrator privileges (root on UNIX servers and Local System on Windows servers), because it performs tasks that require administrator privileges, such as installing patches and rebooting servers.

The Opsware core performs client authentication and, checks to see if the presenting certificate belongs to that particular server. Opsware SAS does this by comparing the certificate to the server's IP address that is generated when the Opsware Agent is initially installed. If the certificate is not valid or the originating IP address does not match the IP address stored in the Opsware Model Repository, authentication fails and the Opsware Agent cannot continue communication with Opsware SAS.

If an unauthorized user were able to log on to a managed server with administrator privileges and compromise a server's security, the user would have only limited access to the following information in the Opsware Model Repository:

- The server's own hardware inventory (already available to someone logged on with administrator privileges)
- The server's own software inventory (already available to someone logged on with administrator privileges)
- The set of assignments from itself to the nodes in the Software Tree
- The custom attributes contained in those nodes

Server Information that the Opsware Agent Tracks

For each managed server, the Opsware Agent reports software, networking, and hardware information, as shown in Figure 6-2 and Figure 6-3. By communicating with the Opsware core and reporting the installed hardware and software for the server, Opsware SAS determines what software should be installed on a server.

Figure 6-2: Manager Servers: Properties Page-Server Information

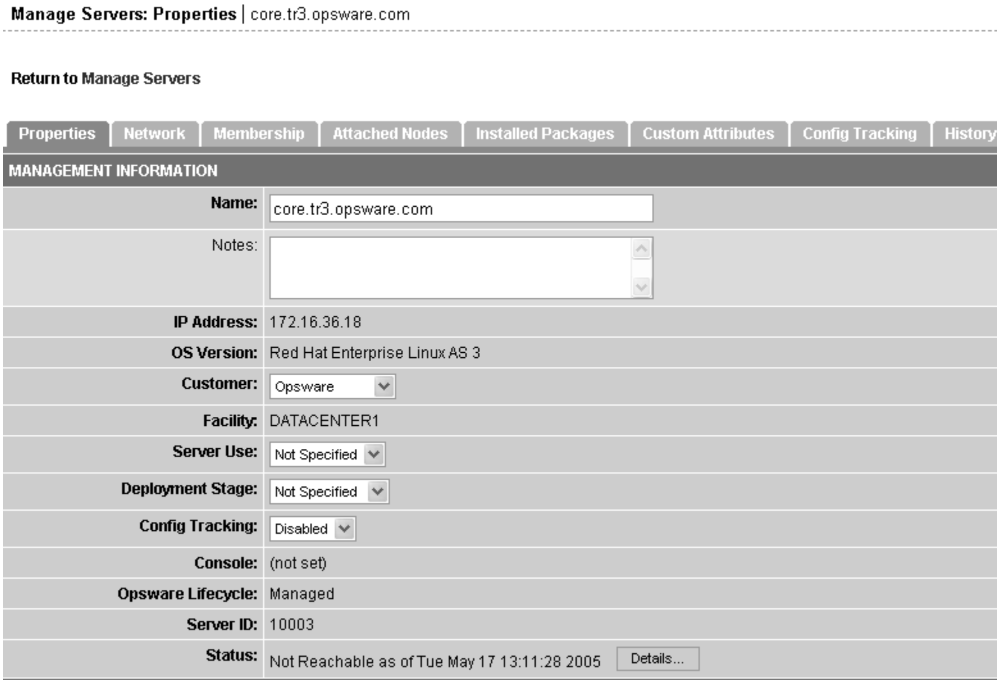

Manage Servers: Properties | core.tr3.opsware.com

Return to Manage Servers

MANAGEMENT INFORMATION

Name:	core.tr3.opsware.com
Notes:	
IP Address:	172.16.36.18
OS Version:	Red Hat Enterprise Linux AS 3
Customer:	Opsware
Facility:	DATACENTER1
Server Use:	Not Specified
Deployment Stage:	Not Specified
Config Tracking:	Disabled
Console:	(not set)
Opsware Lifecycle:	Managed
Server ID:	10003
Status:	Not Reachable as of Tue May 17 13:11:28 2005 Details...

Figure 6-3: Server Properties-Hardware and Additional Information

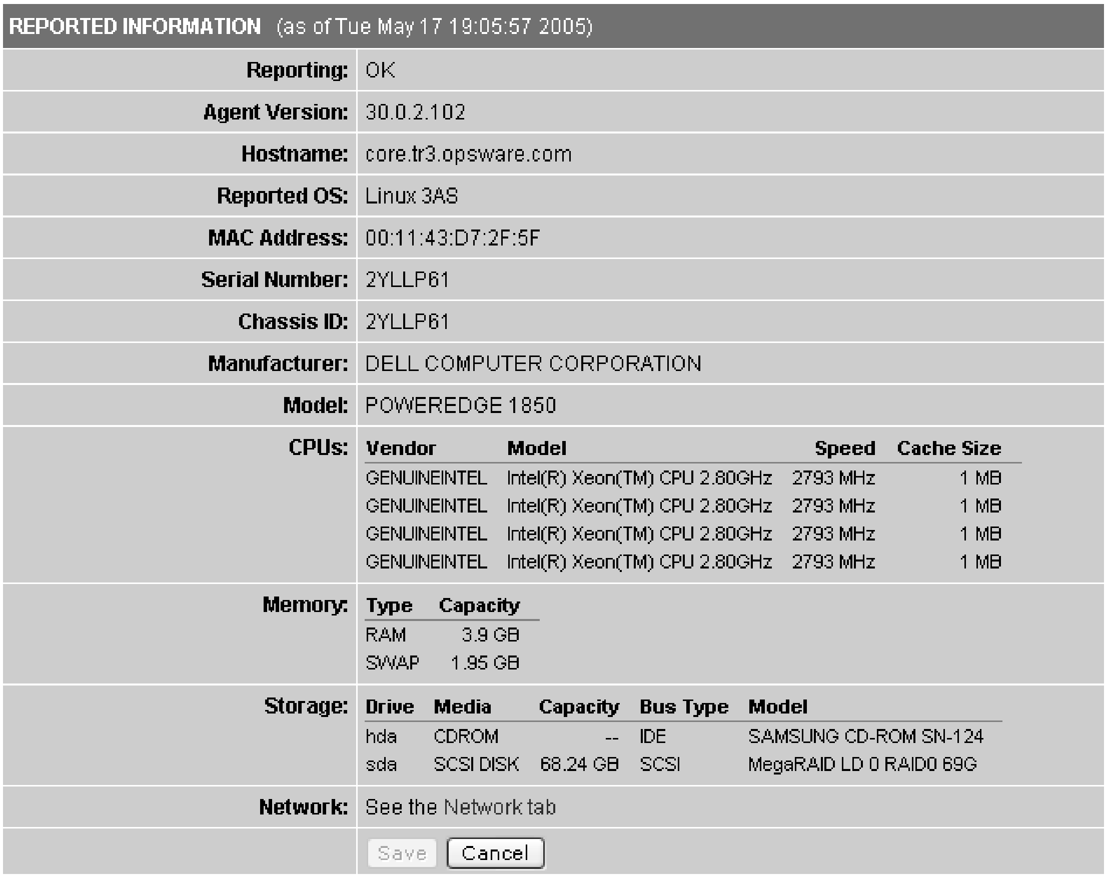

REPORTED INFORMATION (as of Tue May 17 19:05:57 2005)

Reporting:	OK
Agent Version:	30.0.2.102
Hostname:	core.tr3.opsware.com
Reported OS:	Linux 3AS
MAC Address:	00:11:43:D7:2F:5F
Serial Number:	2YLLP61
Chassis ID:	2YLLP61
Manufacturer:	DELL COMPUTER CORPORATION
Model:	POWEREDGE 1850

CPUs:

Vendor	Model	Speed	Cache Size
GENUINEINTEL	Intel(R) Xeon(TM) CPU 2.80GHz	2793 MHz	1 MB
GENUINEINTEL	Intel(R) Xeon(TM) CPU 2.80GHz	2793 MHz	1 MB
GENUINEINTEL	Intel(R) Xeon(TM) CPU 2.80GHz	2793 MHz	1 MB
GENUINEINTEL	Intel(R) Xeon(TM) CPU 2.80GHz	2793 MHz	1 MB

Memory:

Type	Capacity
RAM	3.9 GB
SWAP	1.95 GB

Storage:

Drive	Media	Capacity	Bus Type	Model
hda	CDROM	--	IDE	SAMSUNG CD-ROM SN-124
sda	SCSI DISK	68.24 GB	SCSI	MegaRAID LD 0 RAID0 69G

Network: See the Network tab

Software Information

The software that should be installed is recorded in the Opsware Model Repository. To access that information, click the Install List or Installed Packages tabs. This shows the software that should be installed on the server or all software that is installed on the server.

To display the list of what software packages should be installed on that server by virtue of that server's assigned nodes, select the Install List tab from the Manage Servers: Server Properties page.

To display the list of software that is reportedly installed on the server, select the Installed Packages tab from the Managed Servers: Server Properties page.

Partially-installed Solaris packages do not show up in the Installed Packages list, even though the package was partially installed.

See the *Opsware® SAS Configuration Guide* for information about how to set up the Software Tree.

Hardware Information

Opsware SAS tracks hardware information in a variety of manners. Table 6-1 shows how the Opsware Agent obtains the server and hardware information about each managed server.

Table 6-1: Hardware Information that the Opsware Agent Reports for Servers

ATTRIBUTE	DESCRIPTION	HOW OBTAINED
Name	The user-configurable name for the server. By default, Opsware SAS uses the configured host name of the server until a user edits it.	**Windows**: Uses the fully qualified DNS name of the server. **Linux, Solaris, AIX, HP-UX**: Uses the current host name of the server that the `hostname` command returns.
Reported OS	The version number of the server's operating system.	**Windows**: Uses the Windows version number as reported by the operating system. This information includes the major version number, the minor version number, the Windows build number, and the Service Pack level. **Linux, Solaris, AIX, HP-UX**: Uses the operating system version that the `uname` command returns.
OS Version	The OS version specified for the OS definition.	Specified by the user who prepared the OS with the Prepare Operating System Wizard. See the *Opsware® SAS Configuration Guide* for more information.
Serial Number	The serial number of the system. Opsware SAS attempts to report a chassis ID, if possible.	**Windows, Linux**: Obtained from the system BIOS. **Solaris, AIX, HP-UX**: Obtained from the system ROM.

Table 6-1: Hardware Information that the Opsware Agent Reports for Servers (continued)

ATTRIBUTE	DESCRIPTION	HOW OBTAINED
Manufacturer	The manufacturer of the server if available.	**Windows, Linux**: Obtained from the system BIOS. **Solaris, AIX, HP**: Obtained from the system ROM.
Model	The model of the server if available.	**Windows, Linux**: Obtained from the system BIOS. **Solaris, AIX**: Obtained from the system ROM. **HP-UX**: Output of model command (which is read from the system ROM).
Memory	The amount of physical RAM and the total amount of virtual memory paging space configured.	**Windows**: Uses the Windows 2000 API `GlobalMemoryStatus()`. **Linux:** Obtained from information in the file `/proc/meminfo`. **Solaris**: Obtained from the `sysconf` and `swapctl` APIs. **AIX**: Uses the `lsattr` command for memory information and the `lsps` command for paging space. **HP-UX**: Uses the `pstat` system call.
Processors	Information about each of the processors in the system.	**Windows**: If WMI is available, iterates over all instances of `Win32_Processor`. If WMI is not available, parses the registry key `HARDWARE\DESCRIPTION\System \CentralProcessor`. There is one sub-key for each processor. **Linux**: Obtained from information in the file `/proc/meminfo`. **Solaris, HP-UX**: Uses system APIs to enumerate the processors in the system. **AIX**: Uses the `lscfg` command.

Table 6-1: Hardware Information that the Opsware Agent Reports for Servers (continued)

ATTRIBUTE	DESCRIPTION	HOW OBTAINED
Storage	Information about each installed disk drive or RAID array.	**All Platforms**: Uses system APIs to discover and probe disk drives and RAID arrays.
Server ID	The internal ID that Opsware SAS uses to identify the server.	In most cases, the server ID is the same as the MID.
MID	The MID (Machine ID) is a unique number that Opsware SAS assigns when the server first registers. The server stores the MID and reports it each time the server registers.	**Windows**: The MID is stored in the file `%ProgramFiles%\Common Files\Loudcloud\cogbot\mid` if present. **Linux, Solaris, AIX, HP-UX**: The MID is stored in the file `/var/lc/cogbot/mid`.

In addition to hardware and software reporting, the Opsware Agent reports networking information. See “Network Configuration” on page 280 in Chapter 7 for descriptions of the networking information reported and how you can modify that information using the Network tab in the Opsware Command Center.

Opsware Discovery and Agent Deployment

The Opsware Discovery and Agent Deployment (ODAD) feature helps to deploy Opsware Agents to a large number of servers through the Opsware Command Center Client (OCC Client). This feature helps to identify servers on which to install an Opsware Agent, specify the deployment actions to be carried out on each server, select the login protocols to connect to each server, specify the Opsware Agent Installer options for installing an Opsware Agent, and generate reports on Opsware Agent installation status. See Figure 6-4.

Figure 6-4: Agent Deployment Process

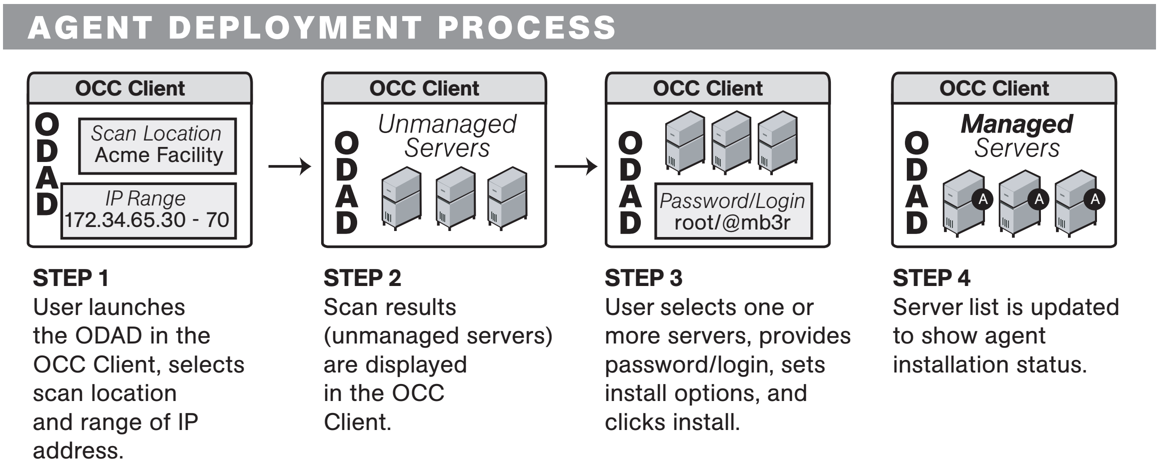

STEP 1
User launches the ODAD in the OCC Client, selects scan location and range of IP address.

STEP 2
Scan results (unmanaged servers) are displayed in the OCC Client.

STEP 3
User selects one or more servers, provides password/login, sets install options, and clicks install.

STEP 4
Server list is updated to show agent installation status.

This overview section contains the following topics:

- Discovering Servers for Installing an Opsware Agent
- Setting Deployment Actions for Each Server
- Specifying Login Settings
- Opsware Agent Installer Options
- Reports on Server Status

See "Installing Opsware Agents Using ODAD" on page 196 in this chapter for information about how to install an Opsware Agent.

Discovering Servers for Installing an Opsware Agent

Using the Opsware Discovery and Agent Deployment feature, you can select a location to scan for servers. After selecting a location, you can specify the IP addresses or IP address ranges to perform a network scan to identify servers in which to install an Opsware Agent. Instead of performing a network scan, you can also import a file containing a list of IP addresses or IP address ranges. When the scan is complete, a list of scanned servers is shown.

For each server, this feature determines the status of the server, its IP address, its host name, detected operating system, and open ports used to connect to the server.

Setting Deployment Actions for Each Server

Once you have identified the servers, you can select the servers and perform the following deployment actions:

- Verify installation prerequisites.

  When you select this action, verification checks are performed to ensure that the Opsware Agent is successfully installed on the server. The verification checks include:

  - Checking for sufficient disk space for Opsware Agent installation on the server
  - Verifying that no other application is using port 1002
  - Verifying if ports to the Opsware Gateway are accessible

- Verify prerequisites, and copy the Opsware Agent Installer to servers.

  When you select this action, verification checks are performed to ensure that the Opsware Agent is successfully installed on the server and the Opsware Agent Installer is copied to the server.

- Verify prerequisites, copy installer, and install Opsware Agent.

  When you select this action, verification checks are performed to ensure that the Opsware Agent is successfully installed on the server, the Opsware Agent Installer is copied to the server, and the Opsware Agent is installed on the server.

Specifying Login Settings

After selecting the deployment action, you can select the network protocols to connect to the server and specify the user name and password to login to each server. The Opsware Agent needs administrator-level privileges (root on Unix servers and administrator on Windows servers) to manage a server. Therefore, Opsware Agent installation is performed as root on Unix operating systems and as administrator on Windows operating systems. ODAD tries to log into each of the selected servers with the specified user name and password and performs the specified deployment actions.

Opsware Agent Installer Options

The Opsware Discovery and Agent Deployment feature allows you to specify the installer options listed in Table 6-2.

Table 6-2: Opsware Agent Installer Options

OPTION	DESCRIPTION
Start the Opsware Agent after Installation	Starts the Opsware Agent after installing it on the server. By default, the Opsware Agent Installer does not start the Opsware Agent.
Remove existing machine IDs and cryptographic material	Removes any existing machine specific identifying material such as Machine ID file (MID), and all machine-specific cryptographic material from the server.
Ignore prerequisite check failures	Ignores the prerequisite check failures and forces Opsware Agent installation.
Set the server's time from the Opsware core	Synchronizes the time on the server in which the Opsware Agent is installed to the Opsware core.
Install Windows Installer (MSI) if required	Installs MSI along with the Opsware Agent. If MSI is already installed, this option has no effect.
Install Windows Management Instrumentation (WMI) if required	Installs WMI along with the Opsware Agent. If WMI is already installed, this option has no effect.
Reboot Windows servers after agent installation	Reboots windows servers after Opsware Agent installation is complete.
Install Red Hat Package Manager (RPM) on AIX and Solaris	Installs the RPM handler with the Opsware Agent. Always include this option when you install Opsware Agents on Solaris and AIX servers.
Reset Opsware Agent configuration, if present	Replaces the existing Opsware Agent configuration.

Table 6-2: Opsware Agent Installer Options (continued)

OPTION	DESCRIPTION
Delete gateway address list, if present	Deletes the Opsware Gateway address list if present and is no longer required.
Overwrite staged Opsware Agent installer	Overwrites the existing Opsware Agent Installer.
Log Level	Sets the log level for log messages. With this option, you can specify the log levels of error, warning, info, and trace.
Use Opsware Gateways	Specifies the Opsware Gateways used during Opsware Agent installation.
Immediately do a full hardware registration	Forces the Opsware Agent Installer to report full hardware information to the Opsware core.
Immediately do software registration	Forces the Opsware Agent Installer to report full software information to the Opsware core.
Suppress Opsware Agent reachability check	By default, the installer triggers the core to check if the server is reachable. This option disables this check during installation.
Disallow anonymous SSL connections if Opsware Agent is dormant	Configures the Opsware Agent so that browsers cannot connect without a valid certificate.
Force creation of new device record if conflict found	During registration, the Data Access Engine creates a new device record. This option suppresses this functionality.
Fail if initial hardware registration fails (do not go dormant)	Does not allow the Opsware Agent to become dormant, if it fails to report hardware information.
Reconcile Type	Reconciles the server against any nodes assigned to the server. The reconcile type can be Full or Add only.

Table 6-2: Opsware Agent Installer Options (continued)

OPTION	DESCRIPTION
Attach to template ID	Assigns the nodes contained in the template to the server.
Extra Installer options	Allows you to specify any other installer options.

Reports on Server Status

After the deployment action is complete, the OCC Client displays the results and updates the status icons for the servers as shown in Table 6-3.

Table 6-3: Server Status

ICONS	SERVER STATUS
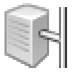	The server is unmanaged.
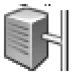	The server is managed by Opsware.
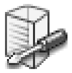	The server failed prerequisite checks.
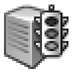	The server passed prerequisite checks.
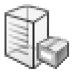	The server passed prerequisite checks and the Opsware Agent Installer was copied to the server.

Table 6-3: Server Status (continued)

ICONS	SERVER STATUS
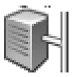	The Opsware Agent was successfully deployed.
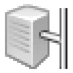	The Opsware Agent was not successfully deployed.

A server is considered to be managed by Opsware when the ODAD determines that the Opsware Agent is listening for TCP connections on port 1002.

For a failed deployment action, you can view the errors on each server. You can also log into the server from this feature and correct the errors.

Using this feature, you can create the following reports:

- All the servers in the current network scan
- Selected servers in the current network scan
- All the servers in the current network scan with successful deployments
- Servers in the current network scan with failed deployments

You can save and export the reports to a CSV, HTML, or text format file.

Opsware Agent Installation Using ODAD

The Opsware Discovery and Agent Deployment (ODAD) feature helps you to deploy Opsware Agents to a large number of servers through the OCC Client. This section contains the following topics:

- Prerequisite Setup for Opsware Discovery and Agent Deployment
- Permissions Required for Opsware Discovery and Agent Deployment
- Installing Opsware Agents Using ODAD

Prerequisite Setup for Opsware Discovery and Agent Deployment

Before you can install Opsware Agents using ODAD, you must install the Windows Agent Deployment Helper package (required for installing Opsware Agents on Windows server only) and install JRE version 1.4.2 06 for running the OCC Client.

See the *Opsware® SAS Deployment and Installation Guide* for more information about installing the Windows Agent Deployment Helper Package.

See "Requirements for Running the OCC Client" on page 97 in Chapter 3 for information about installing JRE.

Permissions Required for Opsware Discovery and Agent Deployment

To use the Opsware Discovery and Agent Deployment feature you must have certain permissions. See *Opsware® SAS Configuration Guide* for more information about the permissions required to use the ODAD feature.

To obtain the required permissions for scanning and deploying Opsware Agents, contact your Opsware administrator.

Installing Opsware Agents Using ODAD

To install Opsware Agents using ODAD, launch the OCC Client from the Opsware Command Center (OCC). See "Installing and Launching the OCC Client" on page 97 in Chapter 3 for more information.

Perform the following steps to install an Opsware Agent:

1. Log into the OCC Client. The OCC Client Page appears.
2. From the navigation pane, select Servers and then select Unmanaged Servers.
3. Select a location to scan for servers from the Scan in drop-down list.
4. Select Supply IP Address Ranges from the drop-down list to specify a set of IP address ranges to scan. Enter the IP address range in the From and To field. Click the plus (+) button (next to the To field) to add another IP address range. You can add a maximum of five IP address ranges. Click the minus (-) button to delete the IP address range field.

   Or

Select Explicit list of IPs from the drop down list to specify the list of IP addresses to scan. Click the ellipsis (...) button as shown in Table 6-5 to display a text editor that allows you to load and scan a file with IP ranges to scan.

Figure 6-5: Loading a File with IP Addresses

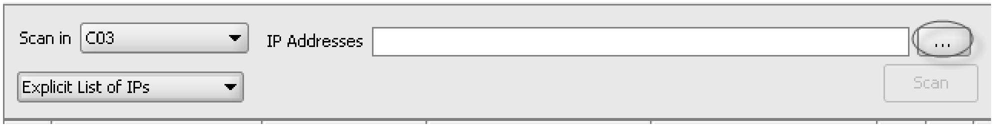

Scan in C03

Explicit List of IPs

IP Addresses

5 Click **Scan** to scan for servers. When the scan is complete, the list of scanned servers is shown. For each server, ODAD determines the status of the server, its IP address, its host name, the detected operating system, the actual operating system, and open ports that can be used to connect to the server. See Figure 6-6. The actual operating system of the server can be only determined if ODAD is able to successfully log into the server.

Figure 6-6: Scan Results

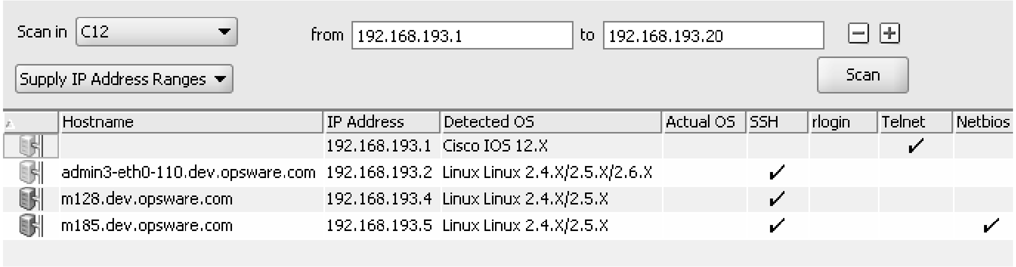

Scan in C12

Supply IP Address Ranges

from 192.168.193.1 to 192.168.193.20

	Hostname	IP Address	Detected OS	Actual OS	SSH	rlogin	Telnet	Netbios
		192.168.193.1	Cisco IOS 12.X				✓	
	admin3-eth0-110.dev.opsware.com	192.168.193.2	Linux Linux 2.4.X/2.5.X/2.6.X		✓			
	m128.dev.opsware.com	192.168.193.4	Linux Linux 2.4.X/2.5.X		✓			
	m185.dev.opsware.com	192.168.193.5	Linux Linux 2.4.X/2.5.X		✓			✓

6 Click on any of the column headings to sort the server list by that column. If you want to hide managed servers, select **Hide Opsware-managed Servers** from the **View** menu.

7 Select servers on which you want to deploy the Opsware Agent. The OCC Client supports hot keys to make multiple selections.

8 From the **Actions** menu, select **Deploy Agent**. The Deploy Agent dialog box appears as shown in Figure 6-7.

Figure 6-7: Deploy Agents

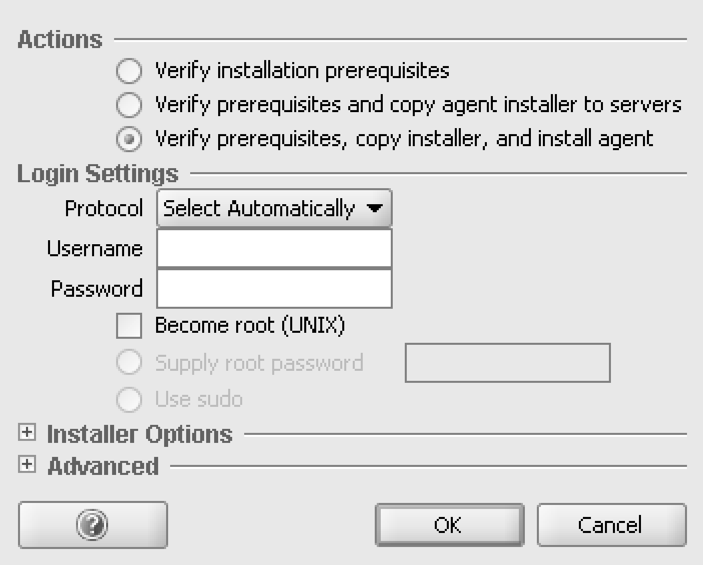

Actions

○ Verify installation prerequisites

○ Verify prerequisites and copy agent installer to servers

● Verify prerequisites, copy installer, and install agent

Login Settings

Protocol: Select Automatically

Username:

Password:

□ Become root (UNIX)

○ Supply root password

○ Use sudo

⊞ Installer Options

⊞ Advanced

9 Select one of the following deployment actions:

- Verify installation prerequisites.
- Verify prerequisites and copy agent installer to servers.
- Verify prerequisites, copy installer, and install agent.

See "Setting Deployment Actions for Each Server" on page 191 in this chapter for more information.

10 Select a network protocol to log in and connect to the server from the drop-down list.

Or

Choose Select Automatically to allow ODAD to select an appropriate protocol for each server.

11 Enter the user name to log into the server. For Windows, log in as administrator. For Unix, log in as root.

12 If root logging is not allowed, select the Become root (UNIX) checkbox. Select "Supply root password" and enter the password or select Use sudo**,** if sudo access is enabled for that account.

See “Specifying Login Settings” on page 191 in this chapter for more information.

13 Specify the Opsware Agent Installer options to control the way the Opsware Agent is installed on a server. See “Opsware Agent Installer Options” on page 192 in this chapter for more information.

14 Click **OK** to deploy Opsware Agents to selected servers.

15 After the deployment action is completed, the OCC Client displays the results and updates the status icons for the servers. You can view information on an unmanaged server and generate reports on Opsware Agent Installation status. See “Reports on Opsware Agent Installation” on page 202 in this chapter for more information.

Viewing Unmanaged Server Information

After the deployment action is completed, you can review the results and generate reports. You can view the summary and history information for an unmanaged server. For a failed deployment action, you can view the errors on each server. You can also log into the server from ODAD and correct the errors.

This section discusses the following topics:

- Summary Information for an Unmanaged Server
- History Information for an Unmanaged Server
- Remote Terminal Sessions on Unmanaged Servers
- Reports on Opsware Agent Installation

Summary Information for an Unmanaged Server

The Unmanaged Server Summary browser shows the following information about the unmanaged server:

- **Host Name**: The host name of the unmanaged server, if defined in the Domain Name System (DNS).
- **IP Address**: The IP address of the unmanaged server.
- **Detected OS**: The operating system detected on the server after performing the network scan.
- **Actual OS**: The actual operating system detected on the server after the deployment action.

- **MAC Address**: The Media Access Control (MAC) address, which is the network interface card's unique hardware number of the unmanaged server. The MAC is used as the server's physical address on the network. The MAC address is only detected if the server is on the same physical network as the Opsware Gateway.
- **NIC Vendor**: The vendor name for the Network Interface Card (NIC) driver. The NIC vendor is only detected if the server is on the same physical network as the Opsware Gateway.
- **Open Ports**: The available open ports on an unmanaged server.
- **Number of Deployment Attempts**: The number of deployment attempts on an unmanaged server.
- **Last Deployment Attempt Date/Time)**: The Date/Time of the last deployment attempt.
- **Last Deployment Attempt Message**: The message stating the possible cause for the failure of the last deployment attempt.

Viewing Summary Information of an Unmanaged Server

Perform the following steps to view the summary of an unmanaged server.

1. Log into the OCC Client. The OCC Client Page appears.
2. From the navigation pane, select Servers and then select Unmanaged Servers.
3. From the Unmanaged Servers page, select the unmanaged server. From the **Actions** menu, select **Open**.

   The Unmanaged Server Browser page appears.

4. Click Summary on the left side navigation panel of the Unmanaged Server Browser page to view the summary on the unmanaged server. See Figure 6-8.

Figure 6-8: Unmanaged Server Summary Page

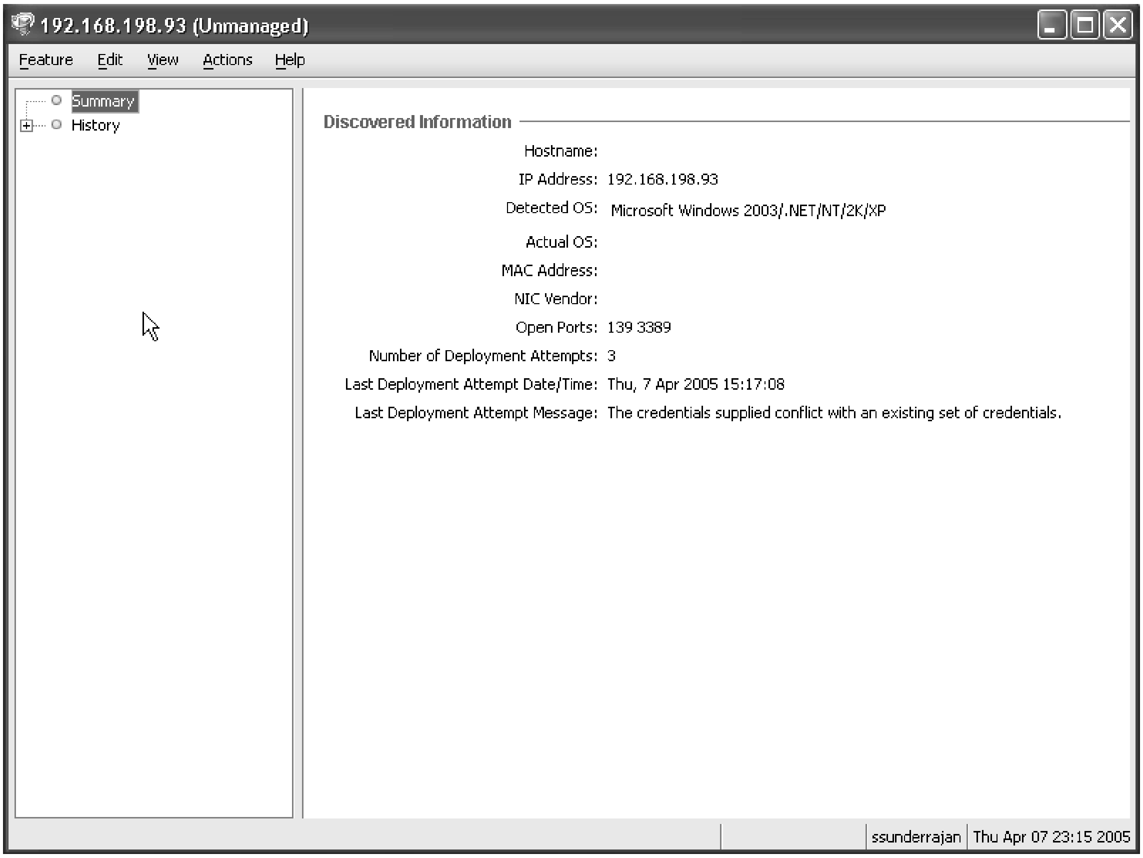

Discovered Information

Hostname:
IP Address: 192.168.198.93
Detected OS: Microsoft Windows 2003/.NET/NT/2K/XP
Actual OS:
MAC Address:
NIC Vendor:
Open Ports: 139 3389
Number of Deployment Attempts: 3
Last Deployment Attempt Date/Time: Thu, 7 Apr 2005 15:17:08
Last Deployment Attempt Message: The credentials supplied conflict with an existing set of credentials.

ssunderrajan | Thu Apr 07 23:15 2005

History Information for an Unmanaged Server

The Unmanaged Server History Browser shows a history of all actions executed on unmanaged servers such as:

- A summary of the actions performed on the unmanaged server
- Details of all the actions performed on the unmanaged server
- Log information

The history is read-only.

Viewing History Information of an Unmanaged Server

Perform the following steps to view the history of an unmanaged server:

1. Log into the OCC Client. The OCC Client Page appears.
2. From the navigation pane, select Servers and then select Unmanaged Servers.
3. From the Unmanaged Servers page, select the unmanaged server. From the **Actions** menu, select **Open**.

   The Unmanaged Server Browser page appears.
4. Click History on the left side navigation panel of the Unmanaged Server Browser page to view the history of all actions executed on unmanaged servers.

Remote Terminal Sessions on Unmanaged Servers

Using the ODAD feature, you can open terminal sessions on an unmanaged server. You can log into the server using the appropriate protocol to correct the errors on the server or perform any other operations.

You can use the OCC Client Remote Terminal and Shell preferences to configure the Terminal (UNIX) and RDP (Windows) clients used to launch terminal sessions. See "Overview of Global Shell" on page 385 in Chapter 10 for more information.

Opening Remote Terminal Sessions on an Unmanaged Servers

Perform the following steps to open remote terminal sessions on a server:

1. Log into the OCC Client. The OCC Client Page appears.
2. From the navigation pane, select Servers and then select Unmanaged Servers.
3. From the Unmanaged Servers page, select the unmanaged server. From the **Actions** menu, select the appropriate log in with protocol.

Reports on Opsware Agent Installation

Using the Opsware Discovery and Agent Deployment feature, you can create the following reports:

- All the servers in the current network scan
- Selected servers in the current network scan
- Servers in the current network scan with successful deployments
- Servers in the current network scan with failed deployments

You can also save and export the reports to a CSV, HTML, or text file format.

Creating Reports on Opsware Agent Installation

Perform the following steps to create reports:

1. Log into the OCC Client. The OCC Client Page appears.

2. From the navigation pane, select Servers and then select Unmanaged Servers.

3. From the Unmanaged Servers page, select the unmanaged server. From the **Actions** menu, select **Export to** the desired report format. The Save Report dialog box appears.

4. From the drop-down list, select the type of report as shown in Figure 6-9.

Figure 6-9: Generating Reports

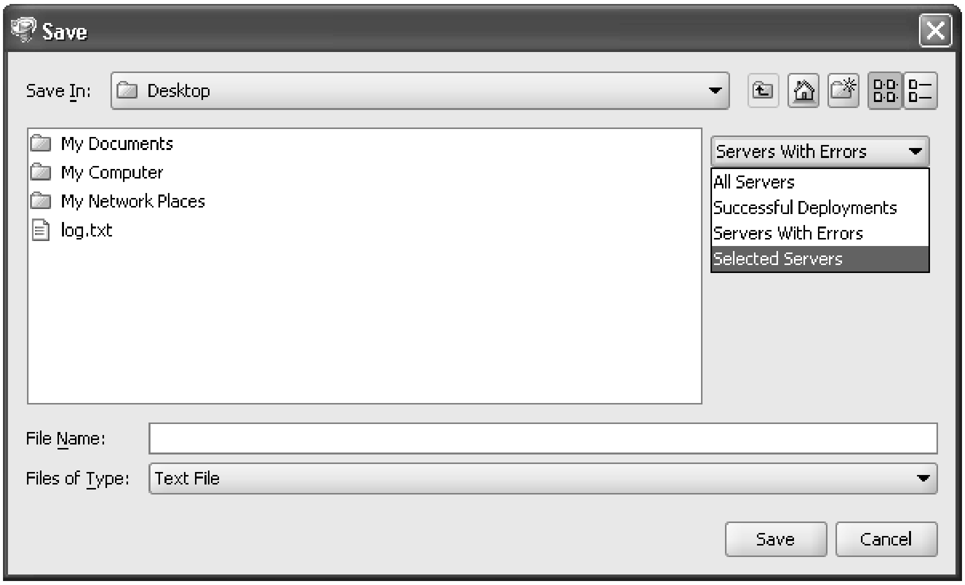

Save In: Desktop

Servers With Errors

File Name:

Files of Type: Text File

5. Enter the location and file name to save the report.

Example Report

The following example shows a report of servers with failed deployments.

```
Servers which the Opsware Agent could not be deployed to:

Hostname : Unknown

IP Address : 192.168.198.93
```

```
Detected Operating System : Microsoft Windows 2003 Server or XP
SP2

Open Ports : 139 3389

MAC Address : Unknown

NIC Vendor : Unknown

# of Deployment Attempts : 2

Last Deployment Message : The credentials supplied conflict with
an existing set of credentials.

Last Deployment Attempt Date/Time : "Wed, 6 Apr 2005 16:18:46"

Failed Phase : Check

Return Code : 2013

Suggested Resolution : The Agent Deployment Helper was unable to
log in to the unmanaged server.

An incorrect login name or password was specified. Try a
different login name and/or password.
```

Agent Reachability Communication Tests

This section provides information on agent reachability Communication Tests within Opsware SAS and contains the following topics:

- Overview of Communication Tests
- Communication Tests and Unreachable Opsware Agents
- Communication Test Types
- Communication Test Errors
- Additional Information on Communication Tests
- Running a Communication Test on an Individual Server
- Running a Communication Test on Multiple Servers
- Viewing Servers by Communication Status
- Searching for Unreachable Servers

- Creating Communication Test DCI Reports
- Viewing My Jobs Communication Tests
- Exporting the Unreachable Server Status List to CSV

Overview of Communication Tests

Sometimes an Opsware Agent can become unreachable, which means that the Opsware Command Center has difficulty communicating with the Opsware Agent. When an Opsware Agent is unreachable, the server it is installed on is considered unmanaged. This section explains how to use the Communication Test to find unreachable Opsware Agents and suggests ways that you can resolve these problems.

To help identify those managed servers that have unreachable agents, the Opsware Command Center runs periodic Communication Tests to verify that Opsware SAS can communicate with all servers under its management. You can always check the reachability of Opsware Agents by looking at the server's properties, or by viewing the current agent reachability status for all managed servers since the last Communication Test was run by choosing the Communication view from the Manage Servers list.

To determine the current reachability of a specific Opsware Agent, you can run a Communication Test to find those servers that have unreachable agents by using the Communication Test feature located in the Server menu of the Manage Server list. A Communication Test lists all servers with unreachable agents, returns specific errors associated with each unreachable Opsware Agent, and provides troubleshooting information to help you get the Opsware Agent back in working order.

You have the ability to check Opsware Agent reachability for individual servers, selected servers, or all servers under management of Opsware SAS. Each time that you run a Communication Test, this test is saved in the My Jobs panel, which allows you to view a history of all the Communication Tests that you have run. You can even export the current reachability status of all managed servers to a CSV file.

Communication Tests and Unreachable Opsware Agents

The Communication Test works by testing communication and data exchange between the specific components of the Opsware core and each managed server. The Opsware core is the entire collection of servers and services that provide Opsware services. In order to successfully manage servers, the Opsware core needs to be able to communicate with each Opsware Agent on all servers under Opsware SAS management.

Communication Test Types

The Communication Test performs the following diagnostics to determine if an Opsware Agent is reachable:

- **Command Engine to Agent Communication (AGT)**: Determines if the Command Engine can communicate with the agent. The Command Engine is the Opsware SAS component that enables distributed programs to run across many servers. The Command Engine handles the entry of scripts into the Opsware Model Repository (the script storage location in Opsware SAS) and the versioning of stored scripts.
- **Crypto Match (CRP)**: Checks that the SSL cryptographic files that the agent uses are valid.
- **Agent to Command Engine Communication (CE)**: Verifies that the agent can connect to the Command Engine and retrieve a command for execution.
- **Agent to Data Access Engine (DAE)**: Checks whether or not the agent can connect to the Data Access Engine and retrieve its device record. The Data Access Engine provides an XML-RPC interface to the Model Repository that simplifies interaction with various clients such as the Opsware Command Center, system data collection, and monitoring agents on servers.
- **Agent to Software Repository Communication (SWR)**: Determines if the agent can establish an SSL connection to the Software Repository. The Software Repository is the central repository for all software that Opsware SAS manages. It contains software packages for operating systems, applications, databases, customer code, and software configuration information.
- **Machine ID Mismatch (MID)**: Checks that the Machine ID (MID) on the server matches the MID registered in the Model Repository for the agent.

When the test finishes, it returns results that show either success or failure for each test run on each server. For each failed test, the nature of the failure is indicated in the Test Summary column of the Communication Test results page. In some cases, the failure of one test might prevent other tests from being executed.

See "Opsware Agent Functionality on Managed Servers" on page 184 in this chapter for information about the Opsware Agent and its relationship to managed servers.

Communication Test Errors

After you run a Communication Test, three icons indicate the success or failure of agent reachability as Table 6-4 shows.

Table 6-4: Agent Unreachability Status Icons

STATUS ICON	DESCRIPTION
●	Communication Test passed. Agent is reachable.
▭	Communication Test unable to be executed.
☒	Communication Test failed. Agent is unreachable.

Table 6-5 describes each type of Communication Test and all possible errors for each test.

Table 6-5: Communication Test Types with Possible Results

TEST	DESCRIPTION	RESULTS
Command Engine to Agent Communication (AGT)	Determines if the Command Engine can communicate with the agent	1 OK 2 Untested 3 Unexpected error 4 Connection refused 5 Connection time-out 6 Request time-out 7 Server never registered 8 Realm is unreachable 9 Tunnel setup error 10 Gateway denied access 11 Internal gateway error 12 Gateway could not connect to server 13 Gateway time-out
Crypto Match (CRP)	Checks that the agent's SSL cryptographic files are valid	1 OK 2 Untested 3 Unexpected error 4 Agent certificate mismatch 5 SSL negotiation failure

Table 6-5: Communication Test Types with Possible Results (continued)

TEST	DESCRIPTION	RESULTS
Agent to Command Engine Communication (CE)	Verifies that the agent can connect to the Command Engine and retrieve a command for execution	1 OK 2 Untested 3 Unexpected error 4 Connection refused 5 Connection time-out 6 DNS does not resolve 7 Old agent version 8 Realm is unreachable 9 No gateway defined 10 Tunnel setup error 11 Gateway denied access 12 Gateway name resolution error 13 Internal gateway error 14 Gateway could not connect to server 15 Gateway time-out 16 No callback from agent

Table 6-5: Communication Test Types with Possible Results (continued)

TEST	DESCRIPTION	RESULTS
Agent to Data Access Engine (DAE)	Checks whether or not the agent can connect to the Data Access Engine and retrieve its device record	1 OK 2 Untested 3 Unexpected error 4 Connection refused 5 Connection time-out 6 DNS does not resolve 7 Old agent version 8 Realm is unreachable 9 No gateway defined 10 Tunnel setup error 11 Gateway denied access 12 Gateway name resolution error 13 Internal gateway error 14 Gateway could not connect to server 15 Gateway time-out

Table 6-5: Communication Test Types with Possible Results (continued)

TEST	DESCRIPTION	RESULTS
Agent to Software Repository Communication (SWR)	Determines if the agent can establish an SSL connection to the Software Repository	1 OK 2 Untested 3 Unexpected error 4 Connection refused 5 Connection time-out 6 DNS does not resolve 7 Old agent version 8 Server identification error 9 Realm is unreachable 10 No gateway defined 11 Tunnel setup error 12 Gateway denied access 13 Gateway name resolution error 14 Internal gateway error 15 Gateway could not connect to server 16 Gateway time-out
MID Match	Checks that the Machine ID (MID) on the server matches the MID registered in the Model Repository for the agent	1 OK 2 Untested 3 Unexpected error 4 MID Mismatch

See “Communication Test Troubleshooting” on page 743 in Appendix C for information about how to troubleshoot Communication Test Errors.

Additional Information on Communication Tests

In the case that the Communication Test cannot be performed on a server, you see an error named Unexpected Error with a small plus (+) button next to it. Click the plus (+) button to see the Additional Information window, which provides traceback information regarding the error. You can send this information to Opsware Customer Support to solve a problem of this nature. See Figure 6-10.

Figure 6-10: Additional Information on an Unexpected Error

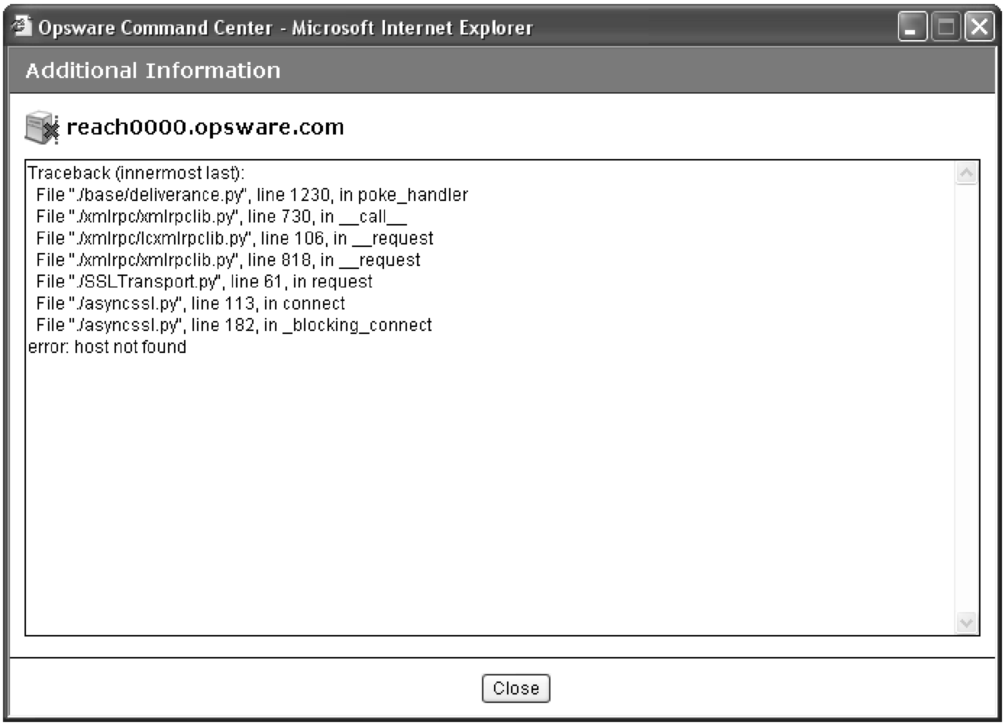

reach0000.opsware.com

```
Traceback (innermost last):
  File "./base/deliverance.py", line 1230, in poke_handler
  File "./xmlrpc/xmlrpclib.py", line 730, in __call__
  File "./xmlrpc/lcxmlrpclib.py", line 106, in __request
  File "./xmlrpc/xmlrpclib.py", line 818, in __request
  File "./SSLTransport.py", line 61, in request
  File "./asyncssl.py", line 113, in connect
  File "./asyncssl.py", line 182, in _blocking_connect
error: host not found
```

Running a Communication Test on an Individual Server

Perform the following steps to run a Communication Test on an individual server to find out if the Opsware Agent on that server is reachable:

1. From the navigation panel, click Servers ➤ Manage Servers.

2. From the Manage Servers list, click the display name of the server that you want to perform a Communication Test on.

   On the Server Property page, look in the Status field and notice that the server is either listed as Reachable or Not reachable. If listed as Not Reachable, a date indicates when the last regularly scheduled Communication Test was performed.

3 To see the results of the last Communication Test for this server, click **Details**. The Communication Test window for the server appears, as Figure 6-11 shows.

Figure 6-11: Communication Test Results on an Individual Server

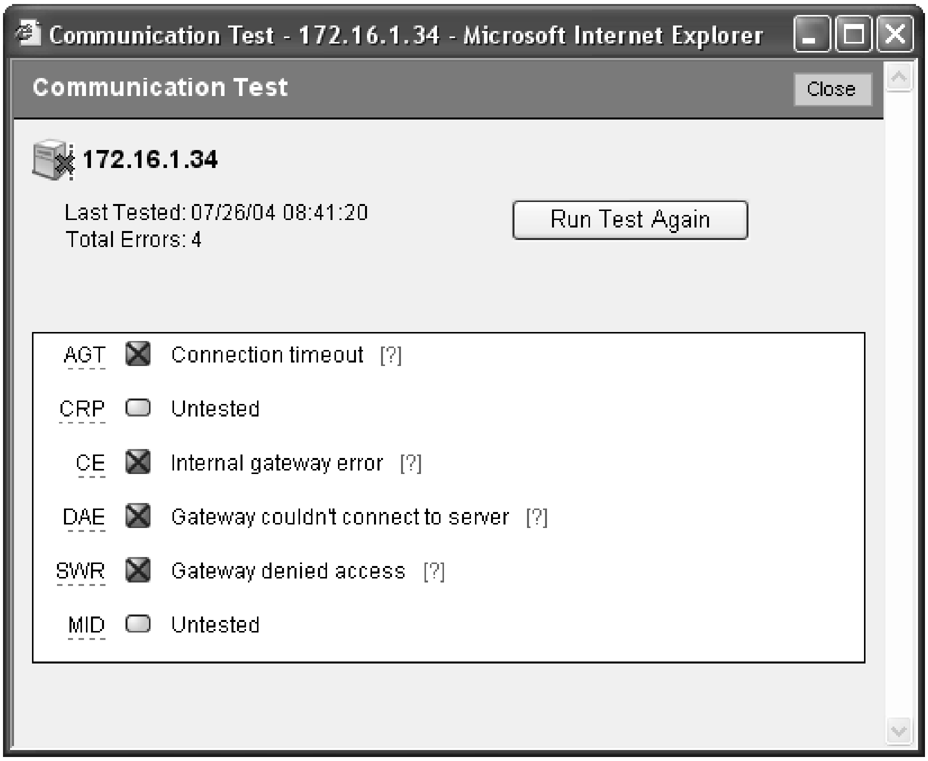

Communication Test - 172.16.1.34 - Microsoft Internet Explorer
Communication Test
Close
Run Test Again

172.16.1.34

Last Tested: 07/26/04 08:41:20
Total Errors: 4

AGT ☒ Connection timeout [?]
CRP ▭ Untested
CE ☒ Internal gateway error [?]
DAE ☒ Gateway couldn't connect to server [?]
SWR ☒ Gateway denied access [?]
MID ▭ Untested

The results listed in this window show details from when the last regularly scheduled Communication Test was run.

4 To view troubleshooting information for any of the test errors, move your mouse over the error name (for example SWR). When your mouse cursor changes to a question mark, click the question mark to view the troubleshooting help.

5 To rerun the Communication Test, click **Run Test Again**. The new results display in the same window when the test finishes.

Running a Communication Test on Multiple Servers

Perform the following steps to run a Communication Test on multiple servers and to find out which managed servers are not reachable:

1 From the navigation panel, click Servers ➤ Manage Servers.

2 In the Manage Servers Summary View page, select the servers for which you want a Communication Test.

3 Choose **Tasks ➤ Run ➤ Communication Test**. The Communication Test window opens and the test is initiated. The top of the test window shows the status report for the Communication Test, which indicates the time of the test, how many servers in the test were reachable, which servers were not reachable, and a progress bar. This summary information is shown in Figure 6-12.

Figure 6-12: Communication Test Summary

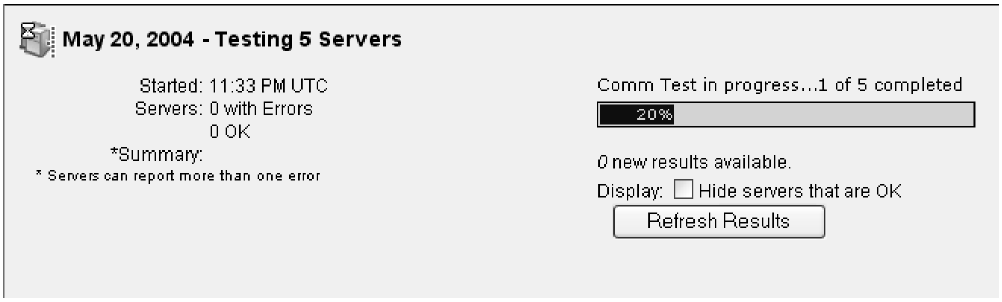

May 20, 2004 - Testing 5 Servers

Started: 11:33 PM UTC
Servers: 0 with Errors
0 OK
*Summary:
* Servers can report more than one error

Comm Test in progress...1 of 5 completed

0 new results available.

Display: □ Hide servers that are OK

- **Date**: Provides the date of the test.
- **Statistics**: Shows start and finish time, total servers OK, total servers with unreachable agents, and a summary of errors.
- **Progress Bar**: Provides live feedback of Communication Test progress. Progress data includes the number of servers completed, the total number of servers to be completed, and the list of servers completed so far.
- **Refresh Results Button**: Refreshes the results screen with new results.

Below the summary section is a list of all Opsware Agents that were not reachable and their details, as Figure 6-13 shows.

Figure 6-13: Communication Test Results on an Unreachable Agent

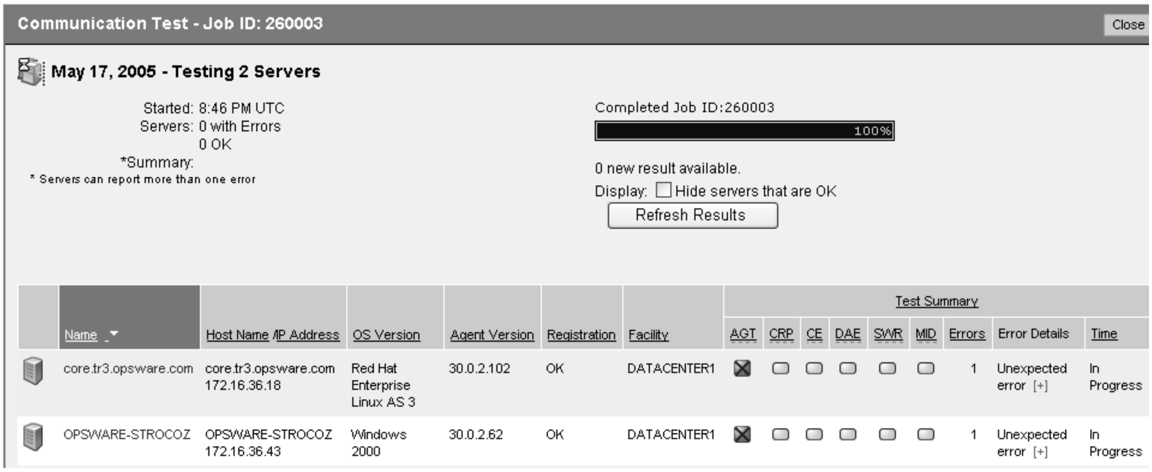

Communication Test - Job ID: 260003

May 17, 2005 - Testing 2 Servers

Started: 8:46 PM UTC
Servers: 0 with Errors
0 OK
*Summary:
* Servers can report more than one error

Completed Job ID:260003

0 new result available.

Display: □ Hide servers that are OK

Name	Host Name /IP Address	OS Version	Agent Version	Registration	Facility	Test Summary								
						AGT	CRP	CE	DAE	SWR	MID	Errors	Error Details	Time
core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	Red Hat Enterprise Linux AS 3	30.0.2.102	OK	DATACENTER1	☒	□	□	□	□	□	1	Unexpected error [+]	In Progress
OPSWARE-STROCOZ	OPSWARE-STROCOZ 172.16.36.43	Windows 2000	30.0.2.62	OK	DATACENTER1	☒	□	□	□	□	□	1	Unexpected error [+]	In Progress

- The results section shows a table of all unreachable servers, detailing server name, host name/IP address, OS, Opsware Agent version, registration (when the Opsware Agent last reported to Opsware SAS), and the time the test was completed.
- Click the title of each column to sort the test information by specific categories.
- The Test Summary section shows a list of all Communication Test types that were run, and which errors were returned. For information about each type of error and how to troubleshoot Opsware Agent reachability problems, click the link on each error to view online help.

Viewing Servers by Communication Status

Perform the following steps to view all manage servers with the most recent Communication Test results for each manage server:

1. From the navigation panel, click Servers ➤ Manage Servers.
2. From the **View** menu, choose **Communication**. A list of the most recent Communication Test that was run on all managed servers appears. All servers listed in this view are listed as unreachable or reachable since the last regularly scheduled Communication Test was run.

3. To sort the information in this view, click any of the column headings. For example, you might want to view the servers by number of errors. Click the Error column title by name, by OS version, and so on.

4. To export this view to the CSV file format, choose **Export to CSV** from the **Resource** menu.

5. To run a new Communication Test for some or all of the servers in this view, select the servers that you want to run a Communication Test on and then choose **Run ➤ Communication Test** from the **Tasks** menu.

Searching for Unreachable Servers

Another way you can discover unmanaged servers (those with unreachable agents) is to search for all servers that have a status of unreachable. See "Using the Search Feature" on page 125 in Chapter 4 for more information.

Perform the following steps to search for unreachable agents:

1. From the navigation panel, click Servers ➤ Server Search. The Server Search page appears.

2 From the Server Search page, choose the Agent Reporting attribute from the first list, as Figure 6-14 shows.

Figure 6-14: Agent Reporting Server Search Attribute

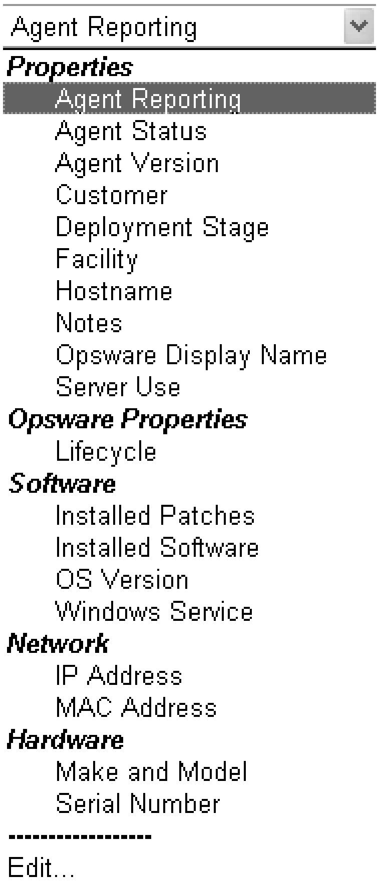

More criteria displays for the search parameters.

3 From the far right list of search attributes, select Not Reporting. Your search parameters are Agent Status is Not Reporting.

4 Click **Search**. Wait a few moments for the results (the speed of the search results depends on how many managed servers are being searched). If any of your managed servers are not reachable, these servers will be listed in the search results.

5 To run a new Communication Test on these servers, select the server (check box next to the server), and choose **Run ➤ Communication Test** from the **Tasks** menu.

Creating Communication Test DCI Reports

To view a Communication Test in a report format, you can use the Data Center Intelligence (DCI) Reporting feature. It enables you to create a printable report of the Communication Test results. After you make a report of the test results, you can export the report to HTML, CSV, Microsoft Word, and other file formats, so that you can exchange the report with others.

For information on how to create a DCI Communication Test report, view the DCI online help by clicking Help in the Opsware Command Center.

Viewing My Jobs Communication Tests

Each time that you run a Communication Test, the information is saved as a My Job. This feature automatically saves a history of all tests that you run. To view saved Communication Tests, perform the following steps:

1. From the navigation panel, click My Job.
2. From the My Job list, click the Communication Test job that you want to view.
3. In the Communication Test window, wait a few moments for the Communication Test information to load, then click **View Details**. You see the Communication Test window.

Exporting the Unreachable Server Status List to CSV

Perform the following steps to export the list of all servers that have a status of unreachable to the CSV file format:

1. From the navigation panel, click Servers ➤ Manage Servers.
2. From the **View** menu, choose **Communication**. You see a list of all servers that are in an unreachable or reachable state.

To export a list of these servers to the CSV file format, select the check box next to each server that you want to include in the report, then from the **Resource** menu choose **Export to CSV**.

Chapter 7: Server Management in Opsware Command Center

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Server Management
- Communication Between Managed Servers and Opsware SAS
- Server Groups
- Server Life Cycle
- Server Management Jobs
- Custom Fields for Servers
- Custom Attributes for Servers
- Service Levels
- Network Configuration

This section does not document how to install operating systems, patches, or applications on servers. However, it does discuss how those tasks fit into the overall server life cycle, and it does provide cross-references to the appropriate topics in other sections.

Overview of Server Management

Opsware SAS manages servers in an operational environment in the following ways:

- Installing Opsware Agents on servers that are already functioning in the operational environment, which allows users to deploy and manage new applications installed on those servers
- Provisioning servers with Microsoft Windows, Red Hat Linux, and Sun Solaris operating systems by using vendor-provided operating system bootstrapping technologies. Additionally, Opsware SAS integrates with AIX NIM and HP-UX Ignite installation

technologies to provide a uniform method for OS provisioning across a heterogeneous environment.

See “Overview of OS Provisioning” on page 569 in Chapter 15 for information about how Opsware SAS provisions Microsoft Windows, Red Hat Linux, and Sun Solaris operating systems on managed servers.

See the *Opsware® SAS Configuration Guide* for information about how Opsware SAS integrates with AIX NIM and HP-UX Ignite installation technologies.

- Automating patch management on Opsware SAS-managed servers by providing the ability to react quickly to newly-discovered security threats and the degree of control required to ensure that new patches are installed in a uniform way.

  See “Patch Management for Windows” on page 453 in Chapter 12 for more information.

  See “Patch Management for Unix” on page 519 in Chapter 13 for more information.

- Managing application provisioning, which enables users to deploy applications and databases across many servers simultaneously, and track what is deployed on each server.

  See “Software Provisioning” on page 549 in Chapter 14 for more information.

- Providing configuration tracking, which allows users to monitor selected configuration files and databases and to take certain actions when change is detected.

  See “Configuration Tracking” on page 643 in Chapter 17 for more information.

To manage servers with Opsware SAS, you do not need root access on Unix or administrator access on Windows. However, you will need permissions to use specific Opsware SAS features, as well as permissions for customers and facilities associated with servers. To obtain these permissions, contact your Opsware administrator. For more information, see the Permissions Reference appendix in the *Opsware® SAS Configuration Guide*.

Communication Between Managed Servers and Opsware SAS

This section provides information about communication between managed servers and Opsware SAS and contains the following topics:

- Network Address Translation (NAT) for Managed Servers
- Key Terms of Opsware SAS Managed Server Communication
- Locating the Management IP Address of a Managed Server
- Code Deployment and Static NAT
- Setting the Primary IP Address of a Server
- NAT Table Mapping and Managed Servers

Network Address Translation (NAT) for Managed Servers

To manage a server, Opsware SAS requires that the server have a unique IP address that is routable from Opsware SAS. However, in a large operational environment, all servers might not have unique IP addresses. In this case, Opsware SAS supports static Network Address Translation (NAT) for managed servers.

Static NAT maps public IP addresses to hosts inside the internal network, which allows Opsware SAS to manage all servers in the environment.

Unlike dynamic NAT, the mapping between Opsware SAS and the servers under management is set ahead of time, not dynamically at runtime.

Key Terms of Opsware SAS Managed Server Communication

To understand how Opsware SAS communicates with managed servers, you must understand these three terms:

- **Management IP**: The IP address that Opsware SAS uses to communicate with the Opsware Agent on the server.

  During hardware registration for a server, the Opsware Agent opens a TCP/IP connection to Opsware SAS. The connection contains the source IP (called peer IP) address of the server. By default, Opsware SAS uses this peer IP address as the management IP for the server.

- **Management Interface**: When a server has more than one network interface, you can designate one of them as the management interface.

- **Primary IP**: The IP address of the management interface. When you change the management interface, the primary IP changes to the IP address of that interface. The primary IP address is a locally-configured IP address.

  During synchronizations, the Code Deployment & Rollback feature uses the primary IP address to communicate with the server. See "Code Deployment and Static NAT" on page 224 in this chapter for more information.

  The Opsware Agents on servers communicate with each other by using the primary IP addresses, even though Opsware SAS uses management IP addresses to communicate with the servers.

Opsware SAS does not support managed servers that have IPv6 addresses.

When static NAT is being used, the management IP address for a server will *not* be the same as the primary IP address. When static NAT is being used, the management IP is the NAT-translated IP address for the server. When static NAT is *not* being used, the management IP address is always the same as the primary IP address.

Locating the Management IP Address of a Managed Server

In the Opsware Command Center, you can find the management IP address of a server and check whether it is using static NAT. You might need this information for troubleshooting any servers marked Not Reachable and to determine whether your NAT configuration is correct. The Opsware Command Center displays the management IP address of a managed server in the following two places:

- The Network tab of the Server Details page
- The Hardware view of the Manage Servers list

The Network tab shows (and allows the user to set) the management interface for the server by selecting it from a drop-down list, as Figure 7-1 shows.

Figure 7-1: Management IP Address Information in the Network Tab

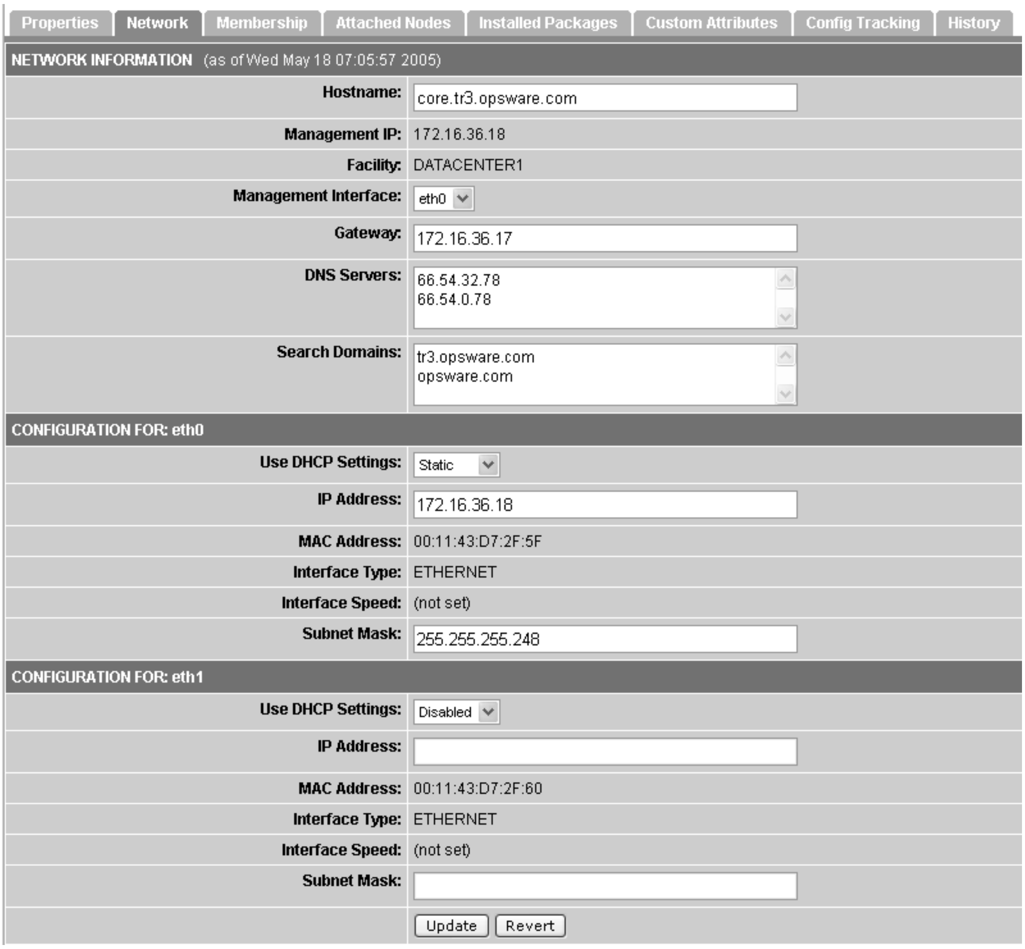

NETWORK INFORMATION (as of Wed May 18 07:05:57 2005)	
Hostname:	core.tr3.opsware.com
Management IP:	172.16.36.18
Facility:	DATACENTER1
Management Interface:	eth0
Gateway:	172.16.36.17
DNS Servers:	66.54.32.78 66.54.0.78
Search Domains:	tr3.opsware.com opsware.com
CONFIGURATION FOR: eth0	
Use DHCP Settings:	Static
IP Address:	172.16.36.18
MAC Address:	00:11:43:D7:2F:5F
Interface Type:	ETHERNET
Interface Speed:	(not set)
Subnet Mask:	255.255.255.248
CONFIGURATION FOR: eth1	
Use DHCP Settings:	Disabled
IP Address:	
MAC Address:	00:11:43:D7:2F:60
Interface Type:	ETHERNET
Interface Speed:	(not set)
Subnet Mask:	

Figure 7-2 shows the Hardware view in the Manage Servers list, which displays the management interface for the server in the Network Info column. (To access the Hardware view, choose **Hardware** from the **View** menu.)

Figure 7-2: Hardware Tab in the Manage Servers List

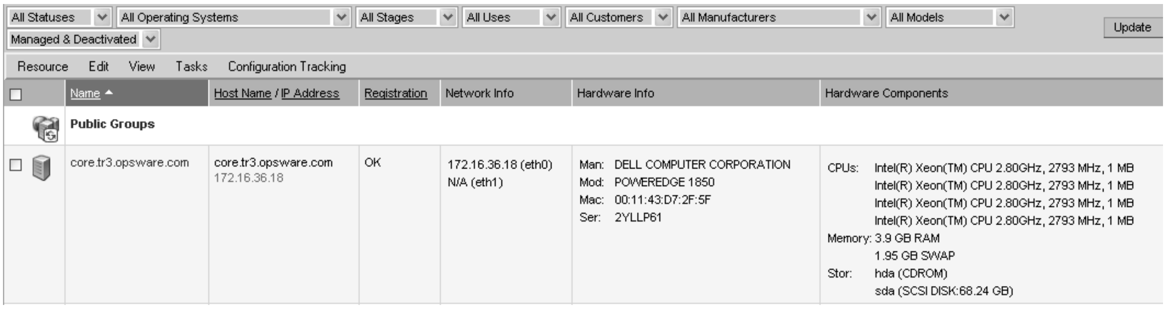

□	Name ▲	Host Name / IP Address	Registration	Network Info	Hardware Info	Hardware Components
	Public Groups					
□	core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	OK	172.16.36.18 (eth0) N/A (eth1)	Man: DELL COMPUTER CORPORATION Mod: POWEREDGE 1850 Mac: 00:11:43:D7:2F:5F Ser: 2YLLP61	CPUs: Intel(R) Xeon(TM) CPU 2.80GHz, 2793 MHz, 1 MB Intel(R) Xeon(TM) CPU 2.80GHz, 2793 MHz, 1 MB Intel(R) Xeon(TM) CPU 2.80GHz, 2793 MHz, 1 MB Intel(R) Xeon(TM) CPU 2.80GHz, 2793 MHz, 1 MB Memory: 3.9 GB RAM 1.95 GB SWAP Stor: hda (CDROM) sda (SCSI DISK:68.24 GB)

The Network Info column shows the IP addresses and interfaces configured for each server in the list. If a server is using static NAT, the management IP is the first entry in this list and (NAT) appears after the IP address. If it is not using static NAT, the management IP is the same as the management interface, so it is already shown.

Code Deployment and Static NAT

Code Deployment and Rollback (CDR) synchronizations can only occur between Opsware Agents in the same NAT domain. Synchronizations cannot be performed between Opsware Agents in different NAT domains.

Opsware SAS uses the primary IP address (instead of the management IP address) during synchronization because it is assumed that static NAT is not occurring between the servers in the synchronization. During CDR synchronizations, two Opsware Agents must communicate directly. Users can override the IP address that Opsware SAS determined and designate a specific network interface as the management interface.

See "Code Deployment and Rollback" on page 671 in Chapter 18 for information about how to use CDR.

Setting the Primary IP Address of a Server

When a server has more than one network interface, users can specify one of them as the management interface and the IP address for this interface is designated the primary IP address. The primary IP address is used for Opsware Agent-to-Opsware Agent communication.

If static NAT is *not* being used, the management and primary IP addresses are the same. If static NAT is being used, the management IP is unaffected when a user changes the management interface.

Perform the following steps to set the primary IP address of a server:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server whose management IP address you want to view.

   Or

   Search for the server whose management IP address you want to view.

2. Click the server name. The Manage Servers: Properties page appears.

3. Select the Network tab. The network information for the server appears.

   The Network tab shows (and allows you to set) the server’s management interface.

4. Set the management interface by selecting it from the Management Interface field. The IP address for this interface is designated the primary IP address.

5. Click **Update**.

See “Using the Search Feature” on page 125 in Chapter 4 for more information. See “Server Searching by IP Address” on page 141 in Chapter 4 for more information.

NAT Table Mapping and Managed Servers

Static, one-to-one NAT tables map routable IP addresses between Opsware SAS and managed servers. Network administrators configure and maintain these NAT tables. After the static NAT tables are configured, you do not have to perform any additional setup for Opsware SAS.

Opsware SAS does not control these NAT tables and errors can occur if they are modified after Opsware-managed servers register their hardware information. The following errors can occur if the IP address mapping of a server changes:

- If the IP address on the Opsware SAS side of the NAT mapping is modified, the server becomes unmanageable and might be marked Not Reachable on the Manage Servers: Status page. It stays in this state until the Opsware Agent requests another hardware registration and the server's management IP is updated.
- If an IP address mapped to a particular server is mapped to a different server, both servers become unmanageable and might be marked Not Reachable on the Manage

Servers: Status page. This problem is resolved when one of the two servers reports its IP address during hardware registration. The other server remains unmanageable until the server registers with the Opsware Agent. Both servers eventually become manageable.

Server Groups

This section provides information on server groups within Opsware SAS and discusses the following topics:

- Overview of Server Groups
- Types of Server Groups
- Public Group Modeling
- Ways to Create Server Groups
- Adding a Server to Static Groups, Both Public and Private
- Removing Servers from Static Groups, Both Public and Private
- Duplicating Server Groups
- Rules for Deleting Server Groups

Overview of Server Groups

The Server Groups feature is useful for gathering servers into collections. These groups can be used as a shortcut for performing the same action on all of the servers simultaneously, instead of performing the action on each individual server, one at a time. Server groups can also be used to simply organize groups of servers.

Server groups can be comprised of individual servers as well as other server groups.

The My Servers feature can also be used to gather servers and server groups, but it has different functionality than the Server Groups feature. You can add individual servers, server groups, and nested groups that you access frequently to My Servers. See "My Servers" on page 122 in Chapter 4 for information about using this feature.

Uses for Server Groups

Some recommended uses for server groups include:

- Grouping servers by OS version
- Grouping servers by customer
- Grouping servers by facility
- Grouping servers by deployment stage
- Grouping servers by use (Server Use in the Server Properties page)
- Grouping servers by operational boundaries, for example, grouping together all servers that require identical application configuration
- Grouping servers by access control boundaries, for example, creating server groups that are associated with a specific User Group

Permissions Required for Working with Server Groups

Users must have the permissions shown in Table 7-1 in order to perform specific tasks related to server groups. Only administrators can set permissions.

Table 7-1: Permissions Required for Working with Server Groups

NAME OF PERMISSION	WHERE SELECTED	ENABLES YOU TO
Manage Servers		Create, edit, and delete private server groups, both static and dynamic.
Manage Public Server Groups	The Manage Servers Permissions section on the Other tab in Users and Groups	Create, edit, and delete public server groups, both static and dynamic.
Model Public Server Groups	The Manage Servers Permissions section on the Other tab in Users and Groups	Model public server groups. See “Public Group Modeling” on page 230 in this chapter for more information.

Table 7-1: Permissions Required for Working with Server Groups (continued)

NAME OF PERMISSION	WHERE SELECTED	ENABLES YOU TO
Allow Run Refresh Jobs	The Manage Servers Permissions section on the Other tab in Users and Groups	Add or remove servers from groups before a scheduled job is run. This permission gives you the option to refresh group membership before a job is run so that it is only run on the servers that belong to the group when the job is actually run.

Characteristics of Server Groups

When using server groups, groups have the following characteristics:

- Individual servers can be included in as many groups as you want, or not included in any groups.
- Adding servers to a group does not remove those servers from the list of all servers that appears when you click Servers ➤ Manage Servers in the navigation panel.
- Groups can contain servers and subgroups.
- Server groups are hierarchical (they can be nested) with these caveats:
    - Private and public groups cannot be mixed in a hierarchy, but static and dynamic groups can.

      See "Types of Server Groups" on page 229 in this chapter for more information about private and public groups.
    - The rules for a dynamic group are not inherited from a parent dynamic group to a dynamic subgroup.

      See "Dynamic Groups" on page 230 in this chapter for more information about the characteristics of dynamic groups.
    - Groups do not inherit modeling data from their parents, including node attachments and custom attributes.
    - When you run an operation on a group that contains nested groups, the operation also applies to all the servers in the nested groups below the current group.

- When an Application Configuration operation within the OCC Client is applied to groups, and those groups contain subgroups, the operation does not apply to all the servers in the subgroups. It only applies to the group upon which the operation was directly applied.

Types of Server Groups

There are private groups and public groups, and each can be either static or dynamic.

Private Groups

If you belong to a user group that has access to the Manage Servers list, you can create groups that you alone can see and work with. Only you see your private groups – other Opsware users cannot see them. Private groups behave the same way as public groups, with the exception that modeling is not available for private groups. See “Public Group Modeling” on page 230 in Chapter 7 for information about how the Opsware model affects groups.

When you create your first group, the default type will be Private Static, which can be changed to Private Dynamic, Public Static, or Public Dynamic. When you create a subgroup, the type of group is private if you are in a private group when you create the new group, and the type is public if you are in a public group when you create the new group. Public and private groups cannot be mixed in a hierarchy. In other words, if the parent group is public, the subgroups must be public, and if the parent group is private, the subgroups are also private.

Public Groups

Public groups can be created, edited, or deleted by anybody who has Manage Public Server Groups permissions. Public groups are visible to all users, and can be used by anybody, regardless of who created them, but only users with the Manage Public Server Groups permissions can change the rules that govern dynamic groups.

A link called Public Groups appears at the top of server lists. Clicking that link displays a list of available public groups.

Only public groups can be used for modeling.

Static Groups

Static groups can be either public or private, and no specific permissions are required for static groups. A static group has servers that are added to and removed from the group manually. When using static groups, you first create the group, and then select the servers to populate it.

Dynamic Groups

Dynamic groups contain servers that are added to or removed from the group based on a set of user-defined rules. If the rules are changed or the servers in the environment change, servers will be added to or removed from the group automatically. Rules apply only to the group being created or modified, not to any subgroups.

Once the rules have been created, Opsware SAS will search for servers that match the criteria of that specific group, and add them to the group. When the rules are changed, Opsware SAS will search again, and the resulting group members reflect the changed criteria. Consequently, as servers are added to or removed from management using Opsware SAS, the members of the group will change automatically.

Opsware SAS calculates server group membership each time any of the following actions occur:

- After users add, delete, or change the rules for dynamic server groups.
- When attributes of servers change such that dynamic group membership could change.

Additionally, Opsware SAS automatically recalculates dynamic group membership every hour.

When a user schedules a job to run, dynamic group membership can be determined in either of the following ways:

- Based on the servers in the dynamic group when the job was scheduled.
- Based on the set of servers in the dynamic group when the job actually runs. The membership is recalculated at that time.

Public Group Modeling

With Opsware modeling, the desired state of a server is defined and then applied to servers. In the case of public server groups – static and dynamic – you can define a model consisting of applications, patches, service levels, and custom attributes, which will be applied to all servers in the group. The modeling information is attached to the group, but not to any subgroups.

If the modeling information changes, the servers in the group are not affected until the reconcile operation runs on those servers. If the model has already been reconciled on that server when it is removed from the group, the installed material will be removed at the next reconcile.

Manage Servers Display for Public and Private Groups

As Figure 7-3 shows, server groups appear at the top of the Manage Server list. Public Groups appear as a link at the top of the list. Private groups, if any have been created, appear next, followed by the list of individual servers.

Figure 7-3: Manage Servers Displaying Public and Private Groups — Servers and Group View

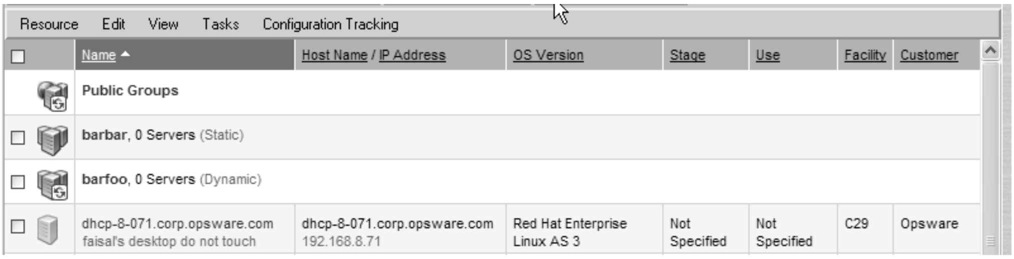

□	Name ▲	Host Name / IP Address	OS Version	Stage	Use	Facility	Customer
	Public Groups						
□	barbar, 0 Servers (Static)						
□	barfoo, 0 Servers (Dynamic)						
□	dhcp-8-071.corp.opsware.com faisal's desktop do not touch	dhcp-8-071.corp.opsware.com 192.168.8.71	Red Hat Enterprise Linux AS 3	Not Specified	Not Specified	C29	Opsware

Manage Servers List - Groups Only

When you select **View ➤ Groups Only** from the menu, information about the number of servers in a group, the number of groups in a group, and whether the group is static or dynamic appears in the list of groups, as Figure 7-4 shows.

Figure 7-4: Manage Servers List — Group Only View

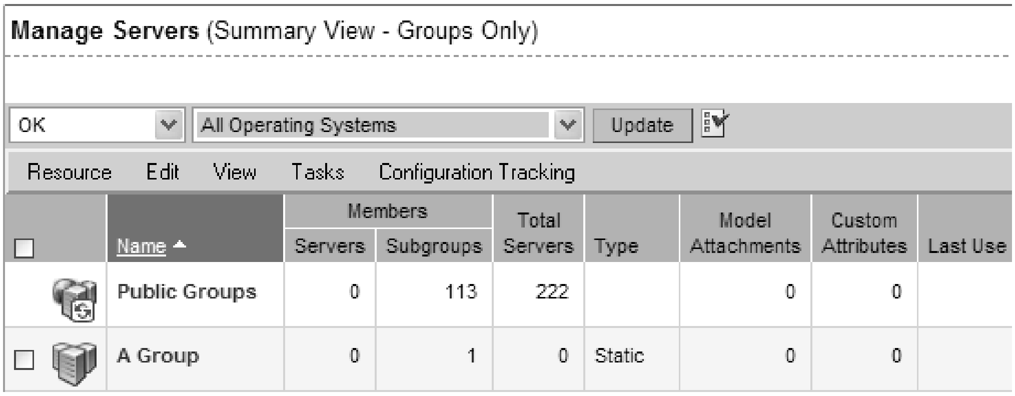
Manage Servers (Summary View - Groups Only)

		Members		Total		Model	Custom	
□	Name ▲	Servers	Subgroups	Servers	Type	Attachments	Attributes	Last Use
	Public Groups	0	113	222		0	0	
□	A Group	0	1	0	Static	0	0	

You can modify your views of servers and server groups by selecting Summary, Hardware, Software, and Communications from the **View** menu. You can also elect to further modify your view by selecting Servers and Groups, Servers Only, or Groups Only.

Operations on Server Groups

Any operation that can be done to a server can be done to a server group, because the group acts as a container for a collection of servers, and so provides a shortcut to avoid having to repeat the same operation on each individual server.

When any of the following operations are performed on a group, they are actually performed on the servers within the group, not on the group itself.

- Install
  - Application
  - Patch (Unix patch)
  - OS
  - By Template
- Uninstall
  - Application
  - Patch (Unix Patch)
- Run
  - Script
  - Control
  - Custom Extension
  - Comm Test
- Reconcile
- Customer
- Usage
- OCC Client operations
  - Configure Application
  - Audit Application Configurations
  - Perform Server Audit
  - Create Server Snapshot

Also, some operations allow users to refresh the list of servers in the group (if the user has the Allow Run Refresh Jobs permissions). Users with the correct permissions schedule a job, and before the job is run, Opsware SAS will update the members of the server groups upon which the operation is performed. The following actions allow refresh when scheduling a job:

- Install Application
- Install Patch
- Install By Template
- Run Script
- Run Custom Extension
- Reconcile

Using Server Groups Tabs

When you view group properties, the information for the group appears in the page, as Figure 7-5 shows.

You can view group properties in either of the following ways:

- By clicking the This Group link at the top of each server group list.
- By selecting a group in the list and choosing **Resource ➤ Properties** from the menu.

Figure 7-5: Tabs Available for Server Groups

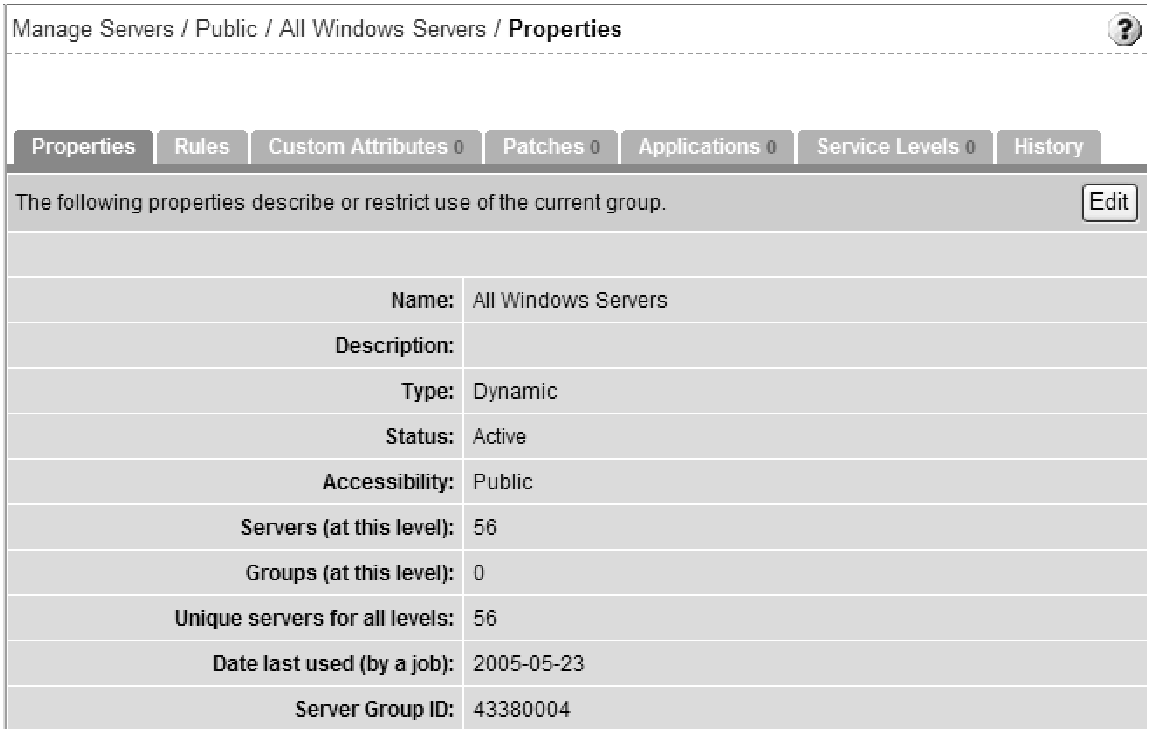

Manage Servers / Public / All Windows Servers / **Properties**

The following properties describe or restrict use of the current group.

Name:	All Windows Servers
Description:	
Type:	Dynamic
Status:	Active
Accessibility:	Public
Servers (at this level):	56
Groups (at this level):	0
Unique servers for all levels:	56
Date last used (by a job):	2005-05-23
Server Group ID:	43380004

You can use these tabs to perform the following actions for a server group.

To perform any of these actions, you must have the correct Opsware SAS permissions. See "Permissions Required for Working with Server Groups" on page 227 in this chapter for more information.

- **Properties**: This displays information about the group, such as the type of group, the number of servers in the group, and the status of the group. Clicking **Edit** allows you to change the group name and description, and to convert a dynamic group to a static group.

A created static group without servers added yet – an empty static group – can be converted to a dynamic group. Conversely, a static group that has servers cannot be converted to a dynamic group.

- **Rules**: This appears when you are viewing a dynamic group. The rules tab displays the rules used by the dynamic group to determine group membership. The rules apply to the current group only and do not apply to subgroups. Click **Edit** to change the rules for the group.

  See "Creating Groups by Using Search" on page 240 in this chapter for more information about specifying rules for a dynamic group.

- **Custom Attributes**: This allows you to set custom attributes for a server group. Click **New** to add an attribute and then click **Edit** to change existing attributes. Custom attributes are not inherited by subgroups within a group hierarchy.

  See "Custom Attributes for Servers" on page 282 in this chapter for more information about how custom attributes affect Opsware-managed servers.

- **Patches**: This allows you to add Unix patches to a server group. Attaching Unix patches to a server group is like adding a patch to an Opsware Template with one exception – groups are not associated with a specific OS or customer. Therefore, you can add a patch for any OS to the group and groups can contain servers running different operating systems. (Patches are always associated with the customer Customer Independent.) Templates and folders are always associated with one OS version and a specific customer.

- **Applications**: Allows you to add applications to a server group. Attaching applications to a server group is similar to adding an application to an Opsware Template with one exception – groups are not associated with a specific OS or customer; therefore, you can add an application for any OS or for any customer to the group.

The Reconcile Software wizard is flexible when you run it on a server group; for example, when you reconcile server groups, Opsware SAS will install the correct applications on the servers in the group, even if the OS versions for the servers and the patches or applications do not all match. Opsware SAS also matches the customer association for applications, operating systems, and templates with the customer association of servers. If Opsware SAS cannot find a match, an error message (no valid devices found) appears at the end of the wizard.

- **Service Levels**: This allows you to attach service levels to a group. Attaching service levels to a server is like attaching a service level to an Opsware Template with one exception – groups are not associated with a specific OS or customer; therefore, you can add a service level for any OS or for any customer to the group.

  See the Application Provisioning Setup chapter in the *Opsware® SAS Configuration Guide* for more information about attaching service levels.

- **History**: This allows you to view the changes to groups. For each action on the group (but not group members), the history displays a description, the date the action occurred, and the user who performed the action (if the group is public).

Ways to Create Server Groups

Server Groups can be created by:

- Using the New Group option from the Resource pull down menu on the Manage Servers page
- Clicking the Create Server Groups icon in the upper right corner of the Copy to Group dialog
- Performing a server search, and saving the resulting list of servers or the rules as a group

Creating Static Groups by Using the New Group Option

Static server groups require servers to be added to them manually. The servers in a static group also must be removed manually.

A created static group without servers added yet – an empty static group – can be converted to a dynamic group. Conversely, a static group that has servers cannot be converted to a dynamic group.

To create static server groups, perform the following steps:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears.

2. From the **Resource** menu, choose **New Group**. The New Private Static Group dialog appears, as shown in Figure 7-6.

Figure 7-6: Default New Private Static Group Dialog

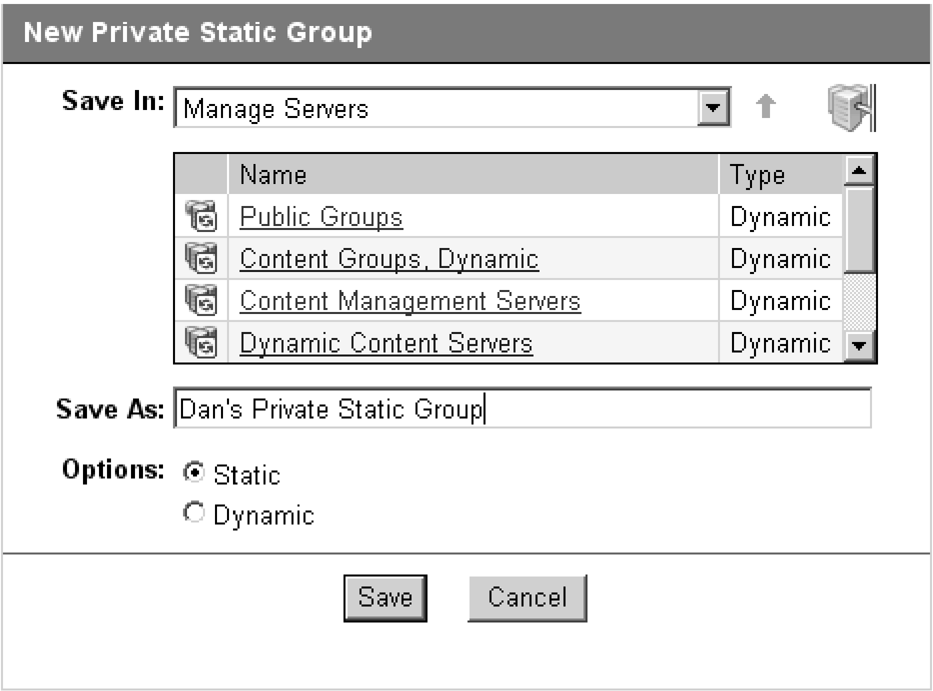

New Private Static Group

Save In: Manage Servers

Name	Type
Public Groups	Dynamic
Content Groups, Dynamic	Dynamic
Content Management Servers	Dynamic
Dynamic Content Servers	Dynamic

Save As: Dan's Private Static Group

Options: ● Static
○ Dynamic

3. To create a Private Static group in the top level of Manage Servers, perform the following steps:

   1. Enter the name of the group in the Save As text box.

   2. Click **Save**.

4. To create a Private Static group below another group, perform the following steps:

   1. Navigate to the group below which you want to create a new group.

2. Enter the name of the group in the Save As text box.
3. Click **Save**.

5 To create a Public Static Group, perform the following steps:

1. Click the Public Groups link in the Name and Type window, and navigate to the location in the group hierarchy where you want to create the group.
2. Enter the name of the group in the Save As text box.
3. Click **Save**.

The Save In drop down list is populated according to the location you drill down to in order to create your group. You can use the Save In drop down list to verify that you are in the correct location in the group hierarchy, and you can move to a different location in the hierarchy by selecting it from the Save In list.

Creating Static Groups by Using the Copy to Group Dialog

To create new server groups by copying existing server groups, perform the following actions:

1 From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears.

2 Navigate to the servers and groups that you want to copy and click the check boxes next to the servers and groups you want to copy.

3 From the **Edit** menu, choose **Copy to Group**. The Copy to Group dialog box appears.

4 Navigate to the place in the group hierarchy where you want to copy the servers and group then by clicking the group links displayed in the Name and Type field.

5 In the Options field, select either of the following options:

- **Maintain Hierarchy**: Select this option to copy any server groups exactly as they are.
- **Expand to a Flat List**: Select this option to copy only the servers within the group to the new group.

Each of those options displays the result of choosing that option by showing the numbers of servers and groups in the new group’s hierarchy, or the number of unique servers in the flat list.

If you select the Expand to a Flat List option, and a server is a member of more than one group, that server will only appear in the list once.

6. (Optional) You can also create an entirely new group for the copied servers and groups by clicking the Create New Group icon in the upper right corner of the dialog.

   A new dialog appears, prompting you to name your new group.

7. Enter the name of your newly-created group.
8. Click **OK**. The name dialog closes, and the Copy to Group dialog reappears.
9. Click **OK** on the Copy to Group dialog.

Creating Dynamic Groups by Using the New Group Option

Dynamic Server Groups are rule-based, and the servers in dynamic groups will be added or removed automatically based on the rules that you define.

The method for creating dynamic server groups is the same as for creating static server groups. The difference is that when dynamic is selected, you are presented with a page that allows you to define the rules for the group.

Dynamic groups can be converted to static groups, and all servers remain in the group, but all rules will be lost when they are converted.

To create dynamic groups using the New Groups Option, perform the following steps:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears.
2. From the **Resource** menu, choose **New Group**. The New Private Static Group dialog appears.
3. To create a private dynamic group, navigate to the place in the group hierarchy where you want to create the group by clicking the group links displayed in the Name and Type field.

   Or

   To create a public dynamic group, click the Public Groups link in the Name and Type field, and navigate to the location in the group hierarchy where you want to create the group.

The Save In drop down list is populated according to the location you drill down to in order to create your group. You can use the Save In drop-down list to verify that you are in the correct location in the group hierarchy, and you can move to a different location in the hierarchy by selecting it from the Save In list.

4 Enter the name of the group in the Save As text box.

5 In the Options field, select Dynamic.

6 Click **Create Rules**. The Manage Servers Properties page appears, with the Rules tab selected, as Figure 7-7 shows.

Figure 7-7: Manage Servers Properties Page with Rules Tab Selected

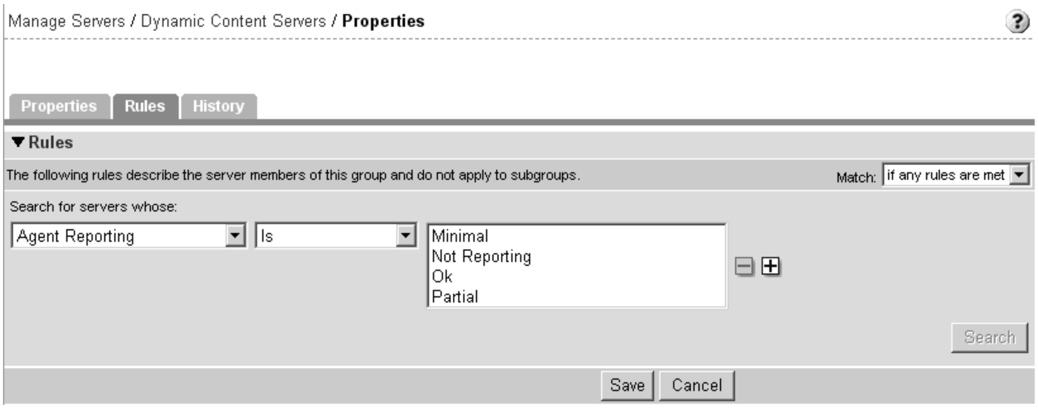

Manage Servers / Dynamic Content Servers / **Properties**

The following rules describe the server members of this group and do not apply to subgroups.

Match: if any rules are met

Search for servers whose:

Agent Reporting | Is | Minimal, Not Reporting, Ok, Partial

7 Select the criteria that apply to the servers you would like the group to include. See "From the navigation panel, click Servers ‰ Search. The Search Rules page appears with the Servers tab displayed." on page 240 in this chapter for more information.

Create as many lines of criteria as required to adequately describe your server group rules by clicking the plus button to add a new line. Conversely, to remove lines of criteria, click the minus button next to the line you want to remove.

8 From the Match drop down list select either "If all rules are met" or "If any rules are met."

9 (Optional) Click **Search** to apply the rules to a server search to validate the results.

10 Click **Save** to save the rules that apply to your group.

Creating Groups by Using Search

To create groups by using search, perform the following steps:

1 From the navigation panel, click Servers ➤ Search. The Search Rules page appears with the Servers tab displayed.

To search for servers, use this tab. To search for groups, select the Groups tab.

2 Use the criteria to create the rules used for a server search and to identify servers for dynamic groups. The options in the user interface for specifying dynamic group rules and for using Search are the same. See "Criteria for Search and Dynamic Group Rules" on page 240 in this chapter for more information.

Create as many lines of criteria as required to adequately describe your server search or your server group rules by clicking the plus button to add a new line. Conversely, to remove lines of criteria, click minus button next to the line you want to remove.

3 If you are defining dynamic server group rules, click **Save** to save your rules. You can also click **Search** to use your rules to perform a server search.

Or

If you are doing a server search, click **Search** to perform the server search, or click **Save** to save the search as the rules for a dynamic group.

The New Dynamic Group dialog appears.

Criteria for Search and Dynamic Group Rules

The following table describes the rules that you can use to search for servers or to create dynamic server groups.

Note that anywhere you can enter text, you can enter a wildcard (*) character to broaden your results.

Table 7-2: Rules for Server Search and for Creating Dynamic Groups

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
PROPERTIES		

Table 7-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Agent Discovery Date: The date that the Opsware Agent was installed.	• Is after • Is before	Drop-down lists with the day, month, and year
	• Is within the last	User-entered text
	• Is today	N/A
Agent Reporting: Whether the Opsware Agent is reporting to Opsware SAS.	• Is • Is not	• Has not reported • OK • Registration in progress • Reporting error
Agent Status: Whether the Opsware Agent is reachable by Opsware SAS.	• Is • Is not	• Not reachable • OK
Agent Version: The version of the Opsware Agent – such as 14.2.3b.	• Contains • Does not contain • Is • Is not	User-entered text
Custom Attribute (any): The name of a custom attribute that is associated with the server through attachment or inheritance.	• Contains • Is	User-entered text
Custom Attribute (local): The name of a custom attribute that is locally attached to the server,	• Contains • Is	User-entered text
Customer: The customer or account that the server is associated with.	• Is • Is not	Popup window of customers

Table 7-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Deployment Stage: The stage the server performs within a lifecycle environment.	• Is • Is not	• In Deployment • Live • Not Specified • Offline The values that appear in this list are customizable; in addition to the values above, values specific to your environment might appear.
Facility: The collection of servers managed by an Opsware SAS installation.	• Is • Is not	Popup window of facilities When Opsware SAS is running multimaster mode, the list can contain many facilities.
Group Membership: Whether the server belongs to a group.	• Is	Popup window of groups
Host name: The host name of the server – such as m004.company.com.	• Contains • Does not contain • Is • Is not	User-entered text
Name (any): This enables searching for any name or IP address associated with a server.	• Contains • Does not contain • Is • Is not	User-entered text

Table 7-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Notes: The contents of the Notes field from the Properties tab for a server.	• Contains • Does not contain • Is • Is not	User-entered text
Opsware Display Name: The user-configurable name for the server.	• Contains • Does not contain • Is • Is not	User-entered text By default, Opsware SAS uses the configured host name of the server until a user edits it.
Server Use: How the server is being used – such as Development, Staging, Production.	• Is • Is not	• Development • Not Specified • Production • Staging You can customize values that appear in this list. In addition to the values above, values specific to your environment might appear.

Table 7-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Service Level: The user-defined category that can be used as an organizational tool.	• Is • Is attached here or below	Popup window of service levels Servers can be associated with multiple service levels. See the *Opsware® SAS User's Guide* for more information about working with service levels.
OPSWARE PROPERTIES		
Application Configuration: Whether the server uses the Application Configuration feature.	• Is not used • Is used	N/A
Code Deployment: Whether the server uses the Code Deployment feature.	• Is not used • Is used	N/A
Configuration Tracking: Whether the Configuration Tracking feature is monitoring or backing up specific files or configurations on a server.	• Is off • Is on	N/A
Lifecycle: The server states that are part of bringing a server into Opsware SAS.	• Is • Is not	• Available • Build Failed • Deactivated • Installing OS • Managed

Table 7-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Server ID: The internal ID Opsware SAS uses to identify the server.	• Is • Is not	User-entered text In most cases, the Server ID is the same as the MID.
SOFTWARE		
Attached Software: The software that is assigned or modeled through the Opsware reconcile operation – the installation process.	• Is • Is attached here or below	Popup window of software
Installed Patches: Whether a patch is installed on the server.	• Contains • Does not contain • Is • Is not	User-entered text
Installed Software: The package reported installed on the server.	• Contains • Does not contain • Is • Is not	User-entered text A package does not have to be installed by Opsware SAS to be reported as installed on a server.
OS Version: The OS version defined by OS definitions in the OS Provisioning feature.	• Is • Is not	Popup window of operating systems
Reported OS: For Windows – the version reported by the OS, for Unix – the version returned by the uname command.	• Contains • Does not contain • Is • Is not	User-entered text

Table 7-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Windows Service: The names of the Windows services that are reported to Opsware SAS.	• Contains • Does not contain • Is • Is not	User-entered text
NETWORK		
DNS Search Domains: The domains configured to be searched in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
DNS Servers: The IP addresses of the DNS servers configured in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
Default Gateway: The IP address of the default router.	• Contains • Does not contain • Is • Is not	User-entered text
IP Address: Any Internet Protocol address for the server.	• Contains • Does not contain • Is • Is not	User-entered text
MAC Address: Any Media Access Control address, which is the network interface card's unique hardware number.	• Contains • Does not contain • Is • Is not	User-entered text

Table 7-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
WINS Servers: The Windows Internet Naming Servers configured in the server's network settings.	• Contains • Does not contain • Is • Is not	User-entered text
HARDWARE		
CPU Make and Model: The vendor name and CPU model for the server – such as GENUINEINTEL Intel(R) Pentium(R) 4 CPU 2.60GHz.	• Is • Is not	Popup window of CPU makes and models
CPU Speed: The Central Processing Unit speed in gigahertz [GHz].	• Does not equal • Equals • Is greater than • Is less than	User-entered text A 600 Mhz machine should be entered as `0.6`.
Make and Model: The vendor name and server model for the server – such as Compaq - DL360.	• Is • Is not	Popup window of server makes and models
Number of CPUs: The number of CPUs on the server.	• Does not equal • Equals • Is greater than • Is less than	User-entered text
RAM: The amount of RAM on the server in megabytes [MB].	• Does not equal • Equals • Is greater than • Is less than	User-entered text To enter 1 Gigabyte, type `1024`.

Table 7-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
Serial Number: The serial number of the server.	• Contains • Does not contain • Is • Is not	User-entered text
Storage Make and Model: The vendor name and storage model for the server – such as WDC - WD800BB-75DKA0.	• Is • Is not	Popup window of storage makes and models
CUSTOM FIELDS		
A **Numeric** field	• Does not equal • Equals • Is greater than • Is less than	User-entered text
A **String** field	• Contains • Does not contain • Is • Is not	User-entered text
A **URI** field	• Contains • Does not contain • Is • Is not	User-entered text
A **File** field	• Contains • Does not contain • Is • Is not	User-entered text

Table 7-2: Rules for Server Search and for Creating Dynamic Groups (continued)

RULE NAME AND DESCRIPTION	OPERATORS	VALUE
A **Date** field	• Is after • Is before	Drop-down lists with the day, month, and year
	• Is within the last	User-entered text
	• Is today	N/A

Adding a Server to Static Groups, Both Public and Private

To add a server or a server group to a static group, perform the following steps:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears.
2. Navigate to the static group that you want to add servers and server groups to.
3. Click the check box next to the group and select **Edit** ➤ **Add Servers**. The Select Servers and Groups to Add to [group name] window opens. The window is populated with the same servers and groups visible on the Manage Servers page. You can use the Status, OS, and Customer filters to change the servers and groups that appear on the list.

When you select **Edit** ➤ **Add Servers** and then select a group, the servers in the group are added. The group itself is not added. If you want to add a group to a group, select **Edit** ➤ **Copy to Group** from the menu.

4. Click the check box next to the servers and groups you want to add, and click **Add** at the bottom of the window.

Adding a Server to Dynamic Groups, Both Public and Private

Servers are added to dynamic server groups automatically, based on the rules created for the group. To change the membership of a dynamic group, add, delete, or update the dynamic group rules.

To update the rules for a dynamic group, perform the following steps:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears.
2. Navigate to the Properties page of the dynamic group that you want to update rules for; navigate by clicking the group links in Manage Server list.
3. Select the Rules tab.
4. Click **Edit**.
5. To add criteria to the existing rules, click the plus (+) button next to existing criteria. The fields for the criteria appear. Enter the values for the rule. See “From the navigation panel, click Servers ‰ Search. The Search Rules page appears with the Servers tab displayed.” on page 240 in this chapter for descriptions of these values.
6. To delete criteria, click the minus (-) button next to the criteria you want to delete from the rules. The criteria are removed from the page.
7. Click **Save**.

Groups can be added manually to a dynamic group, in either of the following two ways:

1. Click the check box next to the name of the server group to which you want to add a server group, and select **Edit ➤ Copy to Group**. Then follow the steps for creating a new group, either static or dynamic. See “Ways to Create Server Groups” on page 235 in this chapter for more information.

   Or

   Navigate to the server group that you want to add a server group to.
2. Click the check box next to “This Group” and select **Resource ➤ New Group**. Then follow the steps for creating a new group, either static or dynamic. See “Ways to Create Server Groups” on page 235 in this chapter for more information.

Removing Servers from Static Groups, Both Public and Private

To remove a server from a static server group, perform the following steps:

A server does not need to be deactivated to be removed from a group. You can remove a server from a Static group at any time. The Membership tab for a server displays all the server groups of which the server is a member.

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears.

2. Navigate to the static server group that you want to remove servers from until you reach the level where the servers in question are located.

Servers can belong to more than one group, so if you want to remove a server from each group it belongs to, you must locate and remove each instance of the server from all groups.

3. Click the check box next to the servers you want to remove and select **Edit ➤ Remove from Group**.

   As soon as you select that menu option, the server is removed and the screen refreshes to display the current members of the group.

Removing Servers from Dynamic Groups, Both Public and Private

Servers are removed from dynamic server groups automatically, based on the rules created for the group. To remove servers, you can create or update rules. To add rules to an existing server group, perform the following steps:

1. From the navigation panel, click **Servers ➤ Manage Servers**. The Manage Servers page appears.

2. Navigate to the Properties page of the dynamic group that you want to add rules to, by drilling down into the group and clicking the “This Group” link.

3. Select the Rules tab.

4. Click **Edit**.

5. Add one or more lines of criteria to the existing rules by following the procedure described in the “Creating Dynamic Server Groups by Using the New Group Option” topic, step 3 on page 240 or see “From the navigation panel, click Servers ‰ Search. The Search Rules page appears with the Servers tab displayed.” on page 240 in this chapter for more information.

Moving Servers from One Static Group to Another

The process of moving servers from one group to another is similar to copying servers from one group to another with the exception that with a move, the servers do not remain in the original group.

The **Edit ➤ Move** menu option is disabled when you select servers in a dynamic group. To move servers in a dynamic group, chose the **Edit ➤ Copy to Group** menu command. This menu command is also disabled for public groups when you do not have permission to manage public groups.

To move servers from one group to another, perform the following steps:

1. From the navigation panel, click **Servers ➤ Manage Servers**. The Manage Servers page appears.
2. Navigate to the group containing the servers you want to move and select those servers.
3. From the **Edit** menu, choose **Move**. The Move [group name] Group dialog box appears.
4. Navigate to the place in the group hierarchy where you want to create the group by clicking the group links displayed in the Name and Type field.
5. Click **Move**.

Duplicating Server Groups

Duplicating a group is similar to copying a group. To duplicate an existing group, perform the following steps:

1. Navigate to the group that you want to duplicate.
2. Click the check box next to the group name.

You can only select one group at a time to duplicate.

3 From the **Edit** menu, choose **Duplicate Group**. The Duplicate [group name] Group dialog appears as shown in Figure 7-8.

Figure 7-8: Duplicate Group Dialog Box

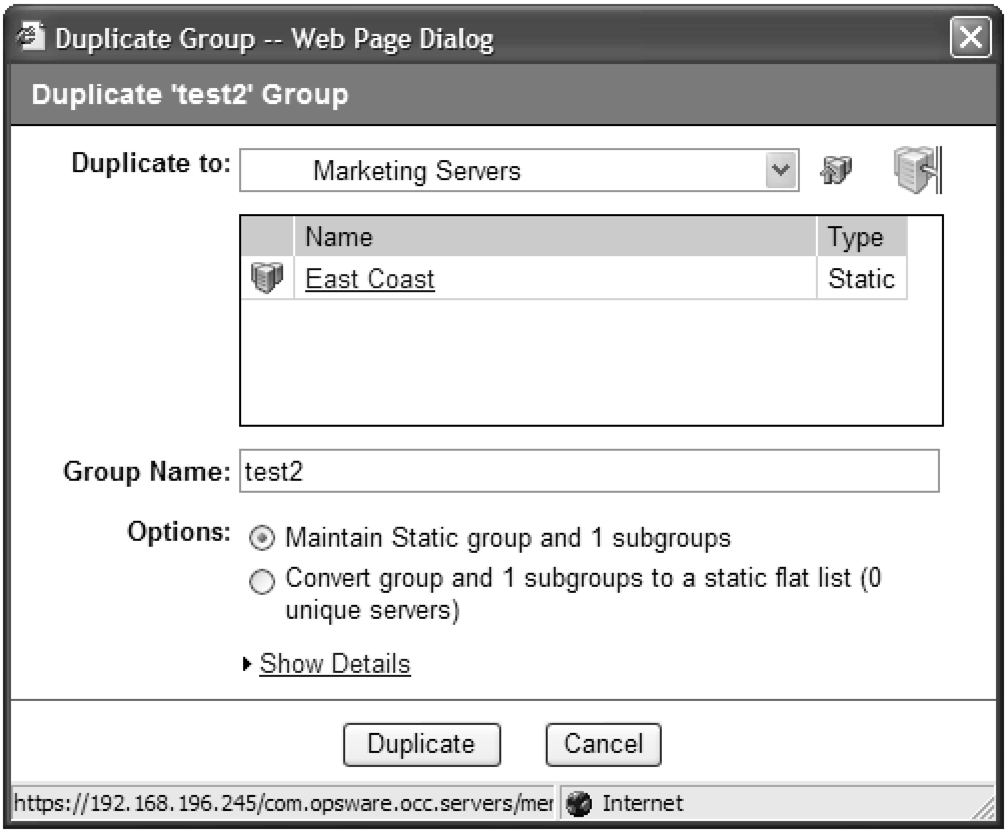
Duplicate Group -- Web Page Dialog
Show Details
Duplicate
Cancel
https://192.168.196.245/com.opsware.occ.servers/mer
Internet

Duplicate 'test2' Group

Duplicate to: Marketing Servers

	Name	Type
	East Coast	Static

Group Name: test2

Options:
● Maintain Static group and 1 subgroups
○ Convert group and 1 subgroups to a static flat list (0 unique servers)

4 Select the location for the newly-duplicated group using the Duplicate In drop down list.

The Duplicate In drop down list's default location is the location of the group you select to duplicate. Navigate through the hierarchy of groups to find the location where you want the newly-duplicated group to reside.

5 Enter the name of the new group.

6 The values on the Options field will vary depending on whether the group duplicated is static or dynamic.

7. If the group is static, select one of the following options:

    - **Maintain Static group and <number> subgroups**: Selecting this option copies the group and its hierarchy as is.
    - **Convert group and <number> subgroups to a static flat list**: Selecting this option copies the group and flattens the hierarchy.

8. If the group is dynamic, select one of the following options:

    - **Maintain rules & group hierarchy**: Selecting this option copies the group and leaves it as a dynamic group, it also copies any subgroups.
    - **Convert to static group, maintain hierarchy**: Selecting this option copies the group, but turns it into a static group with the current servers defined by the rules and also copies any subgroups.
    - **Convert group and <number> subgroups to a static flat list**: Selecting this option copies the group, and turns it into a static group and flattens any hierarchy.

9. Click **Duplicate**. The newly-duplicated group appears in the selected destination.

Rules for Deleting Server Groups

It's important to note the distinction between removing servers from a group, deleting a server, and deleting a group.

Removing servers from a group only removes the server from the selected static group, but the server itself remains in the global list of servers and is still managed by Opsware SAS.

Deleting a server can only be performed on a server whose status is deactivated. Selecting the Delete Server option completely removes the server from within Opsware SAS, although its history remains.

Deleting a server group removes the group, but the servers in the group still remain in the list of servers and in any other groups for which they are members.

A group cannot be deleted when any of the following conditions apply:

- Software is attached to the group or a subgroup of the group.
- Access control boundaries are attached to the groups or a subgroup of the group.
- Servers and groups are selected together for deletion.

Deleting a Server Group

To delete a server group, perform the following steps:

1. Click the check box next to the server group you want to delete.
2. From the **Edit** menu, choose **Delete Group**. A confirmation message appears, detailing the number of servers and server groups in the server group that you want to delete.
3. Click **OK** to complete the deletion of the server group.

   The screen refreshes, showing the list of servers and groups without the deleted server group.

Server Life Cycle

This section provides information on the server life cycle within Opsware SAS and contains the following topics:

- OS Provisioning and the Server Life Cycle
- Server Properties
- Server Management Tasks Related to the Server Life Cycle
- Changing the Use and Stage Values for Servers
- Editing the Properties of a Server
- Tasks Associated with Deactivating a Server
- Deactivating a Server
- Deleting a Server from Opsware SAS
- Overview of Cloning a Server

OS Provisioning and the Server Life Cycle

Opsware SAS is designed to enable multiple teams to work together to provision servers. The OS Provisioning feature allows IT teams to separate the tasks of readying servers for provisioning (such as mounting servers in racks and connecting them to a network) from provisioning the servers with operating systems and applications.

For example, someone mounts a new server in a rack and connects it to the Opsware build network. Next, they boot the server for the first time by using an Opsware Boot Floppy or CD or by using the network.

At a later time, a different system administrator can select the available server from the Server Pool list and provision it with an OS. In the available state, servers do not have the target OS installed and might not have access to disk resources.

During OS provisioning, servers progress through the Opsware SAS life cycle state changes:

Unprovisioned (No OS Build Agent) ➤ Available ➤ Installing OS ➤ Managed

Table 7-3 describes the Opsware SAS server life cycle values.

Table 7-3: Opsware SAS Life Cycle Values for Servers

OPSWARE LIFE CYCLE VALUE	DESCRIPTION
Server Pool Values	
Available	Indicates a server on which the OS Build Agent was installed and is running, but the target OS has not been installed on the server. The OS Build Agent is a small agent that can run in the memory of the bare-metal server. See “Operating System Provisioning” on page 567 in Chapter 15 for more information.
Installing OS	Indicates that a user is installing the target OS on the server. The server stays in the Server Pool list until the installation process finishes successfully. Then the server moves to the Manage Servers list. See “Installing an OS by Using a Template” on page 585 in Chapter 15 for more information. See “Installing an OS by Using a Custom Installation” on page 590 in Chapter 15 for more information.
Build Failed	Indicates a server on which the OS Build Agent was installed and is running, but the installation of the target OS failed. The server remains in the Server Pool list with this status for 7 days before Opsware SAS deletes the entry. See “Recovering when an OS Installation Fails” on page 592 in Chapter 15 for more information.

Table 7-3: Opsware SAS Life Cycle Values for Servers (continued)

OPSWARE LIFE CYCLE VALUE	DESCRIPTION
Managed Server Values	
Managed	Indicates a server that Opsware SAS is managing. Opsware SAS performs periodic reachability checks on managed servers. After a server reaches this life cycle state, the entry for the server moves from the Server Pool list to the Manage Servers list. On managed servers, you can use Opsware SAS to install applications and patches.
Deactivated	Indicates an Opsware-managed server that was removed from service. However, the server's history still exists in Opsware SAS. Deactivated servers are not reachable.

Table 7-4 shows which server icons appear in Opsware SAS and explains what each icon indicates in regard to the server life cycle.

Table 7-4: Server Icons in Opsware SAS

SERVER ICON	DESCRIPTION
🖥	Indicates a server that is available to have a target OS installed on it and on which an Opsware OS Build Agent is installed. Appears in the Server Pool list.
🖥	Indicates a server on which the OS Provisioning feature is in the process of installing the target OS. Appears in the Server Pool list.
⊗	Indicates an available server on which an error occurred while the OS Provisioning Subsystem was installing a target OS. Appears in the My Jobs panel in the home page, in the list in the My Jobs page, and in the Server Pool list.

Table 7-4: Server Icons in Opsware SAS

SERVER ICON	DESCRIPTION
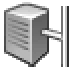	Indicates a server that the Opsware Command Center is managing and that Opsware SAS can communicate with. An Opsware Agent is running on the server. Appears in the My Jobs panel in the home page, in the list in the My Jobs page, in the Manage Servers list, and in the server lists in the Opsware wizards.
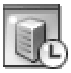	Indicates a server that is scheduled for an operation (install software, uninstall software, and so forth). Appears in the My Jobs panel in the home page and in the list in the My Jobs page.
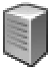	Indicates a managed server that Opsware SAS cannot communicate with (it is Not Reachable) because the Opsware Agent on the server cannot connect to Opsware SAS. If you want to discover reasons why the managed server is unreachable, you can run a Communication Test. See “Agent Reachability Communication Tests” on page 204 in Chapter 6 for more information. Appears in the Manage Servers list.
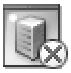	Indicates a managed server on which an error occurred while Opsware SAS was installing or uninstalling software. Appears in the My Jobs panel in the home page and in the list in the My Jobs page.
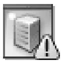	Indicates a managed server on which a warning occurred while Opsware SAS was installing or uninstalling software. Appears in the My Jobs panel in the home page and in the list in the My Jobs page.
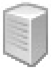	Indicates a server that was deactivated in Opsware SAS so that it is currently not managed and not reachable. Appears in the Manage Servers list and in the server lists in the Opsware wizards (however, it is not selectable in the wizards).

Table 7-4: Server Icons in Opsware SAS

SERVER ICON	DESCRIPTION
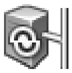	Indicates a managed server on which the configuration file on the server is out of sync with the Application Configuration Template (Opsware model). Appears only in the Application Configuration feature and the server list in the OCC Client.
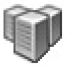	Indicates a static server group. The same states that apply to single servers apply to groups. See "Server Groups" on page 226 in this chapter for information about the different types of server groups.
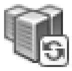	Indicates a dynamic server group. The same states that apply to single servers apply to groups.
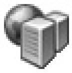	Indicates a public and static server group. The same states that apply to single servers apply to groups.
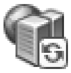	Indicates a public and dynamic server group. The same states that apply to single servers apply to groups.

Server Properties

Figure 7-9 shows the Server Properties columns. Table 7-5, Table 7-6, and Table 7-7 describe the Status, Stage, and Use properties for managed servers.

Figure 7-9: Server Properties Columns in the Manage Servers List

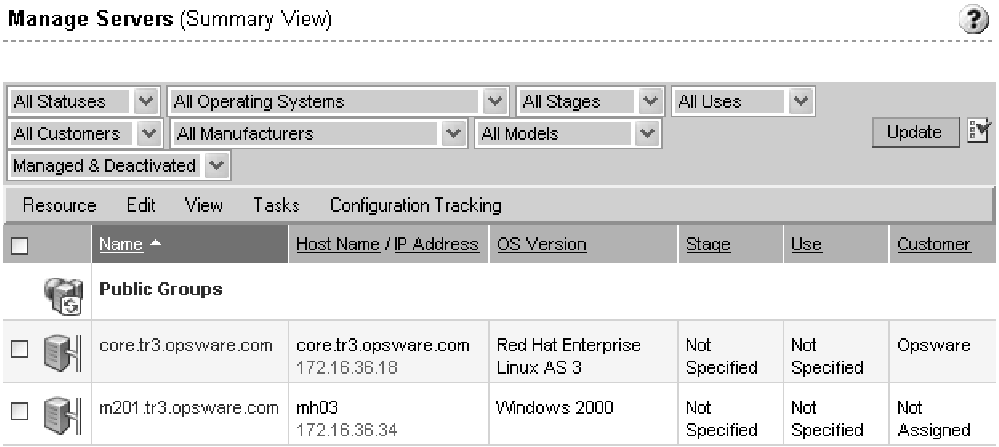

Manage Servers (Summary View)

All Statuses | All Operating Systems | All Stages | All Uses
All Customers | All Manufacturers | All Models
Managed & Deactivated

□	Name ▲	Host Name / IP Address	OS Version	Stage	Use	Customer
	Public Groups					
□	core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	Red Hat Enterprise Linux AS 3	Not Specified	Not Specified	Opsware
□	m201.tr3.opsware.com	mh03 172.16.36.34	Windows 2000	Not Specified	Not Specified	Not Assigned

The Status property is represented by an icon in the first column in the Manage Servers list.

Status (short for Agent Status) is set automatically by Opsware SAS.

Opsware SAS toggles each server between OK and Not Reachable by reachability checks.

The Status value specifies the ability of Opsware SAS to manage servers. Opsware SAS automatically detects the status of servers. To verify the current status of a server, click Update in the Server Properties page for that server.

Table 7-5: Values for the Status Property for Managed Servers

STATUS VALUE	DESCRIPTION
OK	Server is manageable by Opsware SAS. Represented as text (OK) in the properties page for a server. Represented as an icon in the Manage Servers and Server Pool lists:

Table 7-5: Values for the Status Property for Managed Servers (continued)

STATUS VALUE	DESCRIPTION
Not Reachable	Server is unreachable by Opsware SAS due to an error (for example, it cannot connect to the Opsware core); automatically set by Opsware SAS. Represented as text (Not Reachable) in the properties page for a server. Represented as an icon in the Manage Servers list: If you want to discover reasons why the managed server is unreachable, you can run a Communication Test. See "Agent Reachability Communication Tests" on page 204 in Chapter 6 for more information.

Stage (short for Deployment Stage) is set by a user.

The Stage value specifies the stages of deployment for servers; for example, a server is live or offline.

Your Opsware administrator can change the values for the Stage property. By default, Opsware SAS is installed with the following Stage values.

Table 7-6: Values for the Stage Property for Managed Servers

STAGE VALUE	DESCRIPTION
In Deployment	Initial stage after being fully initialized.
Live	Your organization defines the meaning of this stage.
Not Specified	The default value for a server. Cannot be changed by the Opsware administrator.
Offline	Your organization defines the meaning of this stage.
Ops Ready	Your organization defines the meaning of this stage.

Use (short for Server Use) is set by a user.

The Use value specifies how an organization is utilizing servers. For example, a server is a staging server. Users set this property for servers.

By default, Opsware SAS is installed with the following Use values. Except for the Staging, Production, and Not Specified values, an Opsware administrator can change the default values. The CDR feature depends on the Staging and Production values. Therefore, these default values cannot be changed or deleted.

Table 7-7: Values for the Use Property for Managed Servers

USE VALUE	DESCRIPTION
Development	A server that is not being used in production.
Not Specified	The default value.
Production	Fully live in-use servers (includes Opsware core servers).
Staging	A staging server for production.

Server Management Tasks Related to the Server Life Cycle

Managing servers in Opsware SAS involves the following standard tasks:

- Bringing a new server into Opsware SAS so that it appears in the Server Pool

  See “New Server Booting” on page 576 in Chapter 15 for more information.

- Installing an operating system on a server

  See “Ways to Install Operating Systems on Servers” on page 584 in Chapter 15 for more information.

- Installing a patch

  See “Patch Installation and Uninstallation for Unix” on page 535 in Chapter 13 for more information.

  See “Patch Installation” on page 499 in Chapter 12 for more information.

- Installing an application

  See “Software Provisioning” on page 549 in Chapter 14 for more information.

- Reprovisioning a server with a new OS

  See “Reprovisioning a Solaris or Linux Server” on page 596 in Chapter 15 for more information.

  You can reprovision Solaris and Linux servers so that they are running another version of the same OS so long as the hardware supports that new version of the OS.

You can reprovision servers built by Opsware SAS and Opsware-managed servers by using this feature.

You cannot reprovision a Linux server so that it runs a Windows OS.

- Deactivating a managed server so that Opsware SAS no longer manages it.

  See "Tasks Associated with Deactivating a Server" on page 267 in this chapter for more information.

You accomplish server management tasks by using the menus in the Manage Servers list, as Figure 7-10 shows.

Figure 7-10: Menus in the Manage Servers List

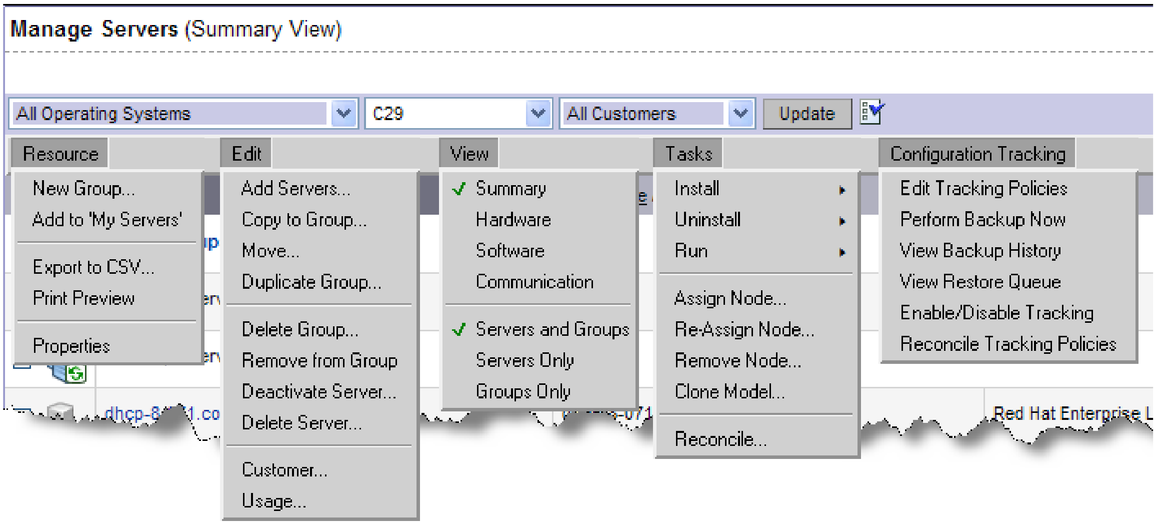

See "Software Provisioning" on page 549 in Chapter 14 for more information.

See "Reconcile" on page 699 in Appendix A for more information.

See "Distributed Script Execution" on page 411 in Chapter 11 for more information about how to run a distributed script on a server.

See "Manage Servers List" on page 119 in Chapter 4 for information about managed servers.

Changing the Use and Stage Values for Servers

Perform the following steps to change the Use and Stage values for multiple servers simultaneously:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server that you want to deactivate.

   Or

   Search for the server that you want to deactivate.

   See "Using the Search Feature" on page 125 in Chapter 4 for more information. See "Server Searching by IP Address" on page 141 in Chapter 4 for more information.

2. Select the servers that you want different Use or Stage values for.

3. Choose **Edit** ➤ **Usage** from the menu above the Manage Servers list. A window prompts you to select different values, as Figure 7-11 shows.

Figure 7-11: Edit Server Popup Window

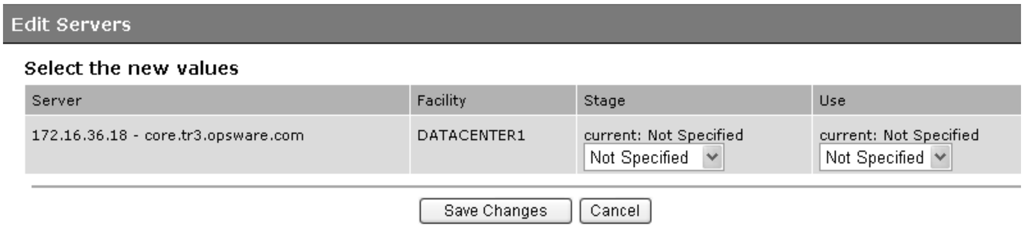

Edit Servers

Select the new values

Server	Facility	Stage	Use
172.16.36.18 - core.tr3.opsware.com	DATACENTER1	current: Not Specified Not Specified	current: Not Specified Not Specified

4. Select the Use and Stage values from the lists.

5. Click **Save Changes**. The window closes and the Manage Servers list refreshes with the updated values.

Editing the Properties of a Server

You can edit a server only if you have permission in the Opsware Command Center to access the customer to whom the server is associated.

When you edit the properties of a server, the server itself does not change; how Opsware SAS views it changes. Perform the following steps to edit the properties of a server:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server whose properties you want to edit.

Or

Search for the server that you want to edit.

See "Using the Search Feature" on page 125 in Chapter 4 for more information. See "Server Searching by IP Address" on page 141 in Chapter 4 for more information.

2. Click the Server Display name. The Properties page for the server appears.

3. Change any of the following properties for the server:

   - To change the name that appears in the Opsware Command Center, edit the text in the Name field.
   - To change the description of the server, edit the Notes field.
   - To change the customer associated with the server, select a different customer from the list. Your Opsware administrator defines the options for customer selections. Contact your Opsware administrator if the list does not contain the customer that you want to associate with this server.

You cannot change the customer associated with a server when the server is part of a CDR service, synchronization, or sequence. See the *Opsware® SAS Configuration Guide* for more information.

   - To change the Use or Stage of the server, make your changes in either of those lists.

     See "Server Properties" on page 261 in this chapter for more information.

   - To change whether configuration tracking is enabled or disabled for the server, select a value from the list.

     See "Configuration Tracking" on page 643 in Chapter 17 for more information.

4. To save your changes, click **Save**.

To change the custom attributes of the server, select the Custom Attributes tab. The Manage Servers: Custom Attributes page appears. See "Managing Custom Attributes" on page 283 in this chapter for more information.

Tasks Associated with Deactivating a Server

You will want to deactivate a server when Opsware SAS removes the server from management. For example, you are moving the server to a warehouse for storage. Additionally, you might choose to deactivate a server when you need to rebuild it from scratch, without using the OS Provisioning feature.

When you deactivate a server, information about the server remains in the Opsware Model Repository for auditing purposes.

After you deactivate a server, you can reactivate it by re-installing an Opsware Agent with the Opsware Agent Installer and the `--clean` command line option.

See "Agent Reachability Communication Tests" on page 204 in Chapter 6 for more information.

When you deactivate a server, you accomplish the following tasks:

- Reset the node assignments in the Software Tree to the defaults.
- Remove custom attributes from the server.
- Delete any configuration tracking policies from the server that are associated with backups.
- Set the server life cycle value to Deactivated.

You cannot deactivate a server when it is part of a CDR service, synchronization, or sequence. See the *Opsware® SAS Configuration Guide* for more information.

Deactivating a Server

Perform the following steps to deactivate a server:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server that you want to deactivate.

   Or

   Search for the server that you want to deactivate.

   See "Using the Search Feature" on page 125 in Chapter 4 for more information. See "Server Searching by IP Address" on page 141 in Chapter 4 for more information.

2. Select the servers that you want to deactivate.

3. Choose **Edit ➤ Deactivate Server** from the menu above the Manage Servers list. A confirmation dialog box prompts you to confirm the deactivation.
4. Click **OK**. The Manage Servers list refreshes and the server appears with a deactivated icon.

Deleting a Server from Opsware SAS

When you want to remove all record of a server from Opsware SAS, you can delete it.

✓ You must deactivate a server before you can delete it from Opsware SAS.

When you delete a server from Opsware SAS, it has these effects:

- Deletes all job information in the My Jobs feature
- Deletes the server from the Model Repository

Perform the following steps to delete a server:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server that you want to delete.

   Or

   Search for the server that you want to delete.

   See "Using the Search Feature" on page 125 in Chapter 4 for more information. See "Server Searching by IP Address" on page 141 in Chapter 4 for more information.
2. Select the servers that you want to delete.
3. Choose **Edit ➤ Delete Server** from the menu above the Manage Servers list. A confirmation dialog box prompts you to confirm the deletion.
4. Click **OK**. The Manage Servers list refreshes and the server disappears from the list.

Overview of Cloning a Server

The Opsware Command Center includes a feature that allows users to copy a server. Copying (referred to as cloning in the Opsware Command Center) is useful when you need to add more capacity to your operational environment.

A user selects a source server (the master server) and copies the configuration of that server to one or more other servers (the target servers). The other servers are assigned to every node to which the master server is assigned. The copied servers contain the same operating system, software applications, and configuration as the original server (though the customer association and facility location remain the same on the target servers). The copied servers also include any changes to the default server configuration made through Opsware SAS.

The only restriction to cloning a server is that both servers need to be on the same platform. However, any existing user-configurable nodes are removed from the target server when you clone servers. The target server is made to look exactly like the master server in terms of node assignments.

Cloning a Server

Perform the following steps to clone a server:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the servers that you want to clone.

2. Select two or more servers. When copying servers, you must select at least two servers so that one server can be the master. The other servers are the ones that you want to copy the nodes to.

   Opsware, Inc. recommends that you select servers with similar hardware architecture (based on the value in the Reported OS column in the Manage Servers list).

3. Choose **Tasks** ➤ **Clone** model from the menu above the Manage Servers list.

   A window prompts you to select the master server, as Figure 7-12 shows.

Figure 7-12: Clone Servers Popup Window in the Opsware Command Center

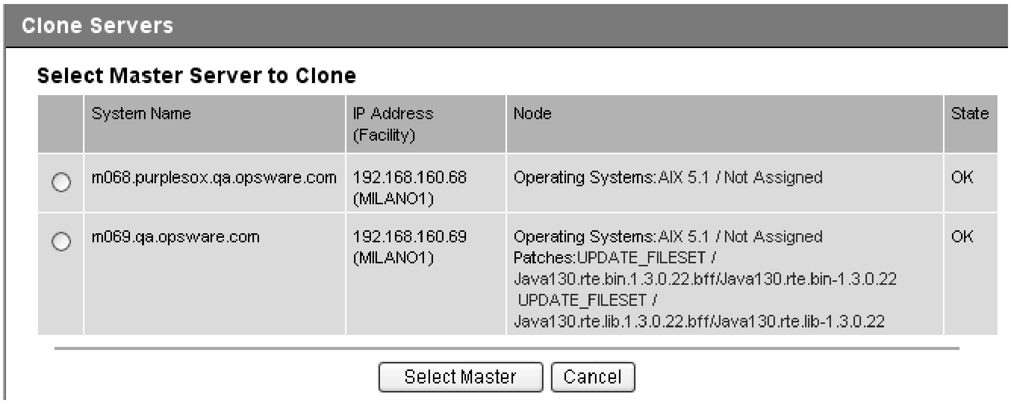

Clone Servers

Select Master Server to Clone

	System Name	IP Address (Facility)	Node	State
○	m068.purplesox.qa.opsware.com	192.168.160.68 (MILANO1)	Operating Systems:AIX 5.1 / Not Assigned	OK
○	m069.qa.opsware.com	192.168.160.69 (MILANO1)	Operating Systems:AIX 5.1 / Not Assigned Patches:UPDATE_FILESET / Java130.rte.bin.1.3.0.22.bff/Java130.rte.bin-1.3.0.22 UPDATE_FILESET / Java130.rte.lib.1.3.0.22.bff/Java130.rte.lib-1.3.0.22	OK

4. Select the option for the master server that you want to use and click **Select Master**.

A confirmation page appears that shows how the nodes on the master server are copied to the server that you select.

5 Click **Commit Clone**.

The Server List reappears and shows the updated target servers.

After you clone a master server to target servers, you must reconcile the target servers for Opsware SAS to install the software on the target servers. See "Reconcile" on page 699 in Appendix A for information about reconciling servers.

Server Management Jobs

This section provides information about server management jobs within Opsware SAS and contains the following topics:

- My Jobs
- My Jobs Display Information
- My Jobs in the OCC
- My Jobs in OCC Client
- Viewing Job Details in the OCC
- Viewing My Job Details from inside the OCC Client

My Jobs

The My Jobs information is available only on a per-user basis. You cannot log in as an Opsware administrator to see the jobs that other Opsware users have run. The My Jobs information appears in two places in the Opsware Command Center:

- A panel on the Opsware Command Center home page that lists your most recent six jobs
- A page (accessed by clicking My Jobs in the navigation panel) that lists all the jobs that you have run

  Opsware SAS maintains information about the server operations that you have run for the last 30 days in the My Jobs list. By default, the jobs are deleted from the Opsware

Model Repository after 30 days. (The bottom of the My Jobs page indicates how long this interval is set for the Opsware SAS installation at your organization.)

My Jobs Display Information

For each job, the My Jobs lists display the following information:

- The name of the job, which is a link to a page that displays more detailed information about the job
- The date and time the job started or is scheduled to start (using your preference for time display)
- The number of servers that the job affects
- The status of the job:
  - Scheduled
  - In Progress
  - Completed
  - Completed with errors
  - Completed with warnings
- You can search for an existing Job by the Job's ID. On the Home page, select Job from the drop-down list and enter a Job ID and click **Go**.

My Jobs in the OCC

The My Jobs feature in the OCC provides information about the following Command Engine scripts:

- Reconcile (install software, uninstall software, install a template, patching servers, and reconcile)
- OS provisioning
- CDR requests
- Distributed script execution
- Custom Extensions

My Jobs in OCC Client

My Jobs feature in the OCC Client displays job logs about the following OCC Client jobs:

- Creating snapshots
- Pushing or Auditing an Application Configuration
- Auditing Servers
- Creating Server Snapshots
- Creating, Installing, and Reconciling Packages
- Any jobs scheduled to be run at a future date

My Jobs in the OCC Client appears in the Jobs Logs feature window.

Viewing Job Details in the OCC

Perform the following steps to view the job details:

1. From the Opsware Command Center home page, click the link in the My Jobs panel for the job that you want to view, as Figure 7-13 shows.

Figure 7-13: My Jobs Panel in the Opsware Command Center Home Page

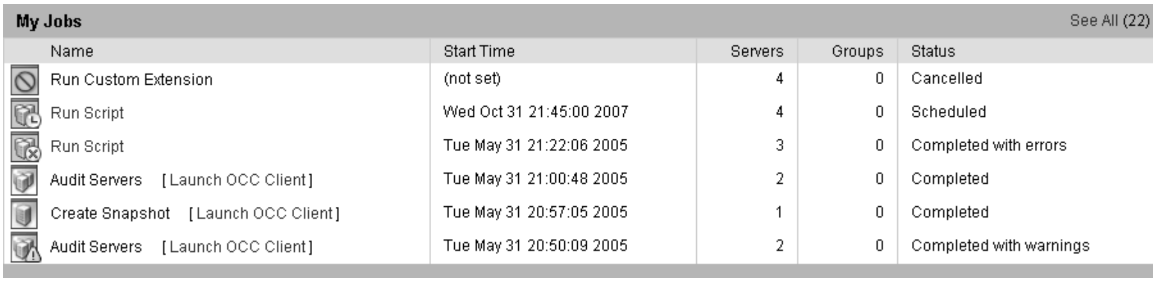

My Jobs See All (22)

Name	Start Time	Servers	Groups	Status
Run Custom Extension	(not set)	4	0	Cancelled
Run Script	Wed Oct 31 21:45:00 2007	4	0	Scheduled
Run Script	Tue May 31 21:22:06 2005	3	0	Completed with errors
Audit Servers [Launch OCC Client]	Tue May 31 21:00:48 2005	2	0	Completed
Create Snapshot [Launch OCC Client]	Tue May 31 20:57:05 2005	1	0	Completed
Audit Servers [Launch OCC Client]	Tue May 31 20:50:09 2005	2	0	Completed with warnings

Or

From the navigation panel, click My Jobs and then click the link for the job to open a window that shows the details of the job, as Figure 7-14 shows.

Figure 7-14: My Jobs Page Accessed from the Navigation Panel

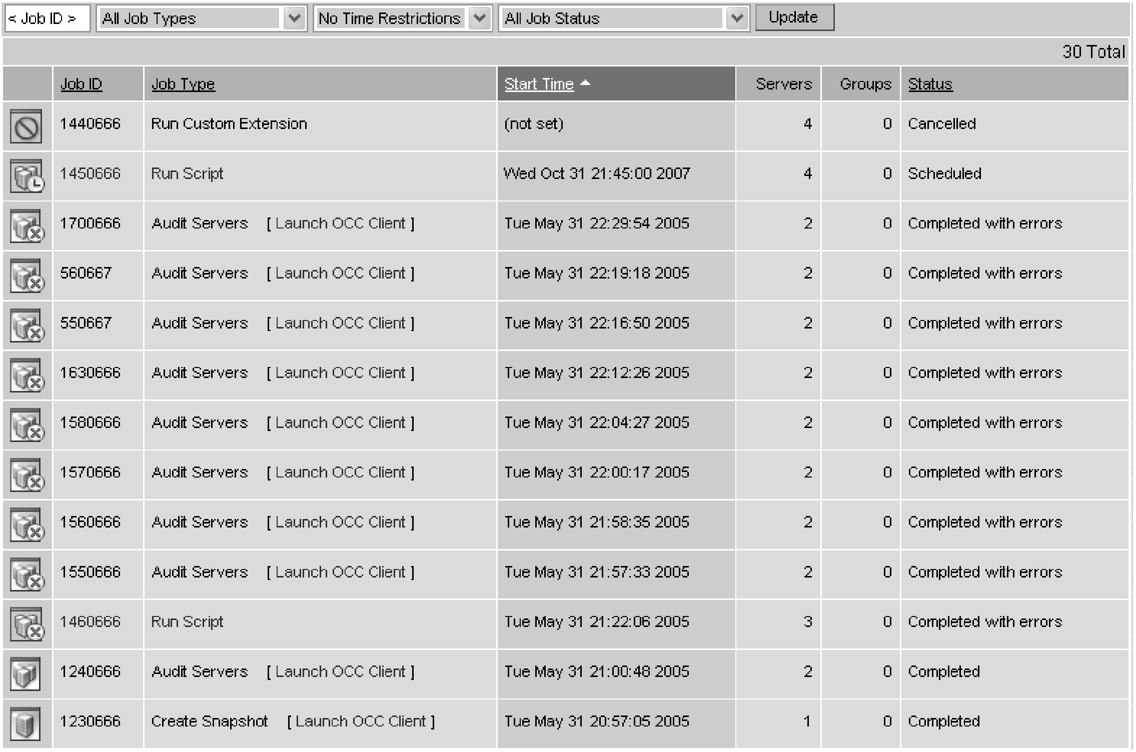

< Job ID > | All Job Types | No Time Restrictions | All Job Status | Update

30 Total

	Job ID	Job Type	Start Time ▲	Servers	Groups	Status
⊘	1440666	Run Custom Extension	(not set)	4	0	Cancelled
	1450666	Run Script	Wed Oct 31 21:45:00 2007	4	0	Scheduled
	1700666	Audit Servers [Launch OCC Client]	Tue May 31 22:29:54 2005	2	0	Completed with errors
	560667	Audit Servers [Launch OCC Client]	Tue May 31 22:19:18 2005	2	0	Completed with errors
	550667	Audit Servers [Launch OCC Client]	Tue May 31 22:16:50 2005	2	0	Completed with errors
	1630666	Audit Servers [Launch OCC Client]	Tue May 31 22:12:26 2005	2	0	Completed with errors
	1580666	Audit Servers [Launch OCC Client]	Tue May 31 22:04:27 2005	2	0	Completed with errors
	1570666	Audit Servers [Launch OCC Client]	Tue May 31 22:00:17 2005	2	0	Completed with errors
	1560666	Audit Servers [Launch OCC Client]	Tue May 31 21:58:35 2005	2	0	Completed with errors
	1550666	Audit Servers [Launch OCC Client]	Tue May 31 21:57:33 2005	2	0	Completed with errors
	1460666	Run Script	Tue May 31 21:22:06 2005	3	0	Completed with errors
	1240666	Audit Servers [Launch OCC Client]	Tue May 31 21:00:48 2005	2	0	Completed
	1230666	Create Snapshot [Launch OCC Client]	Tue May 31 20:57:05 2005	1	0	Completed

The My Jobs page displays the operations that you performed.

Click **View Details** to see detailed information about the job. The My Jobs information contains a build log for the job. This build log contains any error messages that Opsware SAS generates. See Figure 7-15.

Figure 7-15: Communication Test Details Page in the My Job Window

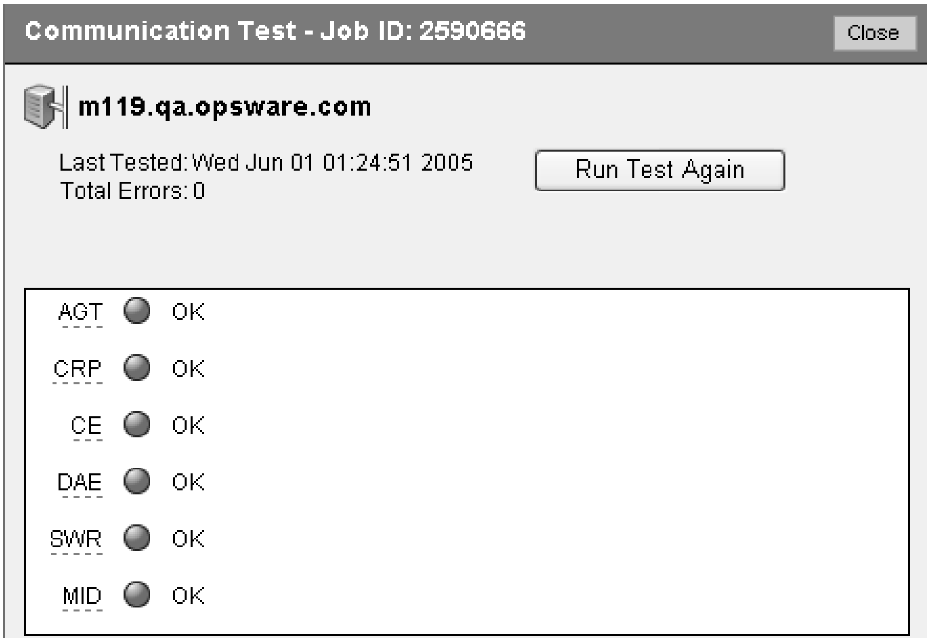

Communication Test - Job ID: 2590666

m119.qa.opsware.com

Last Tested: Wed Jun 01 01:24:51 2005
Total Errors: 0

AGT	●	OK
CRP	●	OK
CE	●	OK
DAE	●	OK
SWR	●	OK
MID	●	OK

Viewing My Job Details from inside the OCC Client

You can also view job details from inside the OCC Client. If the job was run from inside the OCC Client, you will see a link next to the job named Launch OCC Client, as shown in Figure 7-16.

Figure 7-16: My Jobs with link to OCC Client

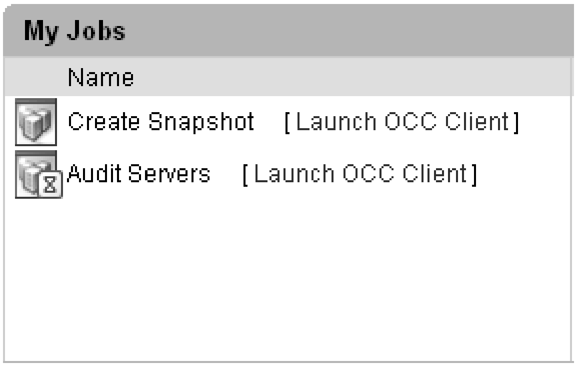

My Jobs

Name
Create Snapshot [Launch OCC Client]
Audit Servers [Launch OCC Client]

To launch the OCC Client and view the job details, click the link.

Server Management Scheduling and Notification

This section contains the following topics:

- Scheduling a Job
- Sending Email Notification

The time used for the scheduled job is specified in the user's preferred time zone (which can be modified in My Profile). If the user does not have a preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

Scheduling a Job

Perform the following steps to schedule server management tasks:

1. In the Schedule and Notify page of an Opsware Wizard, choose the Run Now option to execute the operation immediately or choose the Specify Time Option to schedule the operation at a later date and time. See Figure 7-17

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

2. When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have "Allow Run Refresh Jobs" permissions, you will see additional options. Select either of the following options:

   - **Option 1**: Membership is determined based on the "Time of Confirm Selection." Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.

- **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

Figure 7-17: Scheduling a Job in an Opsware Wizard

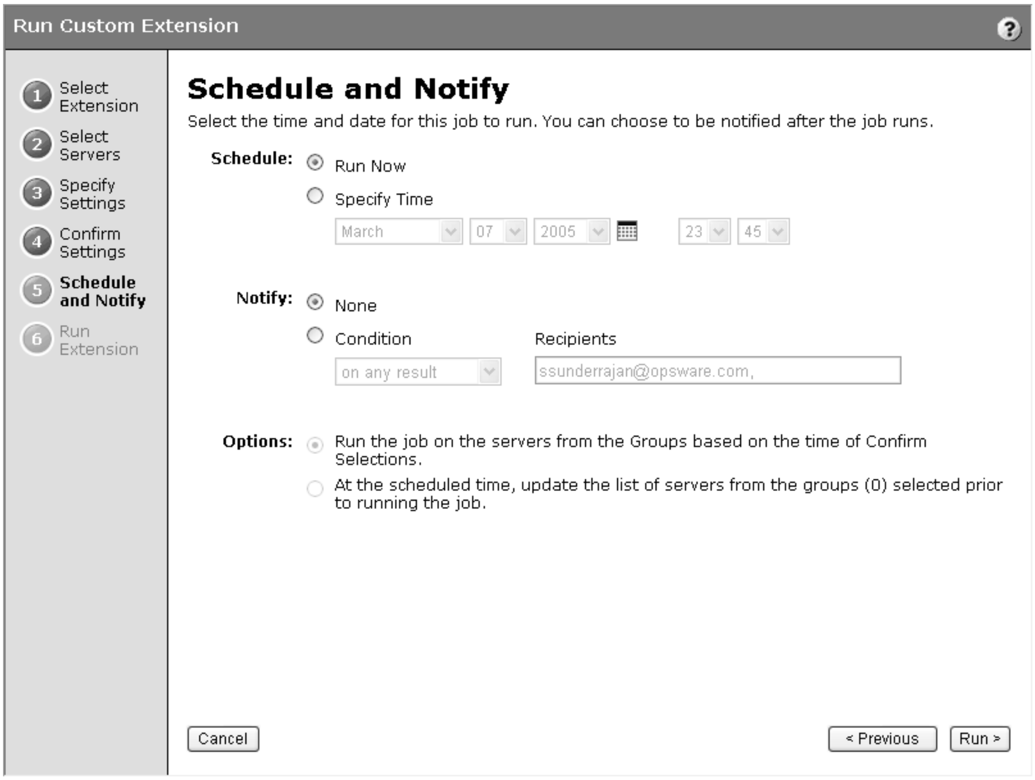

Run Custom Extension

1 Select Extension
2 Select Servers
3 Specify Settings
4 Confirm Settings
5 **Schedule and Notify**
6 Run Extension

Schedule and Notify

Select the time and date for this job to run. You can choose to be notified after the job runs.

Schedule: ● Run Now
○ Specify Time
March | 07 | 2005 | 23 | 45

Notify: ● None
○ Condition — on any result
Recipients: ssunderrajan@opsware.com,

Options: ● Run the job on the servers from the Groups based on the time of Confirm Selections.
○ At the scheduled time, update the list of servers from the groups (0) selected prior to running the job.

3 Additionally, you can view a scheduled job in the My Jobs page and change the date and time for the job to run, or cancel the job entirely. (Click the name of a scheduled job to open a window to change the date or time that the job will run or cancel it.)

Sending Email Notification

The email notification feature provides you with an option to receive an email summarizing the job details when a job is over, and to notify others at the address they have registered with Opsware SAS.

You have the option of sending the notification

- On the job success only
- On the job failure only
- On any result

You can also send notification to various people based on the condition of the job. For example, you can select to send the notification to your manager only on the job success, select to send the notification to yourself on any result of the job, and select to send the notification to support only on the job failure.

To send an email about the job details, choose the Condition option on the Schedule and Notify page and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field. See Figure 7-18.

Figure 7-18: Notifying about a Job in the Opsware Command Center

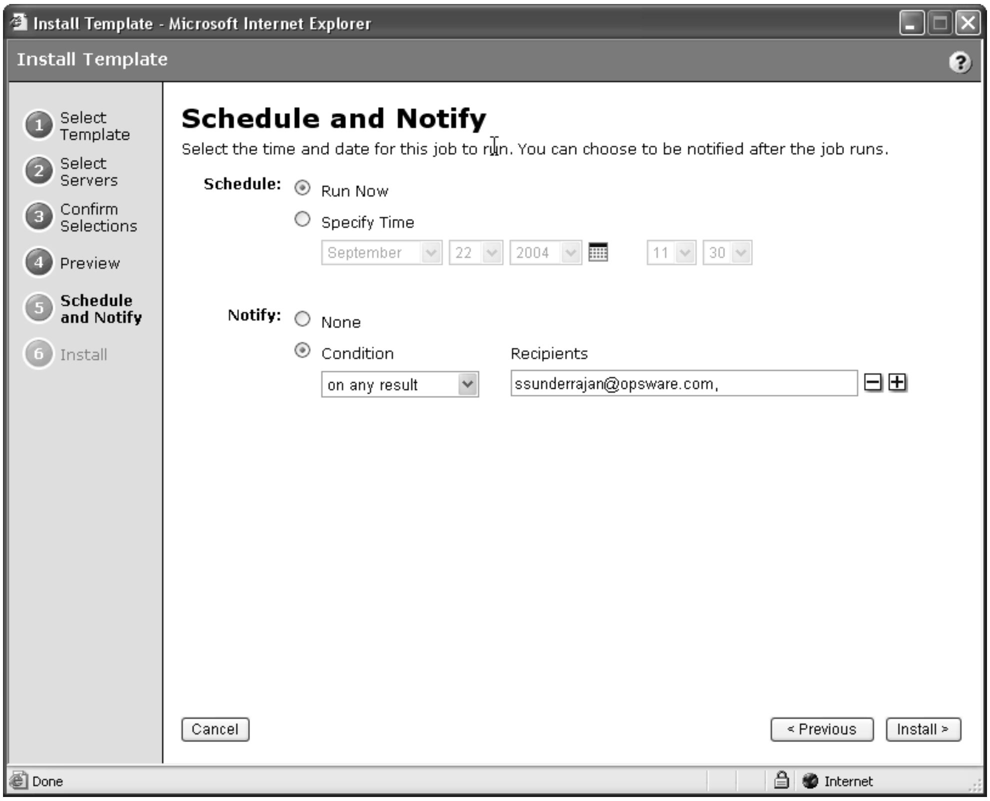

Install Template - Microsoft Internet Explorer

Install Template

1 Select Template
2 Select Servers
3 Confirm Selections
4 Preview
5 **Schedule and Notify**
6 Install

Schedule and Notify

Select the time and date for this job to run. You can choose to be notified after the job runs.

Schedule: ● Run Now
○ Specify Time
September 22 2004 11 30

Notify: ○ None
● Condition: on any result
Recipients: ssunderrajan@opsware.com,

Time-Outs for Server Management Jobs

Table 7-8 describes the time-out values that apply to the server management operations in Opsware SAS.

Table 7-8: Time-Outs in Opsware SAS

TIME OUT (MINUTES)	OPSWARE SAS OPERATION
420 (7 hours)	Reconciling software (installation and uninstallation).
2	Starting Command Engine sessions in response to a command. If the Opsware Agent does not start executing the command within this time, the command will time out and the Command Engine script will continue.
30	Responding to a command (for example, after a reboot, the maximum time to wait until a server responds) or sending a message to the Command Engine from the Opsware Agent. If the Opsware Agent does not respond to the Command Engine at least once during this interval, the command will time out and the Command Engine script will continue. Opsware SAS polls the Opsware Agent every 15 minutes and if the Opsware Agent fails to respond two consecutive times, the command will time out.

Customizing the Monitor Time-Out Duration

If you would like to set a different monitor time-out duration, you can create a custom attribute named `OPSW_reconcile_monitor_timeout` and change the number of minutes before a time out occurs. For each type of hardware running in your operational environment, you can set a custom attribute with the time-out duration that you want. To set a time-out duration for a type of hardware, click Environment ➤ Hardware in the navigation panel. Then, navigate to the type of hardware that you want to add a custom attribute for.

During reconcile, a periodic heartbeat occurs between the Opsware Agent and the Command Engine to ensure that the agent is still responsive. This setting controls the maximum amount of time that can pass between these heartbeat messages. Typically, you only need to increase this setting if you install software that reboots the server, and the time that it takes for the server to reboot and for the agent to restart exceeds the default value.

Custom Fields for Servers

This section provides information on custom fields for servers within Opsware SAS and discusses the following topics:

- Overview of Custom Fields for Servers
- Changing the Value for a Custom Field

Overview of Custom Fields for Servers

In Opsware SAS, you can store custom server data that is specific to your operational environment. These fields were created specifically for your Opsware SAS installation – all servers in an Opsware SAS installation contain the same number and type of custom fields. Custom fields can contain files, URLs, text strings, numbers, and dates.

In addition to adding files, URLs, text strings, numbers, and dates, you can use custom fields to search for servers based on a value stored in a custom field. You can also use a custom field as criteria to create a dynamic server group.

In order for the custom fields to appear in the Manage Server: Properties page, you will have to initially create a custom field. To create a custom field, you will need to install the custFields.py Custom Extension which is available only from the Content Starter Pack. Contact your Opsware Technical Support for assistance in installing the custFields.py Custom Extension in Opsware SAS.

Typed Data Supported for Custom Fields

The custom fields designated for your operational environment will vary. However, the custom fields created for your environment support values with the following data types:

- **Number**: The value provided must be a long integer.
- **Short String**: A text string that must be less than or equal to 4000 characters.
- **Long String**: A text string with no length restrictions. Custom fields with this type cannot be used in search or in the rules for dynamic server groups.

- **Date**: The value will be verified for correctness.
- **File**: Indicates a file attachment.
- **URI**: A Uniform Resource Identifier string, which is validated as a URI.

Opsware SAS validates the value you enter in a custom field based on the data type specified for the field.

Uses of Custom Fields

The custom fields designated for your operational environment will vary. However, you can use custom fields to accomplish any of the following goals:

- To store the date of patch installation
- To assign a severity rating between 1 and 10 to a Hotfix
- To store an ID from an internal bug tracking system with its associated patch
- To store a JPEG image of the back of a server and have that JPEG associated with that server
- To store a Microsoft Word document describing the disaster recovery steps for a server or group of servers
- To search for a server based on the value of a custom field

Changing the Value for a Custom Field

1 From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server whose custom fields you want to change.

Or

Search for the server whose custom fields you want to change.

2 Click the server name. The Manage Servers: Properties page appears. Scroll to the Custom Fields section of the page. The custom fields that appear in the properties page are specific to your operational environment. Figure 7-19 shows an example of the types of custom fields that can appear for a server.

Figure 7-19: Custom Fields That Appear for a Server

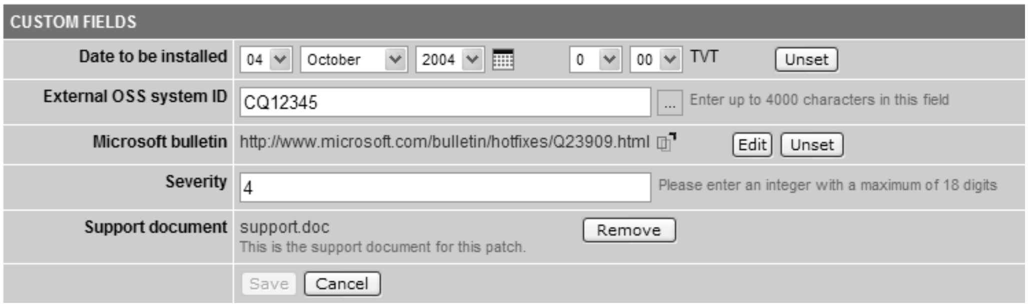

CUSTOM FIELDS

Date to be installed	04 ▾ October ▾ 2004 ▾ 📅 0 ▾ 00 ▾ TVT Unset
External OSS system ID	CQ12345 ... Enter up to 4000 characters in this field
Microsoft bulletin	http://www.microsoft.com/bulletin/hotfixes/Q23909.html Edit Unset
Severity	4 Please enter an integer with a maximum of 18 digits
Support document	support.doc This is the support document for this patch. Remove
	Save Cancel

3 To change a value in a field requires a number or text, enter the value in the field.

4 To add a file, click **Browse** and select the file from the dialog box.

5 To remove a file, click **Remove**.

6 To specify a URL, enter the URL in the field or click **Edit** to change an existing URL.

7 To add or change a date, click the 📅 icon and select the date or enter the date in the field by selecting the appropriate day, month, and year from the drop-down lists.

8 Click **Save**.

Custom Attributes for Servers

This section provides information on custom attributes for servers within Opsware SAS and contains the following topics:

- Overview of Custom Attributes for Servers
- Managing Custom Attributes
- Adding Server Custom Attributes
- Editing Server Custom Attributes
- Deleting Server Custom Attributes

Overview of Custom Attributes for Servers

Users often need to store specific miscellaneous information in the Opsware Model Repository to facilitate server or application installation and configuration, scripting, or other purposes.

The Opsware Command Center provides a data management function by allowing users to set custom attributes for servers. These custom attributes include setting miscellaneous parameters and named data values. Users can write scripts that use these parameters and values when performing a variety of functions, including network and server configuration, notifications, and CRON script configuration.

Custom attributes can be accessed by software packages at installation time to configure settings that might be unique to the installation.

For information about how to set custom attributes required by the software running on a specific server, contact the group responsible for packaging your applications, as Figure 7-20 shows.

Figure 7-20: Custom Attributes Set for a Server

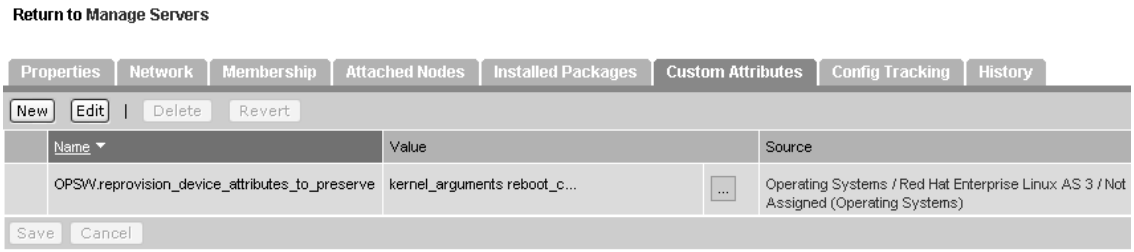

Name	Value	Source
OPSW.reprovision_device_attributes_to_preserve	kernel_arguments reboot_c...	Operating Systems / Red Hat Enterprise Linux AS 3 / Not Assigned (Operating Systems)

Managing Custom Attributes

Do not edit or remove custom attributes without verifying that the change you are making does not impact other users or critical Opsware operations.

To set custom attributes that affect a specific server, use the Manage Servers list. After locating the server and displaying the server properties, select the Custom Attributes tab. The Opsware Command Center displays the currently defined custom attributes for the selected server.

When you add or edit server custom attributes using the Opsware Command Center, Opsware SAS removes leading and trailing whitespace characters from custom attribute values.

See "Adding Server Custom Attributes" on page 284 in this chapter for more information.

To set custom attributes for all the servers assigned to a node in the Software Tree, navigate to the node where you want to set attributes, select the Custom Attributes tab, and add attributes for all the servers assigned to the node. Inheritance applies when you set custom attributes for nodes in the Software Tree.

See the *Opsware® SAS Configuration Guide* for more information.

Additionally, you can set custom attributes that affect every server associated with a specific customer or for every server in a facility. Navigate to the customer or facility where you want to set attributes. Select Environment ➤ Customers or Facilities in the navigation panel, and click the correct name in the list. Select the Custom Attributes tab, and add attributes for all the servers associated with the customer or located in the facility. When you use this option, you define custom attributes at a customer-specific or facility-specific level. The procedure to add custom attributes to a customer or facility is the same as adding custom attributes to individual servers.

Additionally, you can add custom attributes for a server group by viewing a server group, and then selecting the Custom Attributes tab for that group. The procedure to add custom attributes to a server group is the same as adding custom attributes to individual servers.

Adding Server Custom Attributes

For information about how to set custom attributes required by the software running on a specific server, contact the group responsible for packaging your applications.

Perform the following steps to add a custom attribute for a server:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server for which you want to add custom attributes.

   Or

   Search for the server for which you want to add custom attributes.

2. Click the display name of the server. The Manage Servers: Server Properties page appears.

3. Select the Custom Attributes tab. The Manage Servers: Custom Attributes page appears.
4. Click **New**.
5. Enter the name and value for the custom attribute that you want to add.
6. Click **Save**.

See “Using the Search Feature” on page 125 in Chapter 4 for more information. See “Server Searching by IP Address” on page 141 in Chapter 4 for more information

Editing Server Custom Attributes

If you want to change the name of a custom attribute entry, you need to create a new custom attribute and delete the old custom attribute.

Perform the following steps to edit custom attributes for a server:

1. From the navigation panel, click Servers ➤ Manage Servers or Server Pool. The Manage Servers page appears. Browse the list to find the server for which you want to edit custom attributes.

   Or

   Search for the server for which you want to edit custom attributes.
2. Click the display name of the server. The Manage Servers: Server Properties page appears.
3. Select the Custom Attributes tab to change the custom attributes of the server. The Manage Servers: Custom Attributes page appears.
4. Click the attribute name link for the custom attribute that you want to change.
5. Update the value of the custom attribute.
6. Click **Save** to save your changes. The Manage Servers: Custom Attributes page reappears with the updated value.

See “Using the Search Feature” on page 125 in Chapter 4 for more information. See “Server Searching by IP Address” on page 141 in Chapter 4 for more information.

Deleting Server Custom Attributes

Perform the following steps to delete custom attributes for a server:

1. From the navigation panel, click Servers ➤ Manage Servers or Server Pool. The Manage Servers page appears. Browse the list to find the server from which you want to remove custom attributes.

   Or

   Search for the server from which you want to remove custom attributes.

   See "Using the Search Feature" on page 125 in Chapter 4 for more information. See "Server Searching by IP Address" on page 141 in Chapter 4 for more information.

2. Click the display name of the server. The Manage Servers: Server Properties page appears.
3. Select the Custom Attributes tab. The Manage Servers: Custom Attributes page appears.
4. Select the check box for the custom attribute that you want to delete.
5. Click **Delete**. The Opsware Command Center displays a confirmation page.
6. Click **OK** to delete the custom attribute.

See "Using the Search Feature" on page 125 in Chapter 4 for more information. See "Server Searching by IP Address" on page 141 in Chapter 4 for more information.

Service Levels

This section provides information about service levels within Opsware SAS and contains the following topics:

- Overview of Service Levels
- Adding a Service Level to the Opsware Command Center
- Editing a Service Level
- Ways to View the Service Level for Servers
- Assigning a Server to a Service Level
- Removing a Server from a Service Level

Overview of Service Levels

Service levels are user-defined categories that enable you to group servers in an arbitrary way and design your own organizational schemes. For example, you can organize your servers by functionality (finance, engineering, and so forth) or tier (Web, application, and database) or by ontogeny (development, staging, and production).

You can also create service levels to indicate the Service Level Agreement (SLA) for the servers that your IT organization manages. For example, you might create service levels to denote Silver, Gold, and Platinum services.

Please note that assigning servers to service levels does not cause Opsware SAS to operate any differently with respect to those servers. When you first use service levels, the categories will be fairly empty and by default, when an Opsware Agent is installed on a server in the operational environment, the server will be added to the UNKNOWN Service Level.

Adding a Service Level to the Opsware Command Center

Perform the following steps to add a service level to the Opsware Command Center:

1. From the navigation panel, click Environment ➤ Service Levels. The Service Levels page appears.

2 Navigate the hierarchy of service levels until you reach the point in the hierarchy where you want to add a new service level, as Figure 7-21 shows.

Figure 7-21: Service Level Hierarchy

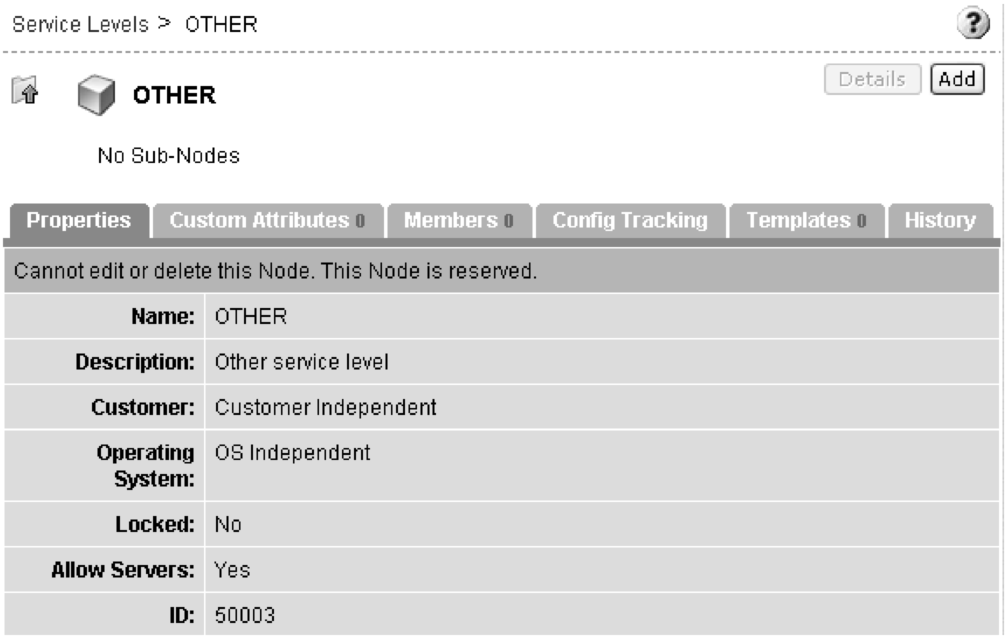

Cannot edit or delete this Node. This Node is reserved.

Name:	OTHER
Description:	Other service level
Customer:	Customer Independent
Operating System:	OS Independent
Locked:	No
Allow Servers:	Yes
ID:	50003

3 Click **Add**. The Service Levels page refreshes and the ADD SUB-NODE TO Service Levels form appears in the page.

4 Enter a name for the service level (required), and (optionally) enter notes and a description for the service level.

5 Click **Save**. The service level is added to the hierarchy of service levels. The Edit Service Level page appears, where you can change the properties of the service level, such as the customer association.

Editing a Service Level

Perform the following steps to edit a service level:

1 From the navigation panel, click Environment ➤ Service Levels. The Service Levels page appears.

2 Navigate the hierarchy of service levels until you reach the point in the hierarchy where you want to edit an existing service level.

3. Click **Edit** in the Properties tab. The page refreshes and an editable form appears for the service level properties.

4. Make changes to the service level name, description, notes, whether servers are allowed to be assigned to the service level, the associated customers and operating systems.

   Unlike nodes in the Software Tree, service levels can have multiple customers and operating systems associated with them.

5. Click **Save**.

Ways to View the Service Level for Servers

Find the server whose service levels you want to view by searching or browsing the Manage Servers list.

- If you are browsing the Manage Server list, you can find the service level for a server by locating the value in the Environment column, as Figure 7-22 shows.

Figure 7-22: Service Level Node Appearing in the Software Tab of the Server List

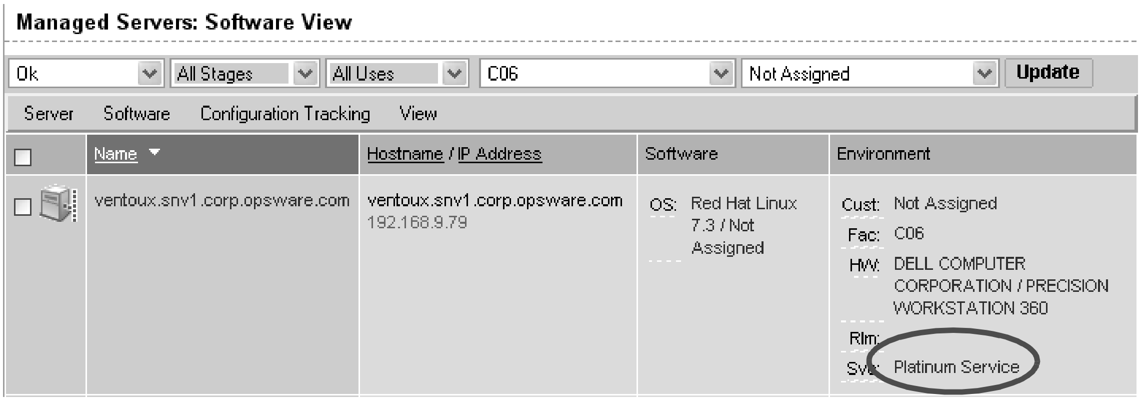

Managed Servers: Software View

□	Name ▼	Hostname / IP Address	Software	Environment
□	ventoux.snv1.corp.opsware.com	ventoux.snv1.corp.opsware.com 192.168.9.79	OS: Red Hat Linux 7.3 / Not Assigned	Cust: Not Assigned Fac: C06 HW: DELL COMPUTER CORPORATION / PRECISION WORKSTATION 360 Rlm: Svc: Platinum Service

- If you searched for the server, click the server name and then select the Attached Nodes tab. You can find the service levels, as Figure 7-23 shows.

Figure 7-23: Nodes Tab That Shows the Service Level to Which a Server is Assigned

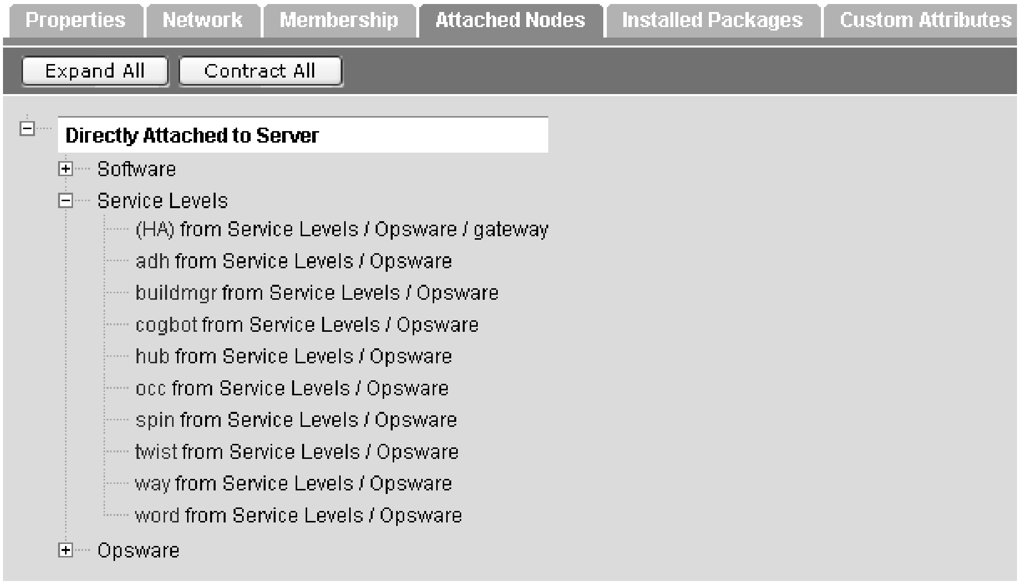

Directly Attached to Server
- Software
- Service Levels
  - (HA) from Service Levels / Opsware / gateway
  - adh from Service Levels / Opsware
  - buildmgr from Service Levels / Opsware
  - cogbot from Service Levels / Opsware
  - hub from Service Levels / Opsware
  - occ from Service Levels / Opsware
  - spin from Service Levels / Opsware
  - twist from Service Levels / Opsware
  - way from Service Levels / Opsware
  - word from Service Levels / Opsware
- Opsware

- To view all the servers assigned to a particular service level, click Environment ➤ Service Levels in the navigation panel. Navigate the hierarchy of service levels until you reach the one for which you want to see which servers are assigned. Select the Members tab, as Figure 7-24 shows.

Figure 7-24: Manage Server Assigned to a Service Level

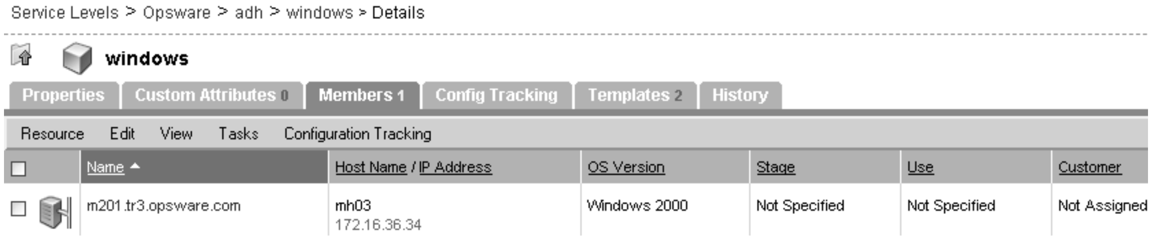

☐	Name ▲	Host Name / IP Address	OS Version	Stage	Use	Customer
☐	m201.tr3.opsware.com	mh03 172.16.36.34	Windows 2000	Not Specified	Not Specified	Not Assigned

Assigning a Server to a Service Level

Perform the following steps to assign a server to a service level:

1 From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server you want to assign to a service level.

Or

Search for the server that you want to assign to a service level.

2. Select the servers that you want to assign to a service level.
3. Choose **Tasks ➤ Assign Node** from the menu above the Manage Servers list. A window displays the categories of nodes, as Figure 7-25 shows.

Figure 7-25: Assign Nodes Popup Window

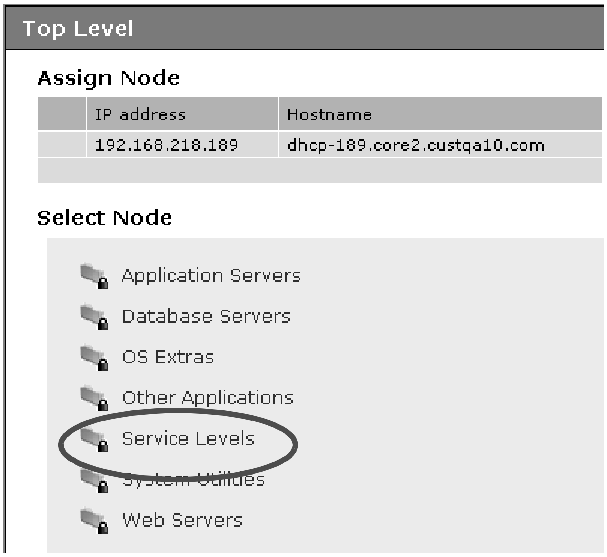

Top Level

Assign Node

	IP address	Hostname
	192.168.218.189	dhcp-189.core2.custqa10.com

Select Node

- Application Servers
- Database Servers
- OS Extras
- Other Applications
- Service Levels
- System Utilities
- Web Servers

4. Click the Service Levels link. The window refreshes to show the service levels created for your operational environment.
5. Navigate to the service level to which you want to assign the server.
6. Click **Assign**. The window closes and you are returned to the Manage Servers list.

See "Using the Search Feature" on page 125 in Chapter 4 for more information. See "Server Searching by IP Address" on page 141 in Chapter 4 for more information.

Removing a Server from a Service Level

Perform the following steps to remove a server from a service level:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server that you want to remove from a service level.

   Or

Search for the server that you want to remove from a service level.

See "Using the Search Feature" on page 125 in Chapter 4 for more information. See "Server Searching by IP Address" on page 141 in Chapter 4 for more information.

2 Select the server that you want to remove from a service level.

3 Choose **Tasks ➤ Remove Node** from the menu above the Manage Servers list. A window displays the nodes to which the server is assigned, as Figure 7-26 shows.

Figure 7-26: Remove Nodes Popup Window

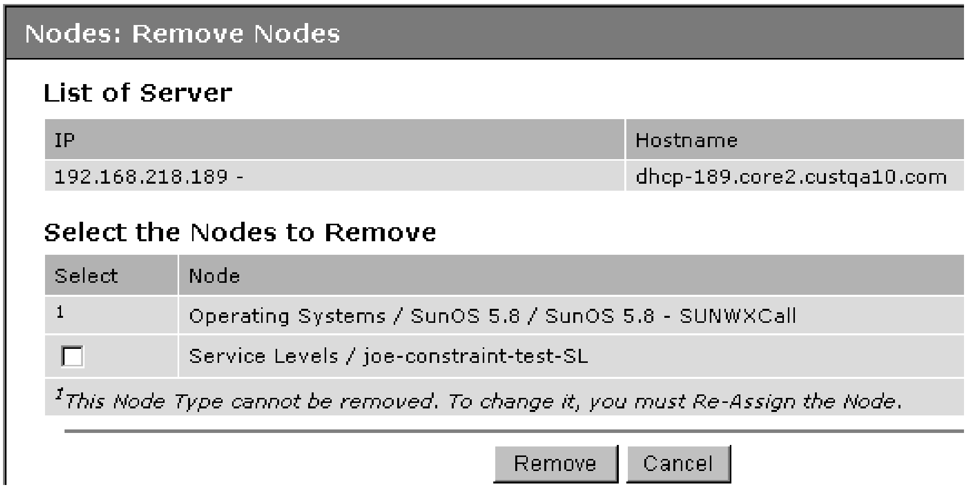

Nodes: Remove Nodes

List of Server

IP	Hostname
192.168.218.189 -	dhcp-189.core2.custqa10.com

Select the Nodes to Remove

Select	Node
1	Operating Systems / SunOS 5.8 / SunOS 5.8 - SUNWXCall
□	Service Levels / joe-constraint-test-SL

[1] *This Node Type cannot be removed. To change it, you must Re-Assign the Node.*

4 Select the service level node from which you want to remove the server and click **Remove**. You are prompted to confirm that you want to remove the server from the service level.

5 Click **Confirm Remove**. The window closes and you are returned to the Manage Servers list.

Network Configuration

This section provides information about network configuration within Opsware SAS and contains the following topics:

- Overview of the Server Network Configuration
- Configuring Networking for an Opsware Managed Server

Overview of the Server Network Configuration

You can use Opsware SAS to automatically configure network settings for a server after you install the OS.

The OS Provisioning feature provisions servers with an OS by using DHCP addresses. Because DHCP servers often assign temporary IP addresses to servers that boot over a network, system administrators typically need to assign static IP addresses (and other network properties) before the servers can be put into service. Opsware SAS enables system administrators to do this through the Opsware Command Center rather than logging onto the server manually after OS provisioning is complete.

Opsware SAS does not support managed servers that have IPv6 addresses.

The Server Network Configuration feature allows you to configure the following settings on a server that are related to its network configuration:

- Host Name
- Domain Name System (DNS) servers
- Management interface (the interface that Opsware SAS should use when managing the server)

  See “Locating the Management IP Address of a Managed Server” on page 222 in Chapter 7 for more information.
- Gateway (the IP address of the default router)
- DNS search domains
- WINS (Windows Internet Naming Service) Servers
- Configuration for each network interface, including whether the interface is configured statically or with a Dynamic Host Configuration Protocol (DHCP) IP address, host name, and subnet mask

You can alter any of these options and then apply the settings to the managed server. Opsware SAS updates the server and reboots it to cause the new settings to take effect.

Configuring Networking for an Opsware Managed Server

You can only use the Network Configuration feature for servers running Sun Solaris, Red Hat Linux, and Microsoft Windows operating systems.

Perform the following steps to configure networking for an Opsware managed server:

1 From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server that you want to configure networking on.

Or

Search for the server that you want to configure networking on.

See "Using the Search Feature" on page 125 in Chapter 4 for more information. See "Server Searching by IP Address" on page 141 in Chapter 4 for more information.

2 Click the name of the server that you want to configure networking for. The Manage Servers: Properties page appears for that server.

3 Select the Network tab. The network information for the server displays.

4 Modify any of the following settings to configure the server networking.

For all fields, the default value is the one currently configured on the server.

- **Host Name**: The host name configured on the managed server. This field only sets the name by which the server knows itself and does not update DNS records for the server.
- **Management Interface**: Instructs Opsware SAS to use a particular network interface when contacting the server. This is useful, for example, when a server has multiple network interfaces, but not all of them are reachable by the Opsware core. Designating a particular interface as the management interface allows Opsware SAS to know which interface to use for managing the server.
- **Gateway**: The IP address of the default router
- **DNS Servers**: A list of DNS nameserver IP addresses
- **Search Domains**: A list of DNS domains to search when attempting to resolve host names
- **WINS Servers**: Set for Windows only; a list of WINS server IP addresses
- **Interface Configuration** (for each network interface in the system):

- **DHCP**: If DHCP is enabled for an interface, the system uses DHCP to configure this network interface. In this case, static configuration settings (IP address, host name, and subnet mask) are not relevant for this interface, and the Opsware Command Center makes those fields not editable. If DHCP is not enabled, then static settings are required.
- **IP Address**: The IP address for this interface (unless DHCP is enabled).
- **Host Name**: The local host name for the server. This item is only required for servers running Solaris. Like the Computer Name field, this setting only affects the name by which the managed server knows itself, and does not update DNS records.
- **Subnet Mask**: The IP network mask to use for this interface.

In addition, Opsware SAS displays the management IP address and MAC address for the server; however, you cannot change these values reported by the Opsware Agent.

5 Click **Update Server** at the bottom of the page.

(If you click **Revert**, it causes any changes that you made to the fields to be discarded.)

A confirmation dialog box appears that shows the changes that will be made to the server. The confirmation dialog box includes a check box that allows you to indicate that the server should revert to its old network configuration when it cannot contact the Opsware core after you save the new network configuration. By default, the Revert check box is selected.

The Opsware Command Center does not validate the network configuration changes that you make in the Network tab. Therefore, it is possible to provide a malformed IP address in the IP address field for an interface.

6 To have the server revert to its previous network configuration if an error occurs, ensure that the check box is selected in the confirmation dialog box.

7 Click **OK** to proceed with the configuration changes.

A progress dialog box appears that shows the progress of the operation. The process of setting a new network configuration involves rebooting the managed server. The operation might take several minutes.

You can wait for the operation to complete or close the progress dialog box and perform other work in the Opsware Command Center. The status of the task is available in the My Jobs user interface if you want to check the status of the network configuration update.

Details About Changing the Domain for Windows Servers

You cannot use the DNS Domain field to change the domain name for a Windows server.

Opsware SAS does not change the domain name of a Windows server because changing the domain name of a server requires password authentication. Changing the domain name of a Windows server is a manual operation. See Figure 7-27.

Figure 7-27: DNS Domain Field Displays in the Network Tab for a Server

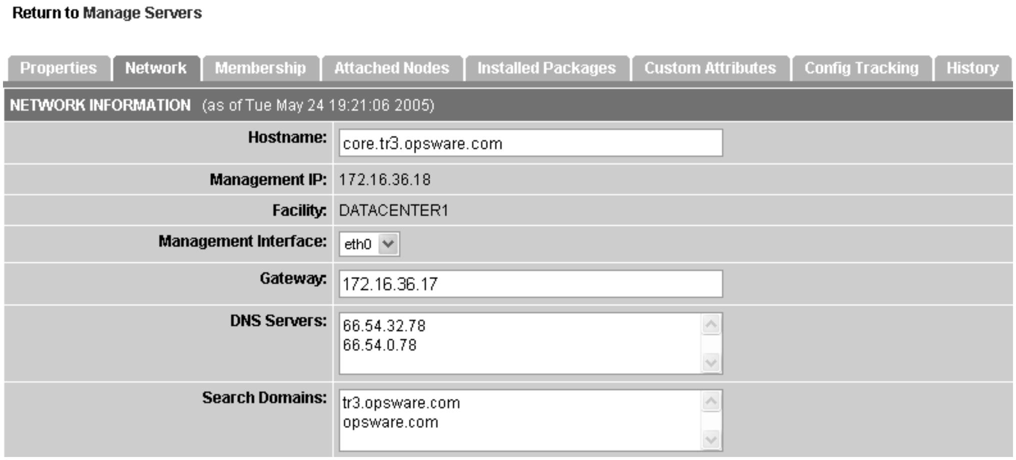

NETWORK INFORMATION (as of Tue May 24 19:21:06 2005)

Hostname:	core.tr3.opsware.com
Management IP:	172.16.36.18
Facility:	DATACENTER1
Management Interface:	eth0
Gateway:	172.16.36.17
DNS Servers:	66.54.32.78 66.54.0.78
Search Domains:	tr3.opsware.com opsware.com

Chapter 8: Server Compliance

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Server Compliance
- Server Objects
- Selection Criteria
- Snapshots
- Ways to Create a Snapshot
- Snapshot Templates
- Snapshot Scheduling
- Uses of Snapshots
- Ways to Find a Snapshot
- Auditing
- Ways to Perform an Audit
- Audit Templates
- Audit Schedules
- Audit Results
- Ways to View Audit Results

Overview of Server Compliance

In your IT environment, it is common for a server to not perform well when changes are made outside of Opsware SAS. These changes include when new patches and packages are installed, when configuration files have been modified, and when processes are started and stopped.

The Opsware Server Compliance feature enables you to keep managed servers up-to-date by comparing them to a known working server. Managed servers are compared to either the current state of a known working server or to a saved state that was recorded in a snapshot. A known working server is a managed server that is compliant (performs as expected) and is also referred to as a reference server or a baseline server.

In Server Compliance, snapshot functionality and audit functionality are used to keep servers up-to-date. A snapshot records how an Opsware-managed server is configured at a particular point in time. To create a reference server, administrators create a snapshot of a managed server deemed to be compliant. An audit compares managed servers, server groups, and snapshots to determine how they differ. By performing an audit, administrators can compare a problematic server with a known working server or with a snapshot of the problematic server when it was previously working properly. When an audit reports a difference between a reference server and a noncompliant server, they can install software and server objects to remediate the discrepancy. For example, administrators can use snapshots to audit the configuration of servers and deploy files and software to correct disparities. They can also define schedules that specify when they want snapshots to be created and audits performed, either once or as a recurring job.

The snapshot functionality is intended to assist system architects and developers who need to record objects that are worth tracking on Opsware-managed servers, such as installed packages, installed patches, hardware, and file system objects. Server Compliance provides selection criteria that let them specify what server objects they want to collect information about and how to collect the server objects. The information collected includes file comparison data, inclusions/exclusions criteria, file modification dates, and user and user group access permissions.

The auditing functionality is intended to be used by system administrators who need to investigate and identify servers that are not performing well. System administrators can use these audit results to remediate servers that are not compliant. By performing an audit, they can compare a problematic server with a known working server or with a snapshot of the problematic server when it was previously working properly.

Administrators can also use Server Compliance in their IT environment to preemptively discover problems by performing routine (scheduled) audits. Routine audits identify changes made to installed software, hardware information, configuration files, system settings, and so on. As a best practice, regularly scheduled audits may significantly reduce the volume of noncompliant servers in your IT environment.

Service Compliance consists of several processes that allow administrators to compare objects across servers: performing an audit, creating a snapshot, and fixing a server that is malfunctioning. These processes and their results enable the following actions:

- Performing an ad hoc audit by comparing a reference server that is behaving normally to managed servers that are not behaving normally. The results from this audit allow administrators to see differences between servers that may affect their ability to operate properly.
- Auditing managed servers to determine whether critical files have been modified by comparing servers.
- Comparing a server snapshot to live servers.
- Comparing a snapshot to a snapshot.
- Packaging differences discovered between servers to create a uniform set of managed servers. See “Visual Packager” on page 363 in Chapter 9 for more information.

Server Objects

Server Compliance creates snapshots to record information about managed servers and server groups, and performs audits to compare server objects. Some objects are captured and audited live and some objects are captured from the Model Repository.

The following server objects are captured and audited live from a managed server:

- File System (files, directories, and symbolic links)
- Windows Registry
- Windows Services
- Windows COM
- Windows IIS Metabase

A Windows COM category (folder) that does not have any objects will not be included in a snapshot or audit, even though Opsware SAS will display an empty COM folder in the Server Explorer. A snapshot or audit will include a Windows COM category when it includes at least one object.

The following server objects are captured and audited from the Model Repository:

- Installed Packages
- Installed Patches
- Installed Hardware (CPU, network interfaces, storage, memory)

Server Compliance does not support device files, door files, and sockets.

Selection Criteria

This section discusses the following topics:

- Inclusion and Exclusion Rules
- Snapshot Selection Criteria Overlap

To create a server snapshot or perform an audit, you must provide instructions for Opsware SAS so that it knows what to collect and compare, how to collect and compare, and what files to include and exclude. In Server Compliance, these instructions are called selection criteria.

In Server Compliance, you can record and compare information about a managed server, such as hardware, packages, and patches that are installed on a server, and services that are running on a server. You can specify detailed selection criteria that allow you to select file system directories for a recursive snapshot or comparison into subdirectories, and to include or exclude files from the snapshot and audit processes. Selection criteria also allow you to record and compare the checksum of files, and store contents of files.

Opsware SAS provides selection criteria recommendations for certain operating systems. See "Selection Criteria for Snapshots" on page 789 in this chapter for more information.

You must have a set of permissions to manage selection criteria. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

The following selection criteria are used in the server snapshot and audit processes:

- **Content**: Specifies how to collect and compare server objects and certain types of object content, such as:
  - Content of subdirectories (Recursive)
  - Full contents of a file (Contents)
  - Windows NT Access Control List (ACL)
  - Checksum for each file (Checksum)

  For the Recursive option, you can specify content that applies to the entire file system or Windows Registry of the managed server. You can also specify content that applies only to selected subdirectories or files in the file system or Windows Registry. If you do not specify options for subdirectories or files, the default content is derived from the options that are set at the top level directory.

  For the Checksum option, you can select detailed options, depending on whether you want the snapshot or audit to include all of the content of the files (Full), only the first megabyte (MB) of each file (Lite), or no content (None). Checksum is performed only on the contents of a file; therefore, file owner and file permission changes are not audited. The default checksum is Full (all of the content of the files).

- **Inclusions/Exclusions**: Identifies the directories and files in the file system of the managed server you want included in and excluded from the snapshot or audit. See “Inclusion and Exclusion Rules” on page 302 in this chapter for more information.

- **Comparison Criteria**: Enables you to select whether you want the audit to compare the file modification date (Date), and the user and group access permissions (Access). For example, if two files have the same checksum, you can further compare these files by the file modification date comparison criteria. Comparison criteria is optional in the audit process. Comparison criteria does not apply to the snapshot process.

- **Override Default Options**: Enables you to select one or more individual directories (folders), files, or registry entries in the file system and then check Override Default Options to override the selection criteria that applies to the entire file system. This option allows you to specify different selection criteria for different directories and files. The selection criteria that is specified at the file system level (default options) will not be applied to these individual objects during the snapshot or audit process. The selection criteria you defined for these individual objects will be applied during the snapshot or audit process.

- **Clear All Overrides**: Enables you to select one or more individual directories (folders), files, or registry entries in the file system, remove options that were previously set, and return them to the default options.

For the Windows Registry object, only the Content tab is enabled and only the Recursive and full Contents options are available. Options are not available for Windows Services objects.

Inclusion and Exclusion Rules

To specify the directories and files that you want included in and excluded from a snapshot or an audit, you need to understand the rationale that Server Compliance applies. It is important to understand what the inclusion and exclusion rules are, and how Server Compliance applies these rules to the relative subset of the absolute path of the file. These concepts are best explained by providing examples and supporting text.

Server Compliance provides the following three types of inclusion and exclusion rules:

- A file-type rule applies to the file name path and contains neither a "/" or a "\".
- A relative-type rule applies to the relative path and can contain a "/" for Unix and a "\" for Windows, and is not fully qualified.
- An absolute-type rule applies to the absolute path. In Unix, an absolute path begins with a "/". In Windows, an absolute path begins with a volume letter that is followed by ":\" and is fully qualified, such as "C:\", "d:\", "f:\", and so on. If you use a "/" (forward slash) for Windows paths, Server Compliance will convert it to a "\" (backslash) to be able to use it as a valid path.

Server Compliance processes all exclusion rules first. After all exclusion rules are applied, then the inclusion rules are applied. The default for include is to include all objects in the file system. In many cases, inclusion rules may not even be processed because, combined with the exclusion rules (which occur first), they may become a moot point.

You can also use the asterisk (*) and the question mark (?) as valid wildcards in inclusion and exclusion rules. The wildcard character is a placeholder for matching to a path, or one or more characters.

Depending on the type of inclusion and exclusion rule, the rule is applied only to the relevant subset of the absolute path of the file. In Server Compliance, there is one top level for each snapshot or audit. Each file that you compare against the inclusion and exclusion rules has an absolute path. In Figure 8-1, the absolute path is /usr/home/abc/defg. A snapshot or an audit looks down the /usr/home/abc/defg absolute path and sees abc/defg as the relative path and defg as the file name. In this example, the inclusion and exclusion rules would apply in the following manner:

- A file-type rule applies to the file name path defg.
- A relative-type rule applies to the relative path abc/defg.
- An absolute-type rule applies to the absolute path /usr/home/abc/defg.

See Figure 8-1 for an illustration of how Server Compliance applies the inclusion and exclusion rules to a relative subset of the path of the file.

Figure 8-1: How Inclusion and Exclusion Rules Apply

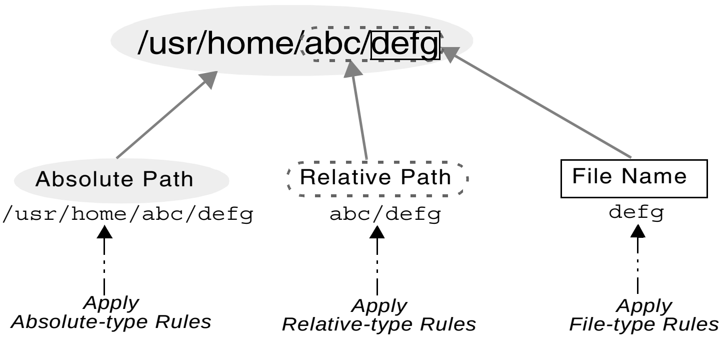

To best explain how these rules are applied, the following examples are provided.

The sample file system structure used in “Example: Including all files with the .txt extension in your snapshot or audit”, “Example: Including only the file a in your snapshot or audit”, and “Example: Including the last temp.txt file and excluding everything else” is:

```
/dir1/dir2/a
/dir1/dir2/b
/dir1/dir2/names.txt
/dir1/dir2/temp.txt
/dir1/dir2/version1.exe
/dir1/dir2/subdir/version2.exe
```

Example: Including all files with the .txt extension in your snapshot or audit

If you want to include all files with the .txt extension in your snapshot or audit, your inclusion and exclusion rules would be:

- /dir1/dir2
- include *.txt (This is a file-type rule.)
- exclude * (This is a file-type rule.)

The following steps explain how Server Compliance iterates through the file structure and applies any corresponding inclusion and inclusion rules:

1. The * causes /dir1/dir2/a to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/a (a) and there is no match. The file is not included.
2. The * causes /dir1/dir2/b to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/b (b) and there is no match. The file is not included.
3. The * matches names.txt, but *.txt matches names.txt as well, which causes the file to be included.
4. Same as step 3.
5. Compare a to *, which is a match; compare a to a, which is a match. The file is included.
6. Compare b to *, which is a match; compare b to a which is not a match. The file is excluded.

These step numbers correspond to the paths in the sample file structure, with the numbering starting with the top-level path.

Example: Including only the file a in your snapshot or audit

If you want to include only the file a in your snapshot or audit, your inclusion and exclusion rules would be:

- /dir1/dir2
- exclude * (This is a file-type rule.)
- include a (This is a file-type rule.)

The following steps explain how Server Compliance iterates through the file structure and applies any corresponding inclusion and inclusion rules:

1. The * causes /dir1/dir2/a to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/a (a) and there is no match. The file is not included.
2. The * causes /dir1/dir2/b to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/b (b) and there is no match. The file is not included.

3. The * matches names.txt, but *.txt matches names.txt as well, which causes the file to be included.
4. Same as step 3.
5. Compare a to *, which is a match; compare a to a, which is a match. The file is included.
6. Compare b to *, which is a match; compare b to a which is not a match. The file is excluded.

These step numbers correspond to the paths in the sample file structure, with the numbering starting with the top-level path.

Example: Including the last temp.txt file and excluding everything else

If you want to include the last temp.txt file and exclude everything else in your snapshot or audit, your inclusion and exclusion rules would be:

- /dir1/dir2
- exclude * (This is a file-type rule.)
- include dir3/temp.txt (This is a relative-type rule.)

The following steps explain how Server Compliance iterates through the file structure and applies any corresponding inclusion and inclusion rules:

1. The * causes /dir1/dir2/a to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/a (a) and there is no match. The file is not included.
2. The * causes /dir1/dir2/b to be excluded. Then *.txt is applied against the file portion of /dir1/dir2/b (b) and there is no match. The file is not included.
3. The * matches names.txt, but *.txt matches names.txt as well, which causes the file to be included.
4. Same as step 3.
5. dir3/temp.txt is dir3/temp.txt is compared against the relative portion of /dir1/dir2/dir3/temp.txt and there is a match.
6. Compare a to *, which is a match; compare a to subdir/version2.exe, which is not a match. The file is excluded.

These step numbers correspond to the paths in the sample file structure, with the numbering starting with the top-level path.

Snapshot Selection Criteria Overlap

When you include a parent directory (with options) in the selection criteria and a child directory (with different options) as additional selection criteria, the parent directory snapshot and the child directory snapshot will overlap each other as one snapshot. This logic also applies to Windows NT ACL collection and content collection options, and Windows Registry content collection options. How selection criteria for a parent and child directory will overlap is best explained by the following examples.

Consider the following file system, where an ending forward slash (/) represents a directory:

```
/cust/app/bin/
/cust/app/bin/file1
/cust/app/bin/conf/
/cust/app/bin/conf/conf1
/cust/app/bin/conf/conf2
/cust/app/bin/conf/dev/
/cust/app/bin/conf/dev/conf3
```

Example A

If you create a snapshot using the following two selection criteria:

Directory `/cust/app/bin` (recursive, no checksum)

Directory `/cust/app/bin/conf` (not recursive, checksum)

The snapshot will record the following file system information:

```
/cust/app/bin/ (directory)
/cust/app/bin/file1 (no checksum)
/cust/app/bin/conf/ (directory)
/cust/app/bin/conf/conf1 (*checksum*)
/cust/app/bin/conf/conf2 (*checksum*)
/cust/app/bin/conf/dev/ (directory)
/cust/app/bin/conf/dev/conf3 (no checksum)
```

As you can see, even though `/cust/app/bin` was recursive and had no checksum, the `/cust/app/bin/conf` directory overrode it and all files in that directory have checksums recorded for them.

Example B

If you create a snapshot using the following two selection criteria (by switching the options used in Example A):

Directory `/cust/app/bin` (recursive, checksum)

Directory `/cust/app/bin/conf` (not recursive, no checksum)

The snapshot will record the following file system information:

```
/cust/app/bin/ (directory)
/cust/app/bin/file1 (checksum)
/cust/app/bin/conf/ (directory)
/cust/app/bin/conf/conf1 (*no checksum*)
/cust/app/bin/conf/conf2 (*no checksum*)
/cust/app/bin/conf/dev/ (directory)
/cust/app/bin/conf/dev/conf3 (checksum)
```

Example C

If you create a snapshot using the following three selection criteria (by adding a file option):

Directory `/cust/app/bin` (recursive, checksum)

Directory `/cust/app/bin/conf` (not recursive, no checksum)

File `/cust/app/bin/conf/conf1` (checksum)

The snapshot will record the following file system information:

```
/cust/app/bin/ (directory)
/cust/app/bin/file1 (checksum)
/cust/app/bin/conf/ (directory)
/cust/app/bin/conf/conf1 (*checksum*)
/cust/app/bin/conf/conf2 (no checksum)
/cust/app/bin/conf/dev/ (directory)
/cust/app/bin/conf/dev/conf3 (checksum)
```

In this example, the very detailed selection criteria for `conf1` override the `/cust/app/bin/conf` selection criteria.

Snapshots

A snapshot is a record of how an Opsware-managed server is configured at a certain date and time. This record is intended to capture information about the state of a known working server. By creating a server snapshot, you can identify attributes of files and directories on a managed server as a baseline record. A snapshot can also be considered as a way to back up a managed server, especially if you plan to make changes to the server and want to keep a record of it before you change anything. You can schedule when you want a snapshot to be created (either once or as a recurring job) and who you want to receive email notification about the status of the job.

In addition to recording information about objects on managed servers, a snapshot can contain the content of some objects. A server snapshot also identifies attributes of other objects on specific types of operating systems, such as the Windows registry and Windows services.

In Server Compliance, a snapshot can be used as a source and a target in the auditing process. See "Auditing" on page 334 in Chapter 8 for more information.

The snapshot process records information about objects on a managed server at a particular point in time. The captured information can also include content of some objects and object attributes on specific types of operating systems, such as the Windows registry and Windows services.

Server Compliance uses snapshots in the audit process to compare Opsware-managed servers to determine how objects, and some object content, may differ.

Each snapshot requires a snapshot template that specifies a target and selection criteria. A target is a managed server or server group that you are recording information about. Selection criteria instruct Opsware SAS to record information about a managed server.

Depending on the server objects that you select for the snapshot, you can set up options for a more detailed snapshot that includes file content and checksum, file inclusions and exclusions, and Windows Registry key options. Figure 8-2 illustrates the server snapshot process.

Figure 8-2: The Server Snapshot Process

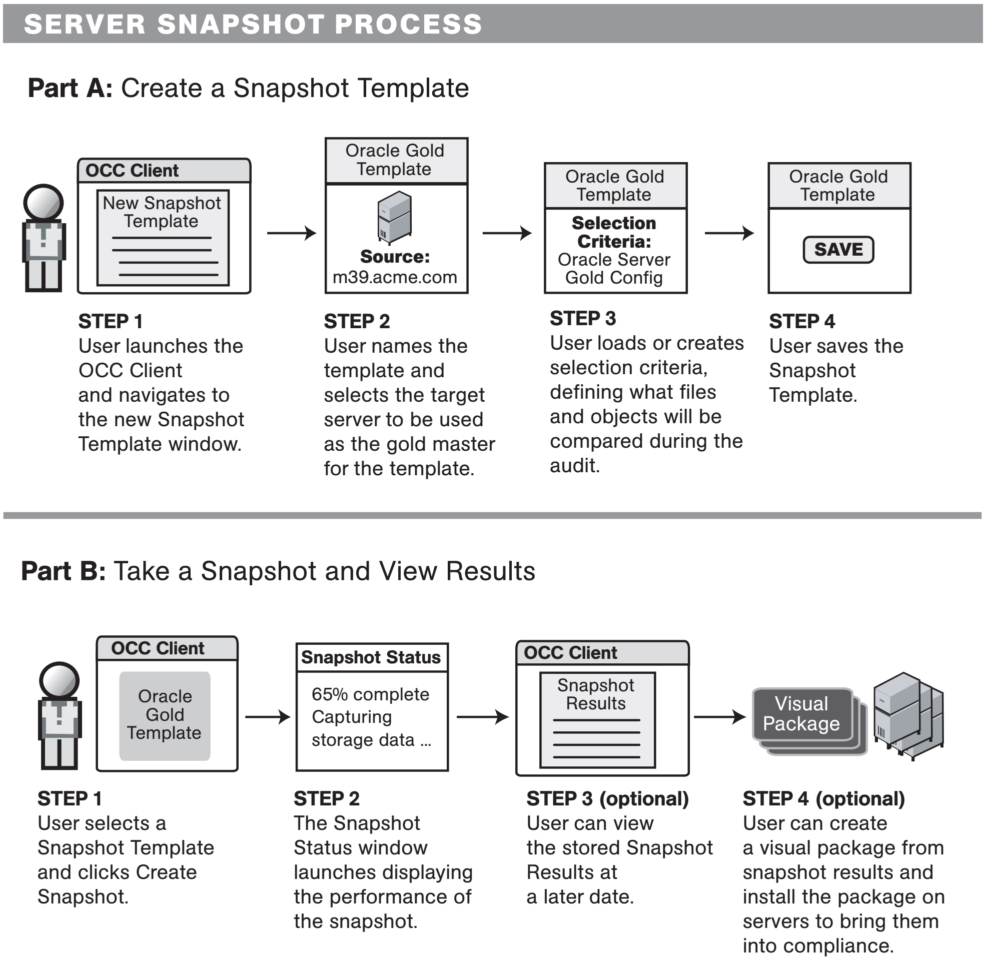

SERVER SNAPSHOT PROCESS

Part A: Create a Snapshot Template

STEP 1
User launches the OCC Client and navigates to the new Snapshot Template window.

STEP 2
User names the template and selects the target server to be used as the gold master for the template.

STEP 3
User loads or creates selection criteria, defining what files and objects will be compared during the audit.

STEP 4
User saves the Snapshot Template.

Part B: Take a Snapshot and View Results

Snapshot Status window:

65% complete
Capturing
storage data ...

STEP 1
User selects a Snapshot Template and clicks Create Snapshot.

STEP 2
The Snapshot Status window launches displaying the performance of the snapshot.

STEP 3 (optional)
User can view the stored Snapshot Results at a later date.

STEP 4 (optional)
User can create a visual package from snapshot results and install the package on servers to bring them into compliance.

Ways to Create a Snapshot

In Server Compliance, you can create a snapshot to record information that is derived from an Opsware-managed server and launch the snapshot from the following different windows in the OCC Client.

Launching a Snapshot from All Managed Servers

1. Launch the OCC Client. From the Navigation pane, select Servers and then select All Managed Servers.
2. From the Content pane, select one or more managed servers.
3. Select **Actions ➤ Create Snapshot**.

Launching a Snapshot from Server Groups

1. Launch the OCC Client. From the Navigation pane, select Servers and then select Server Groups.
2. From the Content pane, select a server group.
3. Select **Actions ➤ Create Snapshot**.

Launching a Snapshot from the Server Explorer

1. Launch the OCC Client. From the Navigation pane, select Servers and then select All Managed Servers.
2. From the Content pane, select a managed server, and then open it.
3. From the Server Explorer window, select **Actions ➤ Create Snapshot**.

Launching a Snapshot from the Compliance Window

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select **Compliance**.
2. Select the Snapshot Template Tab.
3. Select a Snapshot Template and then select **Actions ➤ Create Snapshot**.

Snapshot Templates

This section discusses the following topics:

- Creating a Snapshot Template
- Setting Up Snapshot Selection Criteria
- Creating a Snapshot
- Saving Snapshot Selection Criteria for Reuse
- Loading Snapshot Selection Criteria
- Managing Your Snapshot Selection Criteria
- Editing a Snapshot Template
- Copying a Snapshot Template for Reuse
- Deleting a Snapshot Template

A snapshot template identifies the information (selection criteria) you want to capture about the state and configuration of a managed server or server group (target), at a particular point in time. A snapshot template must define selection criteria and one or more targets.

You can modify snapshot templates that you originally created, and copy and re-use them to routinely record information about the managed servers in Opsware SAS. To conserve disk space, you can remove snapshot templates that you no longer need from the Software Repository.

Since selection criteria are specific to an operating system, such as Unix or Windows, you cannot have one (universal) template for different operating systems of managed servers in Opsware SAS. Figure 8-3 illustrates a sample snapshot template that was used to gather and record information about a Windows 2003 server.

You must have a set of permissions to create and modify snapshot templates. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

Figure 8-3: Sample Snapshot Template

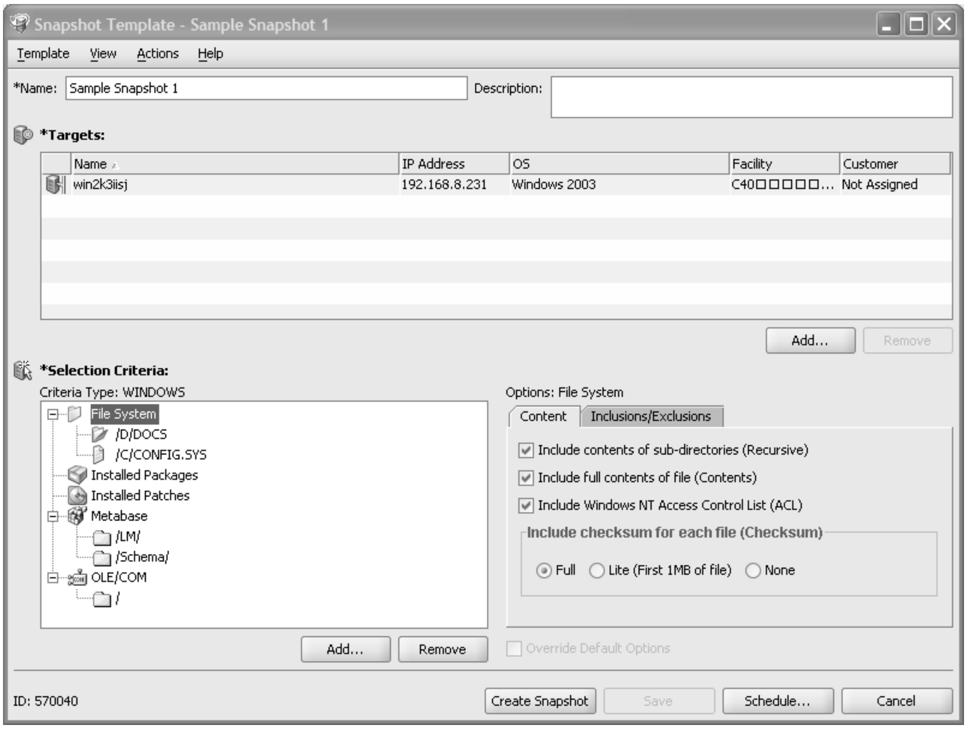

*Name: Sample Snapshot 1

Description:

*Targets:

Name	IP Address	OS	Facility	Customer
win2k3iisj	192.168.8.231	Windows 2003	C40□□□□□...	Not Assigned

*Selection Criteria:

Criteria Type: WINDOWS

- File System
  - /D/DOCS
  - /C/CONFIG.SYS
- Installed Packages
- Installed Patches
- Metabase
  - /LM/
  - /Schema/
- OLE/COM
  - /

Options: File System

Content | Inclusions/Exclusions

☑ Include contents of sub-directories (Recursive)

☑ Include full contents of file (Contents)

☑ Include Windows NT Access Control List (ACL)

Include checksum for each file (Checksum)

● Full ○ Lite (First 1MB of file) ○ None

□ Override Default Options

ID: 570040

Creating a Snapshot Template

You must create a snapshot template to identify your targets and selection criteria for the snapshot process. An asterisk (*) at the beginning of a field name indicates that this information is required.

To create a snapshot template, perform the following steps:

1. From one of the starting points described in "Ways to Create a Snapshot" on page 310, select **Create Snapshot**. The Snapshot Template window opens.

2. In the Name field, enter a name for the snapshot you are creating. This name must be unique in your Opsware System.

3. (Optional) In the Description field, enter a brief description of the type of snapshot that you are creating.

4. The Targets field displays the name, IP address, and operating system of the managed server that you selected for the snapshot.

   - If you want to add to this list of targets, click **Add** and select additional servers and server groups from the Target Chooser window. See “Target Chooser” on page 42 in Chapter 3 for more information.

   🔍 Use the search tool to dynamically filter this list by entering a server name, IP address, or operating system.

   - If you want to remove targets from this list, select one or more servers or server groups and then click **Remove**.

5. Select your selection criteria in the right field, and click **Add** to move it into Criteria Type window. See “Setting Up Snapshot Selection Criteria” on page 313 for more information.

6. When you are finished selecting sources, click the **Add** button that is below the Selection Criteria field.

7. Click **Create Snapshot** to launch the snapshot process and display a job status window that indicates the progress.

8. Click **Save** to save your snapshot template in the Model Repository or click **Cancel** to close this window without saving your changes.

Snapshots are stored in the Software Repository.

Setting Up Snapshot Selection Criteria

To define the selection criteria for the server snapshot process, perform the following steps:

1. From one of the starting points described in “Ways to Create a Snapshot” on page 310, select **Create Snapshot**. The Snapshot Template window opens.

2. To specify your selection criteria, click **Add** below the Selection Criteria field.

3. In the Selection Criteria Editor, select a source and click **Add**. You can also remove Selection Criteria, by selecting the sources and clicking **Remove**.

To include all COM objects in a snapshot of a Windows server, you must select the OLE/COM top-level node in the Selection Criteria Editor. Opsware SAS categorizes the COM objects based on an attribute of the object, where the COM object specifies zero or more categories. Opsware SAS displays all COM objects in one node (OLE/COM) in the server object tree in case you do not know the category or the object did not specify one.

In the Selection Criteria Editor you can select a Windows COM category (folder) whether it contains objects or not. The Selection Criteria Editor and the Server Explorer display all Windows COM folders, whether they are empty or not.

When finished, click the **Add** button that is below the Selection Criteria field. This returns you to the Snapshot Template window.

4 In the Options pane, select the Content tab to specify whether the snapshot will include contents of subdirectories (Recursive), full contents of a file (Contents), Windows NT Access Control List (ACL), and checksum for each file (Checksum). You can also select more detailed Selection Criteria for checksum, you can have the snapshot include all of the content of the files (Full) or only the first megabyte (MB) of each file (Lite). Checksum is performed on the contents of a file. File owner and file permission changes are not included in a snapshot. The default for checksum is Full.

You can specify snapshot content that applies to the entire file system or Windows Registry of the managed server or you can specify snapshot content that applies only to selected subdirectories or files in the file system or Windows registry. If you do not specify options for subdirectories, the default snapshot content is the parent directory. See "Snapshot Selection Criteria Overlap" on page 306 in this chapter for more information.

5 If you have selected objects that have subfolders, select the Inclusions/Exclusions tab. Select **Include** or **Exclude** to specify the directories and files that you want included in and excluded from the snapshot.

- Select **Include** and enter the path or file name in the text box. Click **Add** to display this information in the list of content that will be included in the snapshot.
- Select **Exclude** and enter the path or file name in the text box. Click **Add** to display this information in the list of content that will be excluded from the snapshot.

See “Inclusion and Exclusion Rules” on page 302 in this chapter for more information.

6 (Optional) Select an individual directory or file in the Criteria Type field. Check Override Default Options to override the selection criteria that apply to the entire file system. This option allows you to specify different selection criteria for different directories and files. The selection criteria specified at the file system level (default options) will not be applied to these individual objects during the snapshot process. The selection criteria you defined for these individual objects will be applied during the snapshot process.

For the Windows Registry object, only the Content tab is enabled and only the Recursive and full Contents options are available. Options are not available for Windows Services objects.

7 Click **Save** to save this snapshot template in the Model Repository or click **Cancel** to close this window without saving your changes.

Creating a Snapshot

To create a snapshot, you must run a snapshot template. In Server Compliance, you can create a snapshot in several different ways. You can:

- Create and save a snapshot template, and then create a snapshot.
- Create a snapshot template and then create a snapshot without saving the template. This is known as creating an ad hoc snapshot.
- Select an existing snapshot template and then create the snapshot.
- Edit an existing snapshot template, save your changes, and then create the snapshot.
- Edit an existing snapshot template and then create the snapshot without saving the changes you made to the template.
- Select multiple managed servers, or modify an existing and perform an ad hoc snapshot.

See “Creating a Snapshot Template” on page 312 in this chapter for more information and “Editing a Snapshot Template” on page 318 in this chapter for more information.

To create a snapshot, you must have read permissions for all targets (servers and server groups) that are specified in the snapshot template. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

To create a snapshot and track its progress, perform the following steps:

1. From one of the starting points described in "Ways to Create a Snapshot" on page 310, select **Create Snapshot**. The Snapshot Template window opens.
2. To check the status of a snapshot process, select Jobs and Sessions from the Navigation pane, and then select Job Logs.
3. In the Content pane, select Create Snapshot from the drop-down list.
4. Click **Update**. The status of snapshots will now display. For more information, see "Viewing a Snapshot Job" on page 333.

To prevent runaway processes, the snapshot process will time-out if it exceeds 60 minutes or if the data that is collected from a managed server exceeds 1 gigabyte (GB). If you specify that you want to collect the full contents of files in the selection criteria, the data collected may exceed the maximum size that can be successfully recorded in a snapshot.

Saving Snapshot Selection Criteria for Reuse

You can save selection criteria and use it in other snapshot templates.

To save your selection criteria, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Snapshot Templates tab.
3. Select the template that you want to make a copy of, and then open it.
4. In the Snapshot Template window, select **Actions ➤ Save Selection Criteria As**.
5. In the Save Selection Criteria As window, enter a unique name.
6. (Optional) Enter a description of the Selection Criteria.

7. Click **Save** to save your Selection Criteria or click **Cancel** to close this window without saving your changes.

Loading Snapshot Selection Criteria

When you load previously-saved snapshot selection criteria, you do not create an association between the selection criteria and the snapshot template.

To load selection criteria, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Snapshot Templates tab.
3. Select the template you want to make a copy of, then open it.
4. In the Snapshot Template window, select **Actions ➤ Load Selection Criteria**.
5. In the Load Selection Criteria window, select a selection criteria and click **Load** to load your selection criteria or click **Close** to close this window.

Managing Your Snapshot Selection Criteria

To view, rename, copy, or delete selection criteria, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Snapshot Templates tab.
3. Select the template you want to manage and then open it.
4. In the Snapshot Template window, select **Actions ➤ Manage Selection Criteria**.
5. In the Manage Selection Criteria window, select a selection criteria name and click any of the following buttons:

    - **Open**: To view detailed information about the selection criteria, including the date and time it was last modified and by whom.
    - **Rename**: To change the name of the selection criteria.
    - **Save As**: To make a copy of (clone) the selection criteria.
    - **Delete**: To delete selection criteria. You can select multiple selection criteria names to perform a mass delete.

Editing a Snapshot Template

If you originally created a snapshot template, you can modify the selection criteria and targets.

To modify a snapshot template, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Snapshot Templates tab.
3. Select the template that you want to edit and then open it.
4. Make changes to the target names, snapshot description, or selection criteria as needed, and click **Save**. This saves changes to the current template and will overwrite the previous version of the template.
5. (Optional) You can also perform any of the following actions:
    1. Click **Create Snapshot** to create an another snapshot.
    2. Select **Actions ➤ Save Template As** to save your snapshot template using a different name.
    3. Click **Cancel** if you want to close this window without saving your changes.

Saving Snapshot Selection Criteria for Reuse

You can save snapshot selection criteria to reuse it, or you can modify it, save it, and then reuse it to record objects on different targets (managed servers and server groups). As a best practice, this is something you should do if this is the type of snapshot that you need to create frequently. You can save selection criteria from a new template that you are creating or from an existing template to reuse it in other snapshot templates.

To save selection criteria that is in a snapshot template, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Snapshot Templates tab.
3. Select the template you want to edit and then open it.
4. Make changes to the target names, snapshot description, and selection criteria as desired.
5. Select **Actions ➤ Save Selection Criteria As**.

6. In the Save Selection Criteria As window, enter a unique name.

7. (Optional) Enter a description of the selection criteria.

8. Click **Save** to save your selection criteria or click **Cancel** to close this window without saving your changes.

9. See “Loading Snapshot Selection Criteria” on page 317 in this chapter for information on loading your saved selection criteria.

Copying a Snapshot Template for Reuse

You can copy a snapshot template and reuse the selection criteria to record information about other targets. To copy a snapshot template, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.

2. From the Content pane, select the Snapshot Templates tab.

3. Select the template you want to make a copy of, and then open it.

4. In the Snapshot Template window, select **Actions ➤ Save Template As**.

5. In the Save Template As window, enter a unique name.

6. (Optional) Enter a description of the template.

7. Click **Save** to save the snapshot template in the Model Repository or click **Cancel** to close this window without saving your changes.

Deleting a Snapshot Template

To conserve disk space, you should delete snapshot templates that you no longer need from the Model Repository. You can delete an existing snapshot template only if you originally created it.

To delete an snapshot template, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.

2. From the Content pane, select the Snapshot Templates tab.

3. Select one or more templates and then select **Actions ➤ Delete**.

4. In the Confirmation Dialog, click **Yes** to delete this snapshot template or click **No** if you do not want to delete it.

When you delete a snapshot template, you do not delete any of the snapshots that were created from it. See "Deleting a Snapshot" on page 333 in this chapter for more information. However, when you delete a snapshot template, all schedules that you own (created), that are also associated with that snapshot template, will be deleted. See "Snapshot Scheduling" on page 320 in this chapter for more information.

Snapshot Scheduling

This section discusses the following topics:

- Scheduling a Snapshot
- Viewing a Snapshot Schedule
- Editing a Snapshot Schedule
- Deleting a Snapshot Schedule

A snapshot schedule specifies when you want Opsware SAS to create a snapshot (either once or as a recurring job) and who you want to receive email notification about the status of the job. You can also view, edit, and delete existing snapshot schedules. When you delete a snapshot template, all schedules that you own (created), that are also associated with that snapshot template, will be deleted.

You must have permissions to create, view, edit, and delete snapshot schedules. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

Scheduling a Snapshot

You can define a schedule that specifies when you want a snapshot created after you have completed the Snapshot Template and saved it. To schedule a snapshot job, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and select Compliance.
2. From the Content pane, select the Snapshot Templates tab and then select a snapshot and open it. The Snapshot Template window opens.

3 In the Snapshot Template window, click **Schedule**. The Schedule Job window will display.

Figure 8-4: Snapshot Schedule

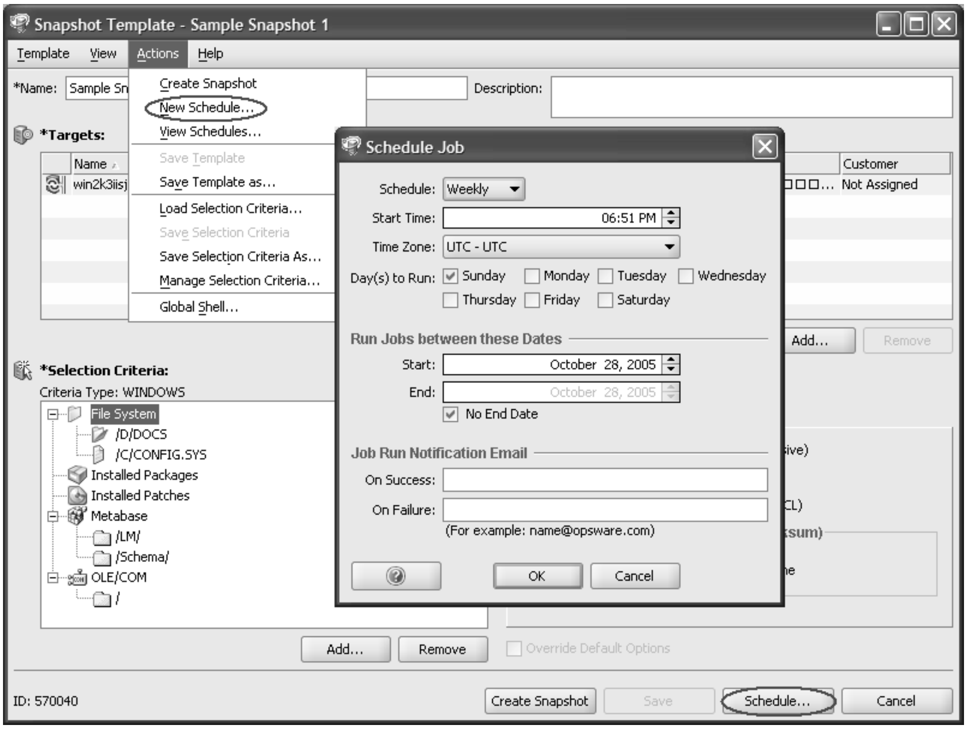

Snapshot Template - Sample Snapshot 1

*Name: Sample Sn

Description:

*Targets:

Name	Customer
win2k3iisj	Not Assigned

*Selection Criteria:

Criteria Type: WINDOWS

- File System
  - /D/DOCS
  - /C/CONFIG.SYS
- Installed Packages
- Installed Patches
- Metabase
  - /LM/
  - /Schema/
- OLE/COM
  - /

□ Override Default Options

ID: 570040

Schedule Job

Schedule: Weekly

Start Time: 06:51 PM

Time Zone: UTC - UTC

Day(s) to Run: ☑ Sunday ☐ Monday ☐ Tuesday ☐ Wednesday ☐ Thursday ☐ Friday ☐ Saturday

Run Jobs between these Dates

Start: October 28, 2005

End: October 28, 2005

☑ No End Date

Job Run Notification Email

On Success:

On Failure:

(For example: name@opsware.com)

4 In the Schedule field, choose whether you want the snapshot to be created once, daily, weekly, monthly, or on a custom schedule.

When you initially specify a snapshot schedule, the Schedule field default is Weekly. If you check every day of the week when the Schedule field is set to Weekly, the Schedule field changes to Daily when you subsequently edit the snapshot schedule.

5 Depending on the type of Schedule you selected, use the following guidelines to enter a Start Time that specifies when you want the job to begin:

- To create a snapshot only once, enter the complete date and time, such as October 5, 2005 12:01AM.
- To create a snapshot daily, weekly, or monthly, enter the time of day, such as 12:01AM.

- To create a snapshot on a custom schedule, use the following syntax to enter a crontab string that specifies the day, date, and time you want the job to begin, as shown in Table 8-5:

Figure 8-5: Crontab

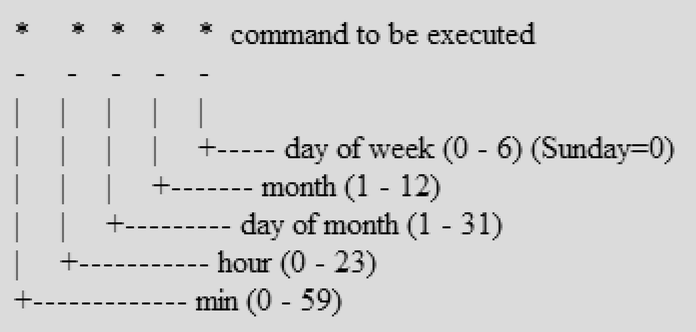

```
*    *  *  *   *  command to be executed
-    -  -  -   -
|    |  |  |   |
|    |  |  |   +----- day of week (0 - 6) (Sunday=0)
|    |  |  +------- month (1 - 12)
|    |  +--------- day of month (1 - 31)
|    +----------- hour (0 - 23)
+------------- min (0 - 59)
```

An asterisk (*) in any of these fields represent all days of the month, all months of the year, all days of the week, and so on. For example, the following crontab string will create the snapshot at midnight every weekday:

0 0 * * 1-5

The crontab string can include serial (1,2,3,4) and range (1-5) values. For more information about crontab entry formats, consult the Unix man pages.

6. Select a default time zone for the job execution time, or accept the default time zone. The default time shown converts the scheduled time to the time zone set in your user preferences. If you do not set a preferred time zone, the time zone is derived from the Opsware SAS core server (typically UTC).

7. Depending on which Schedule you selected, complete one of the following steps:

   - If you selected Weekly, select the days of the week that you want the snapshot to be created in the Days to Run field.

   - If you selected Monthly, select the Day of the month to create the snapshot. You must also select the months of the year to create the snapshot in the Months to Run field.

8. In the Run Jobs Between These Dates section, select a Start and End time to specify when you want the snapshot process to begin and end.

9. (Optional) Select No End Date if you do not want the snapshot schedule to expire.

10 (Optional) In the Job Run Notification Email section, enter email addresses (separated by a comma or a space) of those you want to receive information on the status of the snapshot job, including IDs of servers affected. Enter email addresses in one or both of the following fields:

On Success: Email addresses that will receive notifications of jobs that completed successfully, such as name@opsware.com, name2@opsware.com.

On Failure: Email addresses that will receive notifications of jobs that failed to complete, such as name@opsware.com, name2@opsware.com.

11 Click **OK** to save your schedule information or **Cancel** to close this window without saving any information.

Viewing a Snapshot Schedule

You can view a snapshot schedule after you have created (or edited) and saved it. To view a snapshot schedule, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Software Library and select Compliance.

2 From the Content pane, select the Snapshot Templates tab and then select a snapshot and open it. The Snapshot Template window opens.

3. In the Snapshot Template window, select **Actions** ➤ **View Schedules** or click **Schedule** to display the Scheduled Jobs window.

Figure 8-6: Scheduled Snapshots

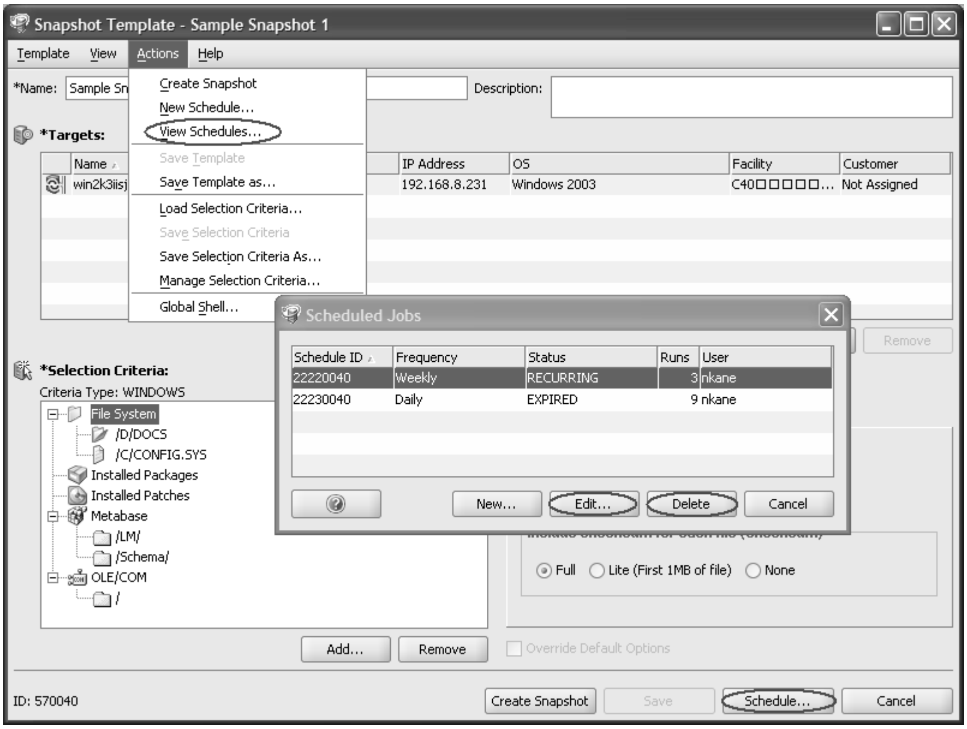

Snapshot Template - Sample Snapshot 1

*Name: Sample Sn

Description:

*Targets:

Name	IP Address	OS	Facility	Customer
win2k3iisj	192.168.8.231	Windows 2003	C40□□□□□□...	Not Assigned

*Selection Criteria:

Criteria Type: WINDOWS

```
File System
    /D/DOCS
    /C/CONFIG.SYS
Installed Packages
Installed Patches
Metabase
    /LM/
    /Schema/
OLE/COM
    /
```

Scheduled Jobs

Schedule ID	Frequency	Status	Runs	User
22220040	Weekly	RECURRING	3	nkane
22230040	Daily	EXPIRED	9	nkane

○ Full ○ Lite (First 1MB of file) ○ None

□ Override Default Options

ID: 570040

4. Open the snapshot job that you want to review.

5. In the Scheduled Jobs window, review the detailed schedule information.

6. Click **Cancel** to close this window.

Editing a Snapshot Schedule

You can edit a snapshot schedule after you have created and saved it. To edit a snapshot schedule, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and select Compliance.

2. From the Content pane, select the Snapshot Templates tab and then select a snapshot and open it. The Snapshot Template window opens.

3. In the Snapshot Template window, select **Actions** ➤ **View Schedules** or click **Schedule** to display the Scheduled Jobs window. See Figure 8-6.

4. Select the Schedule that you want to change and click **Edit**. Make your schedule changes.
5. Click **OK** to save your changes or **Cancel** to close this window without saving your changes.

Deleting a Snapshot Schedule

You can delete a saved snapshot schedule. To delete a snapshot schedule, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and select Compliance.
2. From the Content pane, select the Snapshot Templates tab and then select a snapshot and open it. The Snapshot Template window opens.
3. In the Snapshot Template window, select **Actions ➤ View Schedules** to display the Scheduled Jobs window. See Figure 8-6.
4. Select the snapshot schedule you want to delete and then click **Delete**.
5. In the Confirmation Dialog, click **Yes** to delete the schedule or click **No** if you do not want to delete it.
6. (Optional) Click **Cancel** to close this window without deleting the selected snapshot schedule.

Uses of Snapshots

In Server Compliance, you can use snapshots in several different ways. You can:

- Browse snapshot contents to view detailed information about the objects on a managed server.
- Convert a snapshot to a general snapshot to disassociate it from the managed server.
- Copy file system objects from a snapshot to a managed server.
- Compare a snapshot with servers that are not performing well to discover whether objects on the noncompliant server are missing or out-of-date.
- Delete snapshots that are old or no longer required to conserve disk space.

Ways to Find a Snapshot

After you have created a snapshot, you can find it using different windows in the OCC Client.

In the Compliance window:

1. From the Navigation pane, select Software Library and then select Compliance.
2. Select the Snapshots tab.
3. Select a snapshot and then open it.

In the Jobs and Sessions window:

1. From the Navigation pane, select Jobs and Sessions and then select Job Logs.
2. In the Content pane, select Create Snapshot from the Job Types drop-down list.
3. (Optional) Specify a job ID, date and time range, and job status options to narrow your search for a snapshot job.
4. Click **Update** to view a list of snapshots based on your search criteria.
5. Select a snapshot in the list and then open it.

In the Server Explorer:

1. From the Navigation pane, select Servers.
2. Select All Managed Servers and then select a server from the Content pane.
3. Select a server and then open it.
4. In the Server Explorer window, select Compliance and then select the Snapshots tab. All the snapshots for that server or server group will display.
5. Select a snapshot and then open it.

Browsing Contents of a Snapshot

You can browse the contents of a snapshot and view detailed information about the server objects that were recorded.

To browse the contents of a snapshot, perform the following steps:

1. From one of the starting points described in “Ways to Find a Snapshot” on page 326, open a snapshot.

Figure 8-7: Sample Snapshot Browser of a Windows Server

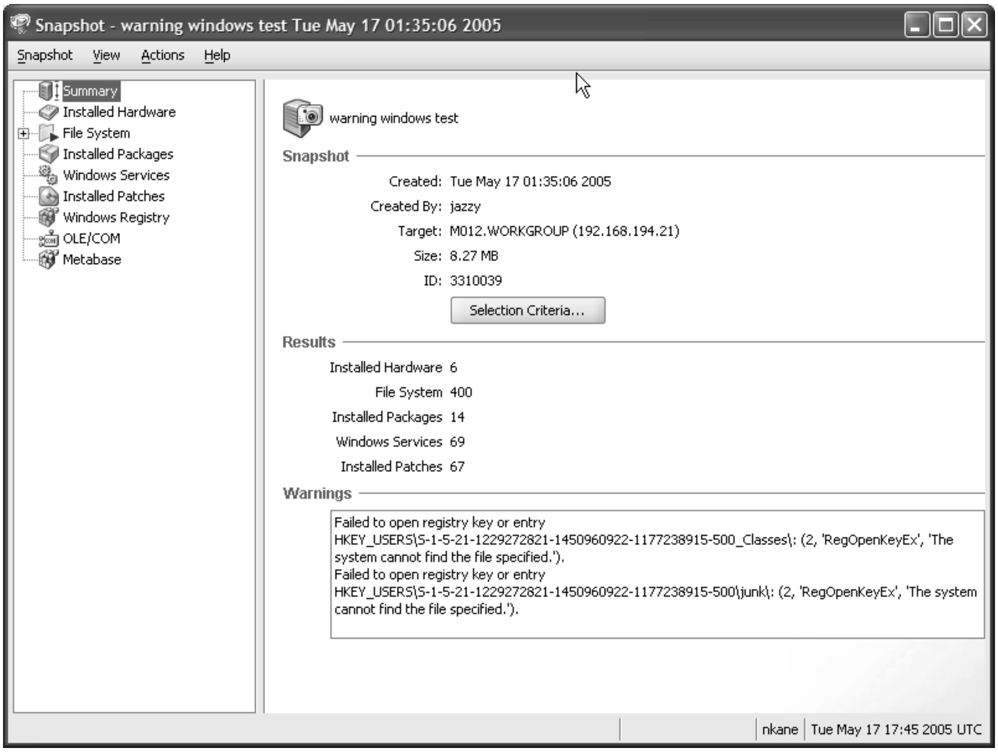

warning windows test

Snapshot

Created: Tue May 17 01:35:06 2005
Created By: jazzy
Target: M012.WORKGROUP (192.168.194.21)
Size: 8.27 MB
ID: 3310039

Results

Installed Hardware	6
File System	400
Installed Packages	14
Windows Services	69
Installed Patches	67

Warnings

```
Failed to open registry key or entry
HKEY_USERS\S-1-5-21-1229272821-1450960922-1177238915-500_Classes\: (2, 'RegOpenKeyEx', 'The
system cannot find the file specified.').
Failed to open registry key or entry
HKEY_USERS\S-1-5-21-1229272821-1450960922-1177238915-500\junk\: (2, 'RegOpenKeyEx', 'The system
cannot find the file specified.').
```

2. In the Snapshot Browser window, you can select:

   - **Summary** to view general information about a snapshot, such as the date and time the snapshot was created and by whom, the snapshot source (name of the managed server), the size of the snapshot file, and a snapshot ID number.

   - **Installed Hardware** to view information about the type of CPU processor and speed, cache size, memory size for SWAP and RAM, and storage devices that were recorded in the snapshot.

- **Installed Patches** to view information about the installed patches that were recorded in the snapshot, such as the patch type.
- **Installed Packages** to view information about the installed packages that were recorded in the snapshot, such as package type, package version, and release number.
- **File System** to view the directories, file properties and attributes, and contents of files recorded in the snapshot. See "Browsing File Properties, Permissions, and Content" on page 329 in this chapter for more information.

If a file in the snapshot exceeds 2MB in file size, Server Compliance cannot display the file contents.

- **Windows Services** to view information about the running services recorded in a snapshot, such as the name, description, startup state, startup type, and log on account. See "Viewing a Snapshot Job" on page 333 in Chapter 8 for more information.
- **Windows Registry** to view information about Windows registry entries in the snapshot, such as the registry key, registry value, and subkey. A registry key is a directory that contains registry values, where registry values are similar to files within a directory. A subkey is similar to a subdirectory. The content area in this window excludes subkeys. Server Compliance supports the following Windows registry keys: HKEY_CLASSES_ROOT, HKEY_CURRENT_CONFIG, HKEY_LOCAL_MACHINE, and HKEY_USERS.

  Valid control characters for the contents of the key entry (Data) include: #x9, #xA, [#xD, #x20-#xD7FF], [#xE000-#xFFFD], and [#x10000-#x10FFFF]. Invalid control characters cannot be stored by Opsware SAS and will be converted to XML entities and will display as &#;. For example, if the data value is 00 00 (in bytes), &#x00; will display in this column.
- **OLE/COM** to view information about Windows COM (Component Object Model) objects in the snapshot, such as the name and GUID (Globally Unique Identifier) of the object, and the path to the in-process server DLL.

  Opsware SAS provides warning messages that explain how Windows COM folders were processed. The following scenarios apply:

  When you create a snapshot where you selected a Windows COM folder that does not contain any objects, the Snapshot Browser displays a summary. Opsware SAS displays a warning that the GUID (Globally Unique Identifier) for that folder is invalid, which means that the Windows COM folder does not contain any objects.

When you create a snapshot where you selected a Windows COM object that does not exist on a target, Opsware SAS displays a warning that the folder is invalid.

When you create a snapshot where you selected a Windows COM folder that does not contain any objects and a Windows COM folder that does contain objects, the Snapshot Browser displays the folder. Opsware SAS displays a warning that the folder is empty.

- **Metabase** to view information about IIS Metabase objects in the snapshot, such as the ID, name, path, attributes, and data of the object.

Browsing File Properties, Permissions, and Content

You can view the properties and permissions of directories and files, and contents of files, that were recorded in a snapshot. Properties include the file path, file type, file size, checksum, user and group permissions, and so on. Permissions are displayed by their corresponding user and user group, such as Full Control, Modify, Read, Write, and so on.

For a snapshot of a Unix file system, you can also view file modes (access permissions), such as rw-r--r--.

For a snapshot of a Windows file system, you can also view file attributes, such as Read-only, Hidden, System, Directory, Archive, Device, Normal, Temporary, Sparse File, Reparse Point, Compressed, Offline, Not Content Indexed, and Encrypted.

To browse the properties and contents of a directory or file, perform the following steps:

1. From one of the starting points described in “Ways to Find a Snapshot” on page 326, open a snapshot.

2. In the Views pane, find a directory or file that you want to view properties and permissions for.

3 In the Content pane, select the object and then open it.

Figure 8-8: Sample Object Browser for a Unix Directory

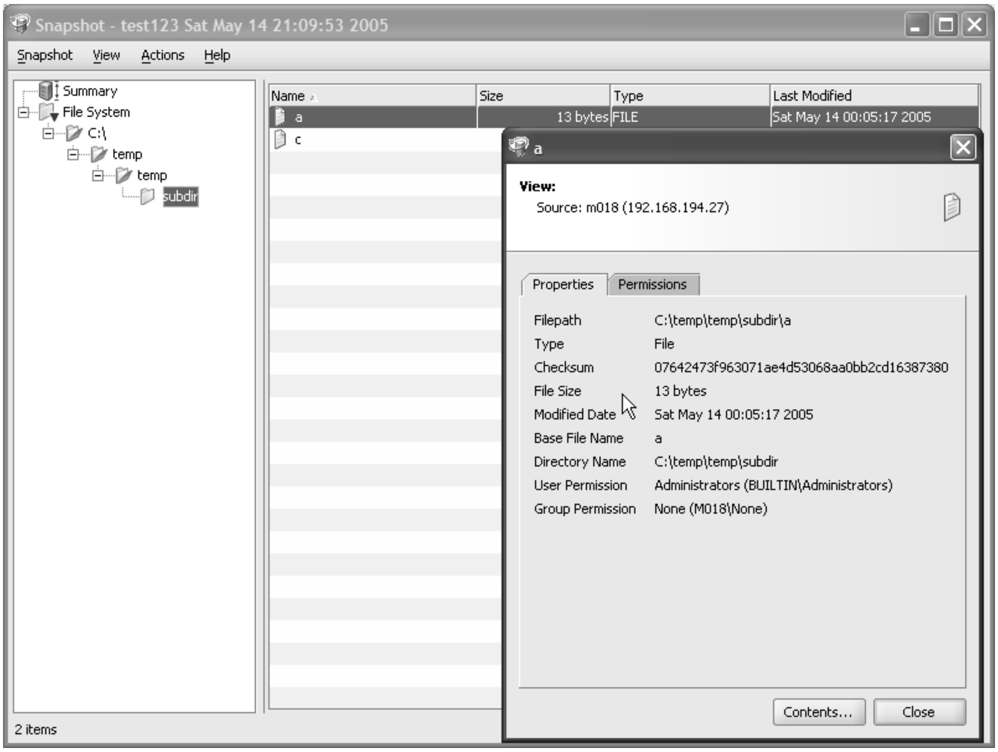

Name	Size	Type	Last Modified
a	13 bytes	FILE	Sat May 14 00:05:17 2005
c			

View:
Source: m018 (192.168.194.27)

Filepath	C:\temp\temp\subdir\a
Type	File
Checksum	07642473f963071ae4d53068aa0bb2cd16387380
File Size	13 bytes
Modified Date	Sat May 14 00:05:17 2005
Base File Name	a
Directory Name	C:\temp\temp\subdir
User Permission	Administrators (BUILTIN\Administrators)
Group Permission	None (M018\None)

4 Click **Close** to close the object browser.

Converting a Snapshot to a General Snapshot

A snapshot is typically associated with the server (source) that it was generated from. If you need to keep the snapshot and decommission its associated server, you must first convert the snapshot to a general snapshot–so that the snapshot is no longer associated with the server it was generated from. This allows the server to be used for other purposes while still allowing the snapshot to be used in an audit.

You must have a set of permissions to convert a snapshot to a general snapshot. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

To convert a snapshot to a general snapshot, perform the following steps:

1. From one of the starting points described in “Ways to Find a Snapshot” on page 326, select a snapshot.
2. From the OCC Client, select **Actions ➤ Save Snapshot**.
3. Click **OK** to save the snapshot in the Software Repository. After you save the snapshot, a general snapshot icon replaces the server snapshot icon.

When you decommission a managed server, all snapshots associated with that server will be deleted from the Software Repository.

Overview of Copying Objects from a Snapshot to a Server

After you have reviewed your snapshot contents, you may need to copy certain objects to a destination server. Server Compliance allows you to copy the following objects to a managed server: directories, files, Windows Services (state only), and Windows Registry keys.

Before you copy these objects over to a managed server, it is important to understand what actually gets copied to or created on the destination server:

- When you select a directory, only the directory will be copied to the destination server, excluding any files in that directory. For example, if dir1 contains file1 and file2, and you select dir1, Server Compliance copies only dir1 (not file1 and file2) to the destination server.
- When you select a file and its parent directory does not exist on the destination server, Server Compliance will create the directory on and copy the files to the destination server. For example, if you select file1 and dir1 does not exist on the destination server, Server Compliance will create dir1 on and copy file1 to the destination server.
- When you copy a Windows Services object, you copy the state of the service, such as started, stopped, paused, and so on. You can select one or more Windows Services objects for a single copy process.
- When you copy a Windows Registry object, you can select one or more registry keys and subkeys for a single copy process.

You must have write permission on the destination server to be able to copy an object to it. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

Copying Objects from a Snapshot to a Server

To copy an object from a snapshot to a managed server, perform the following tasks:

1. From one of the starting points described in "Ways to Find a Snapshot" on page 326, open a snapshot.
2. In the Views pane, select a File System, Windows Services, or Windows Registry.
3. In the Content pane, select one or more objects that you want to copy.
4. Select **Actions ➤ Copy To**.
5. In the Select Server window, select a destination server.

   🔍 Use the search tool to dynamically filter this list by entering a server name, IP address, or operating system.
6. Click **Select** to copy the object to that managed server or click **Cancel** to close this window without saving your changes.

For other types of server objects, such as packages and patches, you can also create installable packages to update a destination server. See "Visual Packager" on page 363 in Chapter 9 for more information.

Viewing a Snapshot Job

You can monitor the progress of a server snapshot process to determine whether it has completed or is still running. To view a snapshot job, perform the following steps:

1. In the Navigation pane, select Jobs and Sessions and then select Jobs Logs.
2. In the Content pane, select Create Snapshot from the drop-down list and select a status, such as In Progress or Completed.

   You can search for snapshots according to job ID, start time, and status of the snapshot process, including the number of managed servers and server groups in the snapshot.
3. Click **Update** to see the snapshots that meet your selection criteria.
4. Open a snapshot to view detailed information.

Deleting a Snapshot

As a best practice, you should delete snapshots that you no longer need from the Software Repository to conserve disk space.

You must have read permissions for the snapshot to be able to delete it. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

To delete snapshots, perform the following steps:

1. From one of the starting points described in “Ways to Find a Snapshot” on page 326, select a snapshot or select multiple snapshots and then select **Actions ➤ Delete**.
2. In the Confirmation Dialog, click **Yes** to delete this snapshot or click **No** if you do not want to delete it.

When you delete a snapshot, you do not delete the snapshot template that was used to create it. See “Deleting a Snapshot Template” on page 319 in this chapter for more information.

Auditing

An audit is a process that compares Opsware-managed servers to determine how objects and object content differ. Each audit requires an audit template that specifies the following requirements:

Selection criteria: These are rules that instruct Opsware SAS. You can set the criteria to select particular server objects to collect information about, select the files to compare, and select the objects to include and exclude. You can also set how the system will collect the server objects.

Source: This is the existing snapshot or server that you are comparing selection criteria from.

Target: This is the existing snapshot, server, or server group that you are comparing selection criteria to.

Depending on the server objects that you select for the audit, you can set up options for a more detailed comparison, such as file content and checksum, file inclusions and exclusions, file last modification date, user permissions, and user group access.

Figure 8-9: The Server Auditing Process

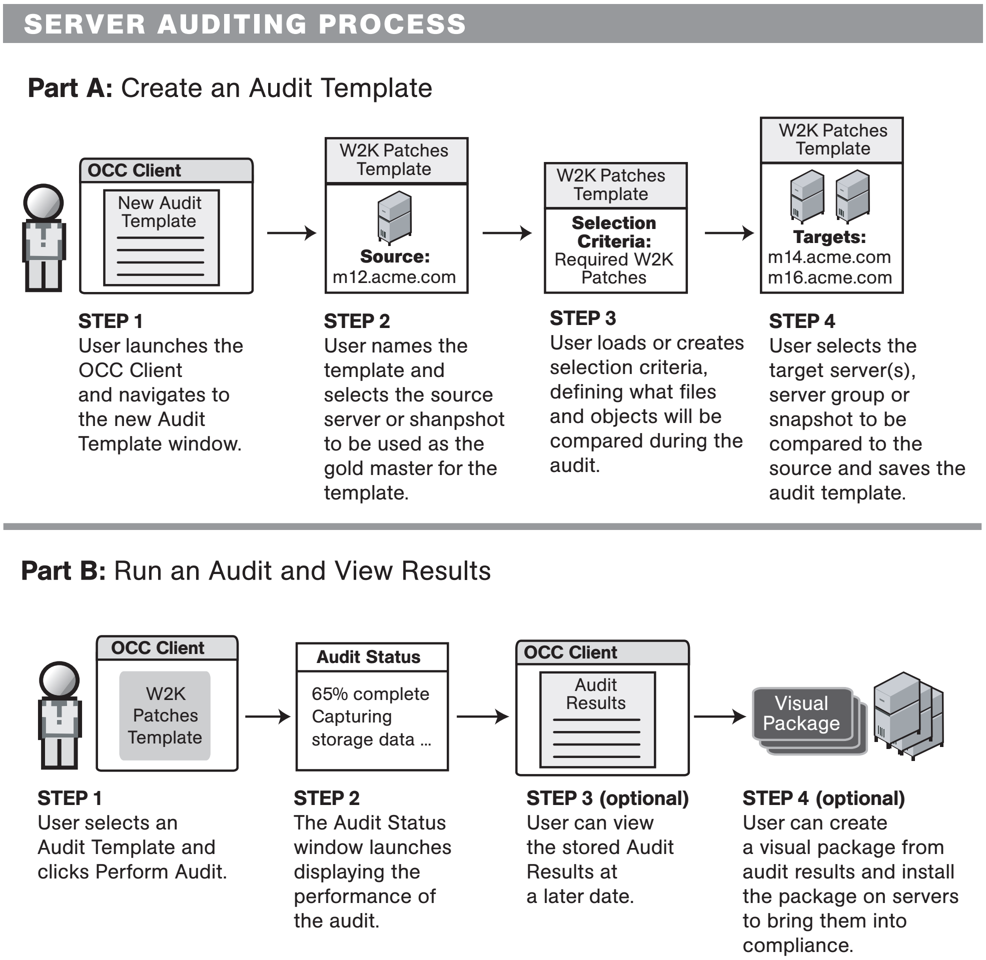

SERVER AUDITING PROCESS

Part A: Create an Audit Template

STEP 1
User launches the OCC Client and navigates to the new Audit Template window.

STEP 2
User names the template and selects the source server or shanpshot to be used as the gold master for the template.

STEP 3
User loads or creates selection criteria, defining what files and objects will be compared during the audit.

STEP 4
User selects the target server(s), server group or snapshot to be compared to the source and saves the audit template.

Part B: Run an Audit and View Results

STEP 1
User selects an Audit Template and clicks Perform Audit.

STEP 2
The Audit Status window launches displaying the performance of the audit.

STEP 3 (optional)
User can view the stored Audit Results at a later date.

STEP 4 (optional)
User can create a visual package from audit results and install the package on servers to bring them into compliance.

Ways to Perform an Audit

In Server Compliance, you can perform an audit in several different ways, based on the Opsware-managed server, a server snapshot, an audit template, or an existing audit result. The section reviews the following actions:

- Performing an Audit from a Managed Server
- Performing an Audit form a Server Group

You can also perform an ad hoc audit by editing an existing template. An ad hoc audit differs from a regular audit, in that you launch the audit process without saving your changes. The results from this audit allow you to see differences between servers that may affect their ability to operate properly. See "Editing an Audit Template" on page 346 in this chapter for more information.

- Performing an Audit from an Audit Template
- Re-running an Audit from Audit Results

Performing an Audit from a Managed Server

To perform an audit from a managed server, the managed server must be reachable.

1. From the Navigation pane, select Servers and then select All Managed Servers.
2. Select a managed server and then select **Actions ➤ Perform Audit**.

Performing an Audit form a Server Group

1. From the Navigation pane, select Servers and then select Server Groups.
2. Select a server group, and the select **Actions ➤ Perform Audit**. (When you perform an audit by selecting a server group, you must select an individual server in that group as the source.)

A server group cannot be used as a source in the audit process. However, a server group can be used as a target in the audit process.

Performing an Audit from a Snapshot

1. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Snapshot Tab.
3. Select a snapshot and then select **Actions ➤ Perform Audit**.

Performing an Audit from an Audit Template

1. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Audit Templates Tab.
3. Select an audit template and then select **Actions ➤ Perform Audit**.

Re-running an Audit from Audit Results

1. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Audit Results Tab.
3. In the Content pane, select audit results and then select **Actions ➤ Re-Run Audit**.

Audit Templates

An audit template identifies the source and targets that will be compared, including the selection criteria that will be used in the audit process. You can modify audit templates, and copy and re-use them to compare other sources and targets in Opsware SAS. When you no longer need an audit template, you can remove it from the Model Repository. Figure 8-10 illustrates a sample audit template.

You must have a set of permissions to create and modify audit templates. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

Figure 8-10: Sample Audit Template

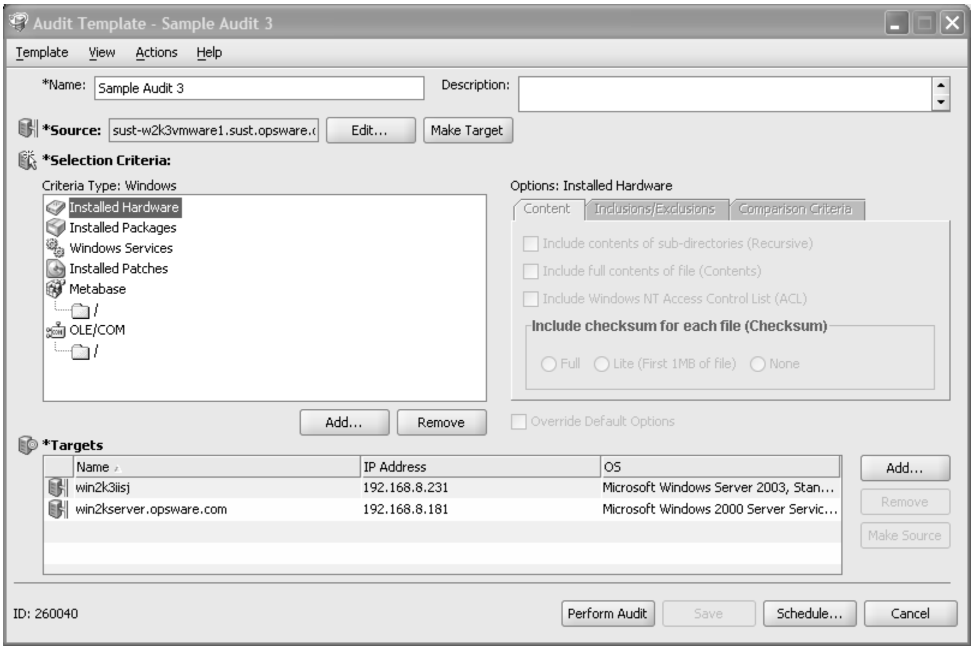

*Name: Sample Audit 3

Description:

*Source: sust-w2k3vmware1.sust.opsware.c

*Selection Criteria:

Criteria Type: Windows

- Installed Hardware
- Installed Packages
- Windows Services
- Installed Patches
- Metabase
  - /
- OLE/COM
  - /

Options: Installed Hardware

□ Include contents of sub-directories (Recursive)

□ Include full contents of file (Contents)

□ Include Windows NT Access Control List (ACL)

Include checksum for each file (Checksum)

○ Full ○ Lite (First 1MB of file) ○ None

□ Override Default Options

*Targets

Name	IP Address	OS
win2k3iisj	192.168.8.231	Microsoft Windows Server 2003, Stan...
win2kserver.opsware.com	192.168.8.181	Microsoft Windows 2000 Server Servic...

ID: 260040

Since selection criteria is specific to an operating system, such as Unix or Windows, you cannot have one (universal) audit template for different operating systems of managed servers in Opsware SAS.

Creating an Audit Template

You must create an audit template to identify your source and targets, and to define the selection criteria used in the comparison audit.

An asterisk (*) at the beginning of a field name indicates that this information is required.

To create an audit template, perform the following steps:

1. From one of the starting points described in “Ways to Perform an Audit” on page 336, select **Actions ➤ Perform Audit**. The Audit Template Window appears.

2. In the Name field, enter a name for the audit you are creating. This name must be unique in Opsware SAS.

3. (Optional) In the Description field, enter a brief description of the type of audit you are creating.

4. Specify a managed server or a snapshot that you want to be used as the source in the audit process. The source is what the audit compares selection criteria from. Specify your selection criteria. See “Setting Up Audit Selection Criteria” on page 340 in this chapter for more information.

5. Specify the servers, server groups, and snapshots that you want to be used as Targets in the audit process.

6. Click **Save** to save your audit template in the Model Repository or click **Cancel** to close this window without saving your changes.

7. When ready to launch the audit, click **Perform Audit**.

Specifying Source and Targets in the Audit Template

Depending on whether you decide to perform an audit by selecting single or multiple managed servers, or by selecting a server group, there are several ways to specify the source and targets that you want to be used in the audit process.

Performing an Audit by Selecting a Single Server

When you perform an audit by selecting a single server, by default that server is defined as the source for the audit process. You can modify the default source as follows:

1. From one of the starting points described in “Ways to Perform an Audit” on page 336, select **Actions ➤ Perform Audit**.

2. In the Audit Template Window, click **Edit** to select a different managed server (using the Source Chooser) that will be used as the source in the audit.

3. Click **Make Target** to change the default source server to a target that will be used in the audit.

Performing an Audit by Selecting Multiple Servers

When you perform an audit by selecting multiple servers, by default they are all defined as targets for the audit process and are listed in the Targets section in the audit template. You can modify this list of default targets as follows:

- In the Audit Template Window, select a server and then click **Make Source** to specify that the server will be used as the source (instead of a target) in the audit. You may be required to do this to correct a setup error and if objects on your reference server have changed.
- Select one or more servers and then click **Remove** so that the servers will not be used as targets in the audit.
- Click **Add** to select one or more servers, server groups, and snapshots (using the Target Chooser) that will be used as targets in the audit.

If you do not make any changes to the default targets, all of the managed servers in the list will be used as targets in the audit process.

Setting Up Audit Selection Criteria

You can set up audit selection criteria in different ways. You can re-use existing (previously-saved) audit or snapshot selection criteria, or you can define new audit selection criteria.

To use existing audit or snapshot selection criteria, you load it into your audit template as is, or you load it into your audit template and then modify it as required.

To define new selection criteria, you choose server objects from the Selection Criteria Editor.

To set up selection criteria for the server audit process, perform the following steps:

1. From one of the starting points described in "Ways to Perform an Audit" on page 336, select **Actions ➤ Perform Audit**. The Audit Template Window appears.

2 In the Audit Template, specify the selection criteria that will be used in the audit process using any of the following methods:

- If you want to define new selection criteria in your audit template, then at the bottom of the Criteria Type pane, click **Add** to open the Selection Criteria Editor window where you can select one or more server objects you want to include in the audit. You can also select one or more objects you want to delete from your selection criteria and then click **Remove**.
- To include all COM objects in an audit of a Windows server, you must select the OLE/COM top-level node in the Selection Criteria Editor. Opsware SAS categorizes the COM objects based on an attribute of the object, where the COM object specifies zero or more categories. Opsware SAS displays all COM objects in one node (OLE/COM) in the server object tree in case you do not know the category or the object did not specify one.
- See “Setting Up Audit Selection Criteria” on page 340 in this chapter for more information.

3 In the Audit Template window, select the Content tab to specify whether the audit will include:

- Contents of subdirectories (Recursive)
- Full contents of a file (Contents)
- Windows NT Access Control List (ACL)
- Checksum for each file (Checksum)

  You can select more detailed selection criteria for checksum, depending on whether you want the audit to include all of the content of the files (Full) or only the first megabyte (MB) of each file (Lite). Checksum is performed on the contents of a file. File owner and file permission changes are not audited. The default checksum is Full.

  You can also specify audit content that applies to the entire file system or Windows Registry of the managed server, or you can specify audit content that applies only to selected subdirectories or files in the file system or Windows Registry. If you do not specify options for subdirectories, the default audit content is the parent directory.

4 Select the Inclusions/Exclusions tab to specify the directories and files in the file system you want included in and excluded from the audit.

- Select Include and enter the path or file name in the text box. Click **Add** to display this information in the list of content that will be included in the audit.
- Select Exclude and enter the path or file name in the text box. Click **Add** to display this information in the list of content that will be excluded from the audit.

See "Inclusion and Exclusion Rules" on page 302 in this chapter for more information.

5 Select the Comparison Criteria tab and select whether you want the audit to compare the file modification date (Date), and the user and group access permissions (Access).

6 (Optional) Select an individual directory or file in the file system and then check Override Default Options to override the selection criteria that apply to the entire file system. This option allows you to specify different selection criteria for different directories and files. The selection criteria that is specified at the file system level (default options) will not be applied to these individual objects during the audit process. The selection criteria you defined for these individual objects will be applied during the audit process.

For the Windows Registry object, only the Content tab is enabled and only the Recursive and full Contents options are available. Options are not available for Windows Services objects.

7 Click **Save** to save this audit template in the Model Repository or click **Cancel** to close this window without saving your changes. See "Performing an Audit" on page 342 in this chapter for more information on launching the audit process.

Performing an Audit

In Server Compliance, you can perform an audit in several different ways. You can:

- Create and save an audit template, and then run the audit process. See "Creating an Audit Template" on page 338 in this chapter for more information.
- Select an existing audit template and then run the audit process. See "Re-running an Audit from Audit Results" on page 337 in this chapter for more information.
- Edit an existing audit template and then run the audit process. See "Editing an Audit Template" on page 346 in this chapter for more information.

- Select multiple managed servers and perform an ad hoc audit.

To perform an audit, you must have read permissions for all sources and targets (servers, server groups, and snapshots) that are specified in the audit template. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

Depending on the source and targets, the audit process can occur quickly or it can require additional time to perform background tasks, as follows:

- If the source and all of the targets defined in the audit template are all snapshots, the auditing process requires little time.

When you perform an audit and an existing snapshot is defined as the source, the selection criteria that was used when the snapshot was created is the selection criteria that will be used in the audit process.

- If the source or any of the targets defined in the audit template are *not* snapshots, Server Compliance will create temporary snapshots of them to be able to perform comparison tasks. The snapshot process runs in the background so that you can continue to use other OCC Client features.
- If a target is a live server, Server Compliance creates a temporary snapshot of it, which requires additional time. When this occurs, the audit process will run in the background (on the Command Engine) so that you can continue to use other OCC Client features.

If a target defined in the audit template is a heterogeneous server group (a server group that contains different operating systems), Server Compliance will filter the server group by the operating system that is specified in the Criteria Type. For example, if there are Windows and Unix servers in a server group, where the Criteria Type in the audit template specifies Unix as the operating system, Server Compliance will filter out all Windows servers during the audit process. Only the Unix servers in the target (server group) will be audited.

Tracking an Audit’s Progress

To track an audit’s progress, perform the following steps:

1. To check the status of an audit, select Jobs and Sessions and then select Job Logs in the Navigation pane.
2. Select Audit Servers from the drop-down menu in the Content pane.
3. Click **Update**. Audit Jobs will be displayed.
4. Select an audit and then open it to view details.

Saving Audit Selection Criteria for Reuse

You can save audit selection criteria to reuse it, or you can modify it, save it, and then reuse it to compare server objects on different targets (managed servers, server groups, and snapshots). As a best practice, this is something you should do if this is the type of audit that you need to perform frequently. You can save selection criteria from a new template you are creating or from an existing template to reuse it in other audit templates.

To save selection criteria in an audit template, perform the following steps:

1. From one of the starting points described in "Ways to Perform an Audit" on page 336, select **Actions ➤ Perform Audit**. The Audit Template Window appears.
2. In the Audit Template window, specify the selection criteria that will be used in the audit process. See "Setting Up Audit Selection Criteria" on page 340 in this chapter for more information.
3. When finished, select **Actions ➤ Save Selection Criteria As**.
4. In the Save Selection Criteria As window, enter a unique name.
5. (Optional) Enter a description of the selection criteria.
6. Click **Save** to save your selection criteria or click **Cancel** to close this window without saving your changes.

Loading Audit Selection Criteria

You can re-use existing (previously-saved) audit or snapshot selection criteria by loading them into your audit template. You can load the selection criteria into your audit template and leave it as is, or you can load them into your audit template and then modify as required.

If the Source in the audit template is a managed server, the Load Selection Criteria window will display a list of existing audit selection criteria.

If the Source in the audit template is a snapshot, the Load Selection Criteria window will display a list of existing snapshot selection criteria. Server Compliance allows you to use snapshot selection criteria in an audit process.

To load audit or snapshot selection criteria into your audit template, perform the following steps:

1. From one of the starting points described in “Ways to Perform an Audit” on page 336, select **Actions ➤ Perform Audit**. The Audit Template Window appears.
2. In the Audit Template window, specify the selection criteria that will be used in the audit process.
3. When finished, select **Actions ➤ Load Selection Criteria** to choose from a list of existing selection criteria in the Load Selection Criteria window. This list will contain only selection criteria that support the type of operating system (Unix or Windows) of the Source (managed server or snapshot) that is specified in the audit template. Server Compliance filters out any selection criteria that does not apply to the operating system of the source.
4. In the Load Selection Criteria window, select a criteria name and then click **Load** to add the selection criteria to your audit template, or click **Close** to close this window without saving your changes.
5. If you want to modify the loaded selection criteria, at the bottom of the Criteria Type pane in the audit template, click **Add** to open the Selection Criteria Editor window where you can select one or more server objects that you want to include in the audit. You can also select one or more objects that you want to delete from your selection criteria and then click **Remove**.

Managing Your Audit Selection Criteria

After you have previously saved your selection criteria, you can view, rename, save as, and delete them as required.

To manage your audit selection, perform the following steps:

1. From the Navigation pane, select Software Library and then select Compliance.
2. In the Compliance window, select **Actions ➤ Manage Selection Criteria**.
3. In the Manage Selection Criteria As window, select a selection criteria name and click any of the following buttons:
   - **Open**: To view the selection criteria.

- **Rename**: To change the name of the selection criteria.
- **Save As**: To make a copy of the selection criteria.
- **Delete**: To delete selection criteria. You can select multiple selection criteria names to perform a mass delete.
- **Close**: To close this window without saving your changes.

Editing an Audit Template

To perform an ad hoc audit, edit a template and run it without saving it.

✓ You must have write permissions to edit an audit template. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

To modify an audit template, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. Select the Audit Templates tab, and then open an audit template.
3. Make changes to the template as necessary.
4. After you make your changes, choose any of the following actions:

   - Click **Perform Audit** to launch the audit process without saving your changes. This is commonly known as performing an ad hoc audit.
   - Click **Save** to save the revised audit template in the Model Repository. This action will overwrite the previous version of the template.
   - Click **Schedule** to schedule the audit.
   - Click **Cancel** if you want to close this window without saving your changes.

Copying an Audit Template for Reuse

You can copy an audit template so that you can reuse the selection criteria to compare other sources and targets.

To copy an audit template, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. Select the Audit Templates tab, and then open an audit template.
3. In the Audit Template window, select **Actions ➤ Save Template As**.
4. In the Save Template As window, enter a unique name.
5. (Optional) Enter a description of the template.
6. Click **Save** to save the audit template in the Model Repository or click **Cancel** to close this window without saving your changes.

Deleting an Audit Template

As a best housekeeping practice, you should delete audit templates that you no longer need from the Model Repository.

✓ You must have write permissions to delete an audit template. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

To delete an audit template, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. Select the Audit Templates tab, and then select one or more audit templates and select **Actions ➤ Delete**.
3. In the Confirmation Dialog, click **Yes** to delete this audit template or click **No** if you do not want to delete it.

Audit Schedules

An audit schedule specifies when you want an audit to be performed (either once or as a recurring job) and who you want to receive email notification about the status of the job. You can also view, edit, and delete existing audit schedules. When you delete an audit template, all schedules that you own (created), that are also associated with that audit template, will be deleted.

You must have permissions to create, view, edit, and delete audit schedules. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

Scheduling an Audit

After you have completed the Audit Template and saved it, you can define a schedule that specifies when you want the audit performed.

To schedule an audit job, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. Select the Audit Templates tab, and then open an audit template.
3. In the Audit Template window, select **Actions ➤ New Schedule** to display the Schedule Job window.

Figure 8-11: Audit Schedule

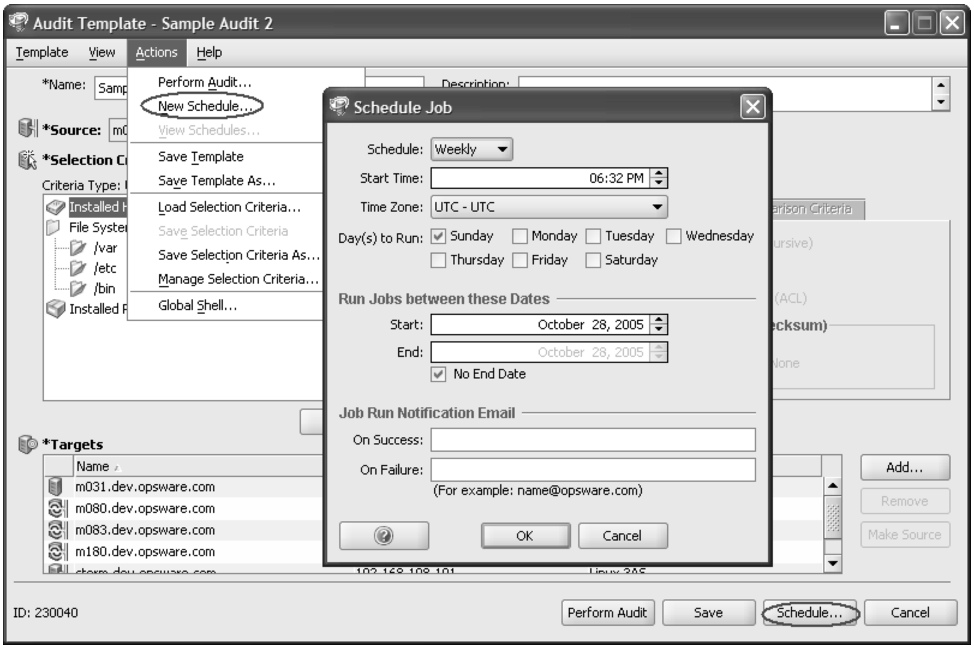

Audit Template - Sample Audit 2

Schedule Job

Schedule: Weekly

Start Time: 06:32 PM

Time Zone: UTC - UTC

Day(s) to Run: ☑ Sunday ☐ Monday ☐ Tuesday ☐ Wednesday ☐ Thursday ☐ Friday ☐ Saturday

Run Jobs between these Dates

Start: October 28, 2005

End: October 28, 2005

☑ No End Date

Job Run Notification Email

On Success:

On Failure:

(For example: name@opsware.com)

ID: 230040

4. In the Schedule field, choose whether you want the audit to be performed once, daily, weekly, monthly, or on a custom schedule.

When you initially specify an audit schedule, the Schedule field default is Weekly. If you check every day of the week when the Schedule field is set to Weekly, the Schedule field changes to Daily when you subsequently edit the audit schedule.

5 Depending on the type of Schedule that you selected, use the following guidelines to enter a Start Time that specifies when you want the job to begin:

- To perform an audit only once, enter the complete date and time, such as October 5, 2005 12:01AM.
- To perform an audit daily, weekly, or monthly, enter the time of day, such as 12:01AM.
- To perform an audit on a custom schedule, use the following syntax to enter a crontab string that specifies the day, date, and time you want the job to begin, as shown in Table 8-12:

Figure 8-12: Crontab

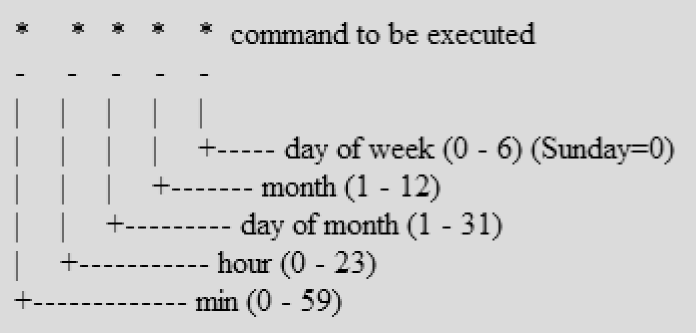

```
*   *  *  *  *  command to be executed
-   -  -  -  -
|   |  |  |  |
|   |  |  |  +----- day of week (0 - 6) (Sunday=0)
|   |  |  +------- month (1 - 12)
|   |  +--------- day of month (1 - 31)
|   +----------- hour (0 - 23)
+------------- min (0 - 59)
```

An asterisk (*) in any of these fields represents all days of the month, all months of the year, all days of the week, and so on. For example, the following crontab string will perform the audit at midnight every weekday:

0 0 * * 1-5

The crontab string can include serial (1,2,3,4) and range (1-5) values. For more information about crontab entry formats, consult the Unix man pages.

6 Select a default time zone for the job execution time, or accept the default time zone. The default time shown converts the scheduled time to the time zone set in your user preferences. If you do not set a preferred time zone, the time zone is derived from the Opsware SAS core server (typically UTC).

7 Depending on the Schedule that you selected, complete one of the following steps:

- If you selected Weekly, select the days of the week you want the audit to be performed in the Days to Run field.
- If you selected Monthly, select the Day of the month to perform the audit. You must also select the months of the year to perform the audit in the Months to Run field.

8 In the Run Jobs Between These Dates section, select a Start and End time to specify when you want the audit process to begin and end.

9 (Optional) Select No End Date if you do not want the audit schedule to expire.

10 (Optional) In the Job Run Notification Email section, enter email addresses (separated by a comma or a space) of those that you want to receive information about the status of the audit job, including IDs of servers affected. Enter email addresses in one or both of the following fields:

On Success: Email addresses that will receive notifications of jobs that completed, such as name@opsware.com, name2@opsware.com.

On Failure: Email addresses that will receive notifications of jobs that completed, such as name@opsware.com, name2@opsware.com.

11 Click **OK** to save your schedule information or **Cancel** to close this window without saving any information.

Viewing an Audit Schedule

You can view an audit schedule after you have created (or edited) and saved it. To view an audit schedule, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.

2. Select the Audit Templates tab, and then open an audit template.

3. In the Audit Template window, select **Actions ➤ View Schedules** or click **Schedule** to display the Scheduled Jobs window.

Figure 8-13: Scheduled Audits

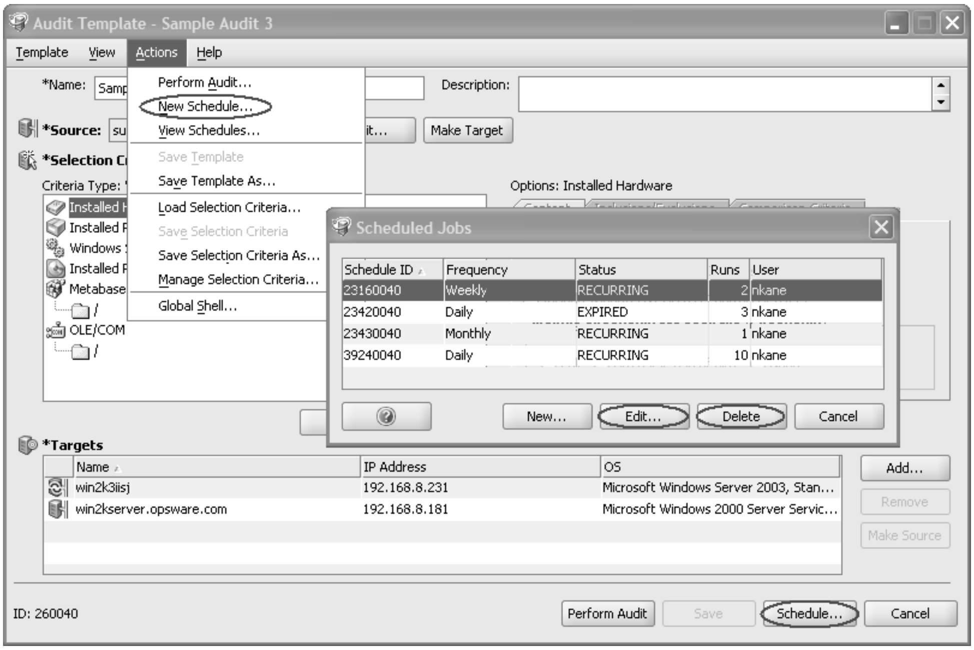

*Name: Samp

Options: Installed Hardware

Schedule ID	Frequency	Status	Runs	User
23160040	Weekly	RECURRING	2	nkane
23420040	Daily	EXPIRED	3	nkane
23430040	Monthly	RECURRING	1	nkane
39240040	Daily	RECURRING	10	nkane

Name	IP Address	OS
win2k3iisj	192.168.8.231	Microsoft Windows Server 2003, Stan...
win2kserver.opsware.com	192.168.8.181	Microsoft Windows 2000 Server Servic...

ID: 260040

4. Open the audit job you want to review.

5. In the Scheduled Jobs window, review the detailed audit schedule information.

6. Click **Cancel** to close this window.

Editing an Audit Schedule

You can edit an audit schedule after you have created and saved it. To edit an audit schedule, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. Select the Audit Templates tab, and then open an audit template.
3. In the Audit Template window, select **Actions ➤ View Schedules** to display the Scheduled Jobs window. See Figure 8-13.
4. Select a job and then click **Edit** to display the Schedule Job window.
5. Make your schedule changes.
6. Click **OK** to save your changes or **Cancel** to close this window without saving your changes.

Deleting an Audit Schedule

You can delete a saved audit schedule. To delete an audit schedule, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Compliance.
2. Select the Audit Templates tab, and then open an audit template.
3. In the Audit Template window, select **Actions ➤ View Schedules** to display the Scheduled Jobs window. See Figure 8-13.
4. Select the audit schedule you want to delete and then click **Delete**.
5. In the Confirmation Dialog, click **Yes** to delete the schedule or click **No** if you do not want to delete it.
6. (Optional) Click **Cancel** to close this window without deleting the selected audit schedule.

Audit Results

Audit results contain only the differences that were discovered during the server audit process. You can view differences between servers and snapshots, and differences between two files (side by side and line by line).

You can use these audit results to determine what type of remediation tasks are required for the server that is not compliant. To conserve disk space, you should delete audit results that you no longer need from the Software Repository.

You must have read write permissions to delete audit results. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

Ways to View Audit Results

You can review audit results in the Audit Result windows in the OCC Client. The Audit Result window provides visual cues that indicate whether server object differences were discovered in the audit process. In the server object tree in the Audit Result window, objects that differ are displayed in blue text and are followed by an asterisk (*). See Figure 8-14 on page 354 for an illustration.

In the Jobs and Sessions window:

1. From the Navigation pane, select Jobs and Sessions and then select Job Logs.
2. In the Job Types drop-down list, select Audit Servers.
3. (Optional) Specify a job ID, date and time range, and job status options to narrow your search for audit results.
4. Click **Update** to view a list of audit results based on your search criteria.
5. Select an audit in the list and then open it.

In the Compliance window:

1. From the Navigation pane, select Software Library and then select Compliance.
2. Select the Audit Results tab.
3. Select an audit and then open it.

In the All Managed Servers window:

1. From the Navigation pane, select Servers and the All Managed Servers.
2. From the **View** Menu, select **Software Library ➤ Compliance**.
3. In the Compliance window, select the Audit Results tab.
4. Select an audit and then open it.

Viewing Audit Results on a Server by Server Basis

To view audit results on a server-by-server basis, perform the following steps:

1. From one of the starting points described in “Ways to View Audit Results” on page 353, open an audit result. The Audit Results Window appears.

Figure 8-14: Sample Audit Result Window

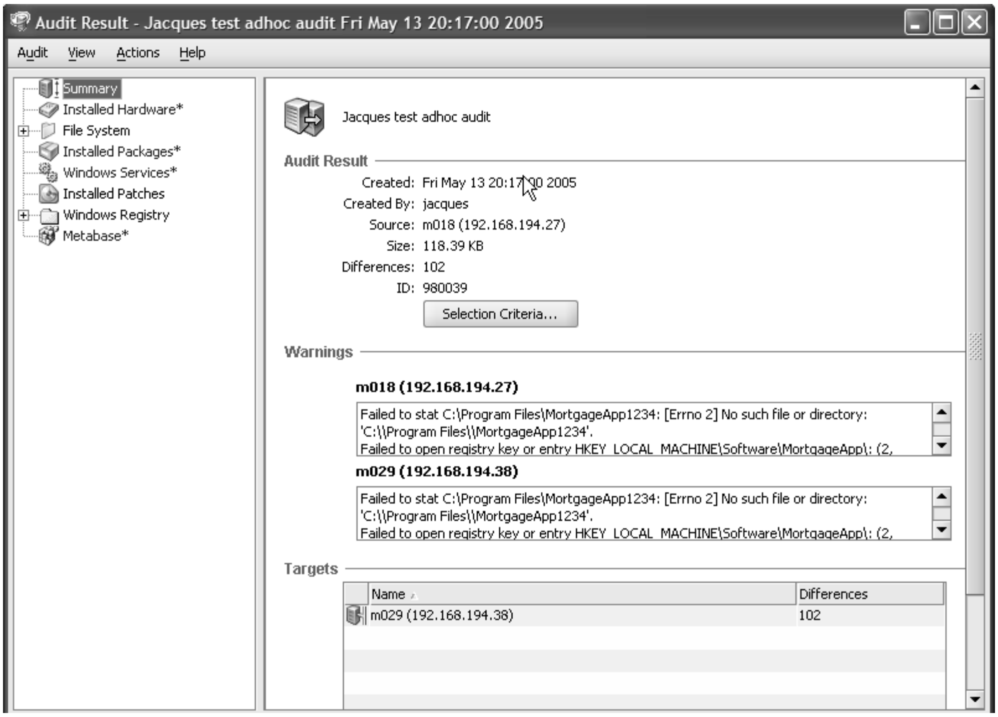

Jacques test adhoc audit

Audit Result

Created: Fri May 13 20:17:00 2005
Created By: jacques
Source: m018 (192.168.194.27)
Size: 118.39 KB
Differences: 102
ID: 980039

Warnings

m018 (192.168.194.27)

```
Failed to stat C:\Program Files\MortgageApp1234: [Errno 2] No such file or directory:
'C:\\Program Files\\MortgageApp1234'.
Failed to open registry key or entry HKEY_LOCAL_MACHINE\Software\MortgageApp\: (2,
```

m029 (192.168.194.38)

```
Failed to stat C:\Program Files\MortgageApp1234: [Errno 2] No such file or directory:
'C:\\Program Files\\MortgageApp1234'.
Failed to open registry key or entry HKEY_LOCAL_MACHINE\Software\MortgageApp\: (2,
```

Targets

Name	Differences
m029 (192.168.194.38)	102

2. In the Audit Result window, you can select:

   - **Summary** to view a summary of your audit results, such as the date and time the audit was created, the audit source (an existing snapshot or a server) that was compared to the targets, the size of the audit file, the total number of differences found for each target in the audit, an audit ID number, warning messages, and a list of targets in the audit. The warning messages are useful for searching the database to troubleshoot problems.

   - **Installed Hardware** to view the hardware audit results for your source and targets, such as CPU model, amount of RAM, storage differences, different cache size, processor speed, owner, permissions, and interfaces. The number in parentheses next to the tab name indicates how many differences were discovered for that type of source and target comparison. Browse the tabs to view hardware differences that were found only on the source, only on the targets, or on both the source and the targets. See “Viewing Contents of Objects that Differ” on page 356 in this chapter for more information.

- **Installed Packages** to view the installed packages differences for your source and targets, such as the package version and package type. Server Compliance supports the following types of packages: Slowpoke for Unix, MSI for Windows, LPP for AIX, Depot for HPUX, and RPM. The number in parentheses next to the tab name indicates how many differences were discovered for that type of source and target comparison. Browse the tabs to view installed packages differences that were found only on the source, only on the targets, or on both the source and the targets.
- **Patches** to view the installed patches differences for your source and targets, such as the patch type. The number in parentheses next to the tab name indicates how many differences were discovered for that type of source and target comparison. Browse the tabs to view installed patch differences that were found only on the source, only on the targets, or on both the source and the targets.
- **File System** to view the file system audit results for your source and target servers, such as checksum, permissions, and file size. You can also view differences between two files (side by side and line by line). The number in parentheses next to the tab name indicates how many differences were discovered for that type of source and target comparison. Browse the tabs to view installed file system differences that were found only on the source, only on the targets, or on both the source and the targets. In the View Contents window, select an item from the Encoding drop-down list to specify the character encoding of the data displayed.
- **Windows Services** to view the running services audit results for your source and targets, such as the name, status, startup type, and log on (as a particular user, such as LocalSystem, NT Authority, and so on) of the service. The number in parentheses next to the tab name indicates how many differences were discovered for that type of source and target comparison. Browse the tabs to view differences that were found only on the source, only on the targets, or on both the source and the targets.
- **Windows Registry** to view the registry key audit results for your source and targets, such as HKEY_CLASSES_ROOT, HKEY_CURRENT_CONFIG, HKEY_LOCAL_MACHINE, and HKEY_USERS. The number in parentheses next to the tab name indicates how many differences were discovered for that type of source and

target comparison. Browse the tabs to view Windows registry key differences that were found only on the source, only on the targets, or on both the source and the targets.

Valid control characters for the contents of the key entry (Data) include: #x9, #xA, [#xD, #x20-#xD7FF], [#xE000-#xFFFD], and [#x10000-#x10FFFF]. Invalid control characters cannot be stored by Opsware SAS and will be converted to XML entities and will display as &#;. For example, if the data value is 00 00 (in bytes), &#x00; will display in this column.

- **OLE/COM** to view Windows COM (Common Object Model) audit results for your source and targets, such as the name and GUID (Globally Unique User ID) of the object, and the path to the in-process server DLL. The number in parentheses next to the tab name indicates how many differences were discovered for that type of source and target comparison. Browse the tabs to view OLE/COM object differences that were found only on the source, only on the targets, or on both the source and the targets. See "Viewing OLE/COM Objects that Differ" on page 357 in this chapter for more information.
- **Metabase** to view IIS Metabase object audit results for your source and targets, such as the ID, name, path, attributes, and data of the object. The number in parentheses next to the tab name indicates how many differences were discovered for that type of source and target comparison. Browse the tabs to view Metabase object differences that were found only on the source, only on the targets, or on both the source and the targets. See "Viewing Metabase Objects that Differ" on page 359 in this chapter for more information.

Viewing Contents of Objects that Differ

For some objects that were audited, you can view content differences side by side and line by line. You can see which lines in a file were added, deleted, or modified.

To view the contents of two objects that differ, perform the following steps:

1. From one of the starting points described in "Ways to View Audit Results" on page 353, open an audit result. The Audit Results Window appears.
2. In the Views pane, select an object in the File System.
3. In the Content pane, select the On Both but Different tab and select two objects.
4. Select **Actions ➤ View Differences**.

5. In the Comparison window, select an item from the Encoding drop-down list to specify the character encoding of the data displayed.

6. Click the arrows to find the first, next, previous, or last lines that were added, deleted, or modified in the file that is in the source or in the file that is in a target. Differences are highlighted according to the following color scheme:

   - **Green**: This content was added.
   - **Blue**: This content was modified.
   - **Red**: This content was deleted.
   - **Black**: No changes were made to this content.

7. Click **Close** to close this window.

If one of the files you are comparing exceeds 2MB in file size, Server Compliance cannot display the file differences.

Viewing OLE/COM Objects that Differ

For OLE/COM objects that were audited, you can view differences in object properties and user group permissions, side by side.

To view the contents of two objects that differ, perform the following steps:

1. From one of the starting points described in “Ways to View Audit Results” on page 353, open an audit result. The Audit Results Window appears.

2. In the Views pane, select the OLE/COM object.

3. In the Content pane, select the On Both but Different tab.

4 In the Content pane, select objects and then select **Actions ➤ View Differences**.

Figure 8-15: Sample OLE/COM Object Browser

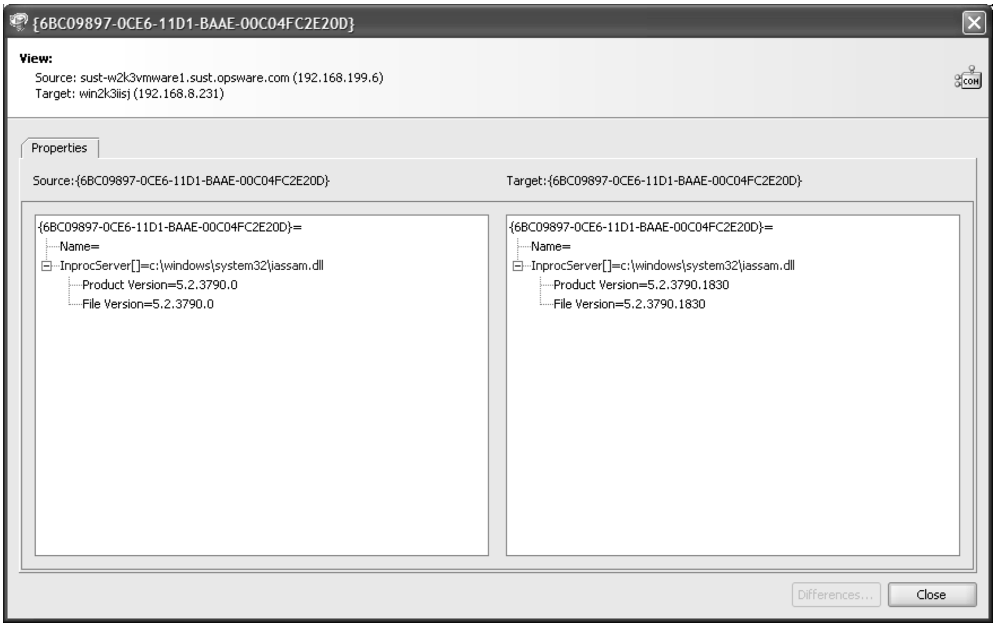

{6BC09897-0CE6-11D1-BAAE-00C04FC2E20D}

View:
Source: sust-w2k3vmware1.sust.opsware.com (192.168.199.6)
Target: win2k3iisj (192.168.8.231)

Source:{6BC09897-0CE6-11D1-BAAE-00C04FC2E20D}

```
{6BC09897-0CE6-11D1-BAAE-00C04FC2E20D}=
  Name=
  InprocServer[]=c:\windows\system32\iassam.dll
    Product Version=5.2.3790.0
    File Version=5.2.3790.0
```

Target:{6BC09897-0CE6-11D1-BAAE-00C04FC2E20D}

```
{6BC09897-0CE6-11D1-BAAE-00C04FC2E20D}=
  Name=
  InprocServer[]=c:\windows\system32\iassam.dll
    Product Version=5.2.3790.1830
    File Version=5.2.3790.1830
```

5 In the View window, select one of the tabs (Properties, Access Permissions, or Launch Permissions) to view those types of differences for the selected object. Object properties and permissions that differ are displayed in blue text.

When you perform an audit of two servers using selection criteria that includes installed hardware where the RAM size is different on the source and target, this information is displayed in the file object browser as the same RAM size, even though the RAM size is really different. The OCC Client collects RAM size in bytes and converts it to megabytes (MB). If the RAM size difference is very slight, it will not be recorded as a difference in megabytes. However, the difference will be displayed in kilobytes (KB) in parenthesis next to the MB display.

In the object browser, only the parent (top-level) nodes are displayed in blue text. Child node differences are *not* displayed in this tree structure. To identify differences in child nodes, you must individually inspect them.

6 Click **Close** to close the object browser window.

Viewing Metabase Objects that Differ

For Metabase objects that were audited, you can view differences in connection information, bandwidth, server state and size, IISAdmin Extensions, cache extensions, Allow Keep Alives, and so on.

To view the contents of two objects that differ, perform the following steps:

1. From one of the starting points described in “Ways to View Audit Results” on page 353, open an audit result. The Audit Results Window appears.
2. In the Views pane, select the Metabase object.
3. In the Content pane, select the On Both but Different tab.
4. Select an object and then select **Actions ➤ View Differences**.

Figure 8-16: Sample IIS Metabase Object Browser

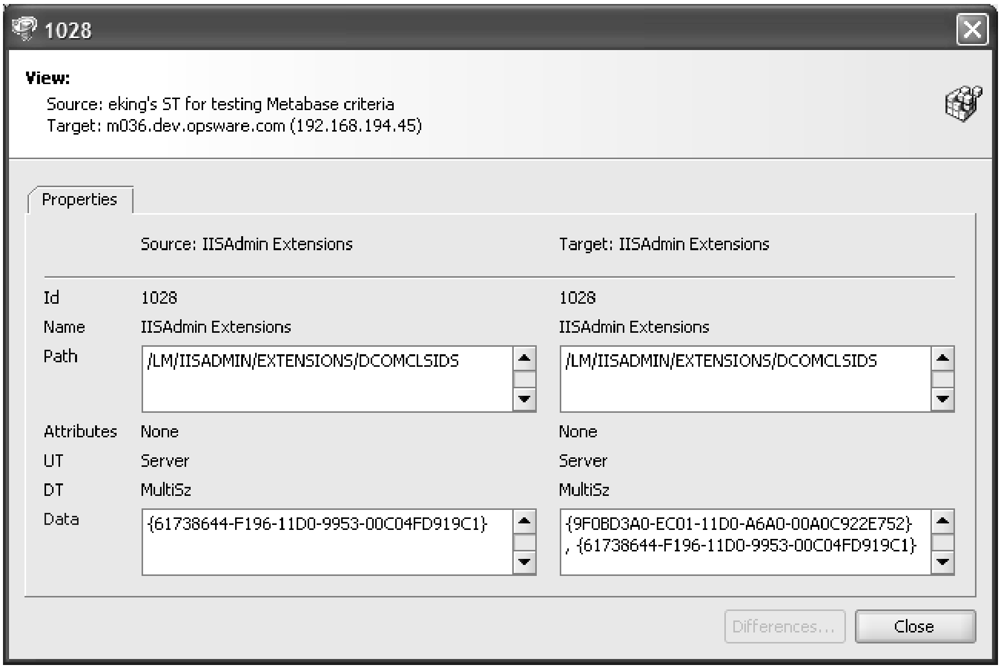

1028

View:
Source: eking's ST for testing Metabase criteria
Target: m036.dev.opsware.com (192.168.194.45)

Properties

	Source: IISAdmin Extensions	Target: IISAdmin Extensions
Id	1028	1028
Name	IISAdmin Extensions	IISAdmin Extensions
Path	/LM/IISADMIN/EXTENSIONS/DCOMCLSIDS	/LM/IISADMIN/EXTENSIONS/DCOMCLSIDS
Attributes	None	None
UT	Server	Server
DT	MultiSz	MultiSz

Data	{61738644-F196-11D0-9953-00C04FD919C1}	{9F0BD3A0-EC01-11D0-A6A0-00A0C922E752}
		, {61738644-F196-11D0-9953-00C04FD919C1}

5. In the View window, review the blue text for object properties that differ.
6. Click **Close** to close the object browser window.

Overview of Copying Objects to a Destination Server

After you have reviewed your audit results, you may need to copy certain objects to a destination server. Server Compliance allows you to copy the directories, files, Windows Services (state only), and Windows Registry keys. If an audit discovers that any of these objects exist only on the source or are on both the source and the target but differ, you can copy them to a destination server. The process works in the following manner:

- When you select a directory, only the directory will be copied to the destination server, excluding any files in that directory. For example, if dir1 contains file1 and file2, and you select dir1, Server Compliance copies only dir1 (not file1 and file2) to the destination server.
- When you select a file and its parent directory does not exist on the destination server, Server Compliance will create the directory on and copy the files to the destination server. For example, if you select file1 and dir1 does not exist on the destination server, Server Compliance will create dir1 on and copy file1 to the destination server.
- When you copy a Windows Services object, you copy the state of the service, such as started, stopped, paused, and so on. You can select one or more Windows Services objects for a single copy process.
- When you copy a Windows Registry object, you can select one or more registry keys and subkeys for a single copy process.

✓ You must have write permissions on the destination server to copy an object to it. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

Copying an Object from a Source or Target

To copy an object from a source or a target, perform the following steps:

1. From one of the starting points described in "Ways to View Audit Results" on page 353, open an audit result. The Audit Results Window appears.
2. In the Views pane, select an object in the File System.
3. In the Content pane, select the Only on Source or the On Both but Different tab.
4. Select one or more objects and select **Actions ➤ Copy To**.
5. In the Select Server window, select the server you want to copy the object to.

Use the search tool to dynamically filter this list by entering a server name, IP address, or operating system.

6 Click **Select** to copy the object or click **Cancel** to close this window without saving your changes.

For other types of server objects, such as packages and patches, you can also create installable packages to update a destination server. See "Visual Packager" on page 363 in Chapter 9 for more information.

Exporting Audit Results to HTML or CSV

If you need to use audit results in a different file format, you can export them to an HTML report or to a CSV file.

To export audit results to HTML or CSV, perform the following steps:

1 From one of the starting points described in "Ways to View Audit Results" on page 353, open an audit result. The Audit Results Window appears.

2 In the Audit Results window, select **Audit** ➤ **Export** ➤ **HTML**.

Or

Select **Audit** ➤ **Export** ➤ **CVS**.

3 In the Export window, enter a file name and specify the location to save this file.

4 Click **Open** (for HTML) or **Export** (for CVS) to save the file, or click **Cancel** to close the Export window without saving your changes.

Deleting Audit Results

As a best practice, you should delete audit results that you no longer need from the Software Repository to conserve disk space. You can delete audit results only if you performed the audit.

To delete audit results, perform the following steps:

1 From one of the starting points described in "Ways to View Audit Results" on page 353, select one or more audit results and select **Actions** ➤ **Delete**.

2 In the Confirmation Dialog, click **Yes** to delete or click **No** to close the window.

Chapter 9: Visual Packager

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Visual Packager
- Packaging Server Setup
- Overview of Packages
- Packaging Process
- Ways to Create a Package
- Creating a Package
- Adding New Package Content
- Specifying Options for New Package Content
- Viewing Package Details

Overview of Visual Packager

Visual Packager is an Opsware feature that helps you install software on managed servers. It guides you through the process of creating installable software packages using server compliance information, such as server snapshots and audit results. File system objects recorded in a snapshot and compliance information produced by an audit help you define the content of packages, and packages, in turn, can be used to update servers with new server objects.

Server objects can be selectively packaged according to the operating system of the servers that the package will be distributed to. Visual Packager supports Unix and Windows Operating Systems by allowing packages to contain the following objects:

- A Unix package can contain files (including attributes), directories, packages, patches, and patch clusters.
- A Windows package can contain files (including attributes), directories, packages, patches, Windows registry, and Windows services.

When you create a package, Visual Packager creates an application node in the Software Tree and verifies that all other packages and non-package content are attached to this same application node. Packages can consist of other packages, patches, and non-package content such as the Windows registry, Windows services, and file system objects. When you create a package that contains any of these items, Visual Packager analyzes the objects that you have selected and characterizes them in the following ways:

- When a package contains a package or a patch that does not exist in the Software Repository, Visual Packager lets you provide the missing information for selected packages and patches.
- When a package contains a package or a patch that already exists in the Software Repository and you want to overwrite them, Visual Packager lets you select the source location of the preferred version.
- When a package contains non-package content or patch data, Visual Packager creates a new package that contains them. You can specify a set of options that apply only to the non-package content, such as reboot requirements and pre/post install scripts. This non-package content is considered to be a package that is part of the same application node that the other packages and patches are attached to.

To get started using Visual Packager, select **Create Package** under the **Actions** menu in the OCC Client.

You must have RPM installed on your packaging server to enable the Visual Packager feature to create an RPM package for Solaris and AIX. The Visual Packager feature does not verify whether RPM is available on the packaging server.

Packaging Server Setup

Visual Packager requires a packaging server for each type of operating system for the packages you plan to create. For example, for Solaris packages you need a Solaris packaging server, and for MSI (Microsoft® Installer Utility) packages you need a Windows packaging server.

When you are using Visual Packager to create a package for a Red Hat Linux operating system, the operating system version and architecture of the packaging server must be identical to the operating system version and architecture that you want the package created for and installed on. Table 9-1 illustrates these requirements.

Table 9-1: Red Hat Linux Packaging Servers and Packages

OPERATING SYSTEM VERSION AND ARCHITECTURE OF THE PACKAGING SERVER	OPERATING SYSTEM VERSION AND ARCHITECTURE OF THE PACKAGE
Red Hat Enterprise Linux 3 AS 32 bit x86	Red Hat Enterprise Linux 3 AS 32 bit x86
Red Hat Enterprise Linux 3 AS 64 bit x86	Red Hat Enterprise Linux 3 AS 64 bit x86
Red Hat Enterprise Linux 3 ES 32 bit x86	Red Hat Enterprise Linux 3 ES 32 bit x86
Red Hat Enterprise Linux 3 ES 64 bit x86	Red Hat Enterprise Linux 3 ES 64 bit x86
Red Hat Enterprise Linux 3 WS 32 bit x86	Red Hat Enterprise Linux 3 WS 32 bit x86
Red Hat Enterprise Linux 3 WS 64 bit x86	Red Hat Enterprise Linux 3 WS 64 bit x86

The following installation and configuration tasks are required to set up a packaging server:

- You must first install the ISM Development Kit 2.0 (IDK) on a packaging server by using the Opsware Command Center (OCC). The packaging server must already be a managed server.
- You can then configure preferences for a packaging server by using the OCC Client itself.

The IDK must be installed on the packaging server by using a template in the OCC. Installing the IDK by using a template also enables the Visual Packager feature. See the *Opsware® SAS ISM Development Kit Guide* for more information.

Installing the IDK on a Packaging Server

To install the IDK on a packaging server, perform the following tasks:

1. Log into the OCC and click **Manage Servers** to select a server that you want to designate as your packaging server.

Figure 9-1: Installing the IDK Using a Template in the OCC

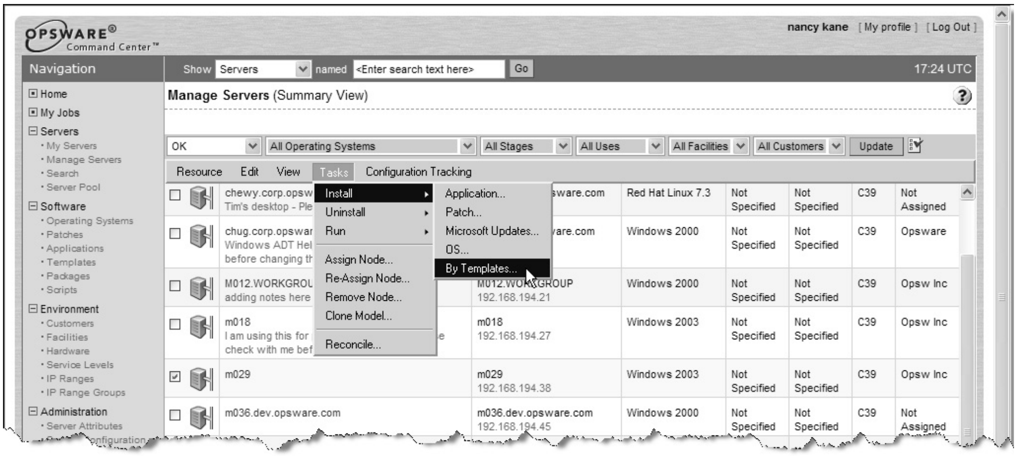

Show Servers named <Enter search text here>

Manage Servers (Summary View)

OK | All Operating Systems | All Stages | All Uses | All Facilities | All Customers

☐	chewy.corp.opsw Tim's desktop - Ple	sware.com	Red Hat Linux 7.3	Not Specified	Not Specified	C39	Not Assigned
☐	chug.corp.opswar Windows ADT Hel before changing th	are.com	Windows 2000	Not Specified	Not Specified	C39	Opsware
☐	M012.WORKGROU adding notes here	M012.WORKGROUP 192.168.194.21	Windows 2000	Not Specified	Not Specified	C39	Opsw Inc
☐	m018 I am using this for check with me bef	m018 192.168.194.27	Windows 2003	Not Specified	Not Specified	C39	Opsw Inc
☑	m029	m029 192.168.194.38	Windows 2003	Not Specified	Not Specified	C39	Opsw Inc
☐	m036.dev.opsware.com	m036.dev.opsware.com 192.168.194.45	Windows 2000	Not Specified	Not Specified	C39	Not Assigned

2. From the **Tasks** menu, select **Install ➤ By Templates**. Find the appropriate Visual Packager template in the Templates ➤ Opsware Tools ➤ Visual Packager ➤ <platform> Visual Packager node, and then install it. This declares the server that you selected in Step 1 as your packaging server by installing the IDK on it.

The Templates ➤ Opsware Tools ➤ Visual Packager ➤ <platform> Visual Packager node contains the platform-specific application node, including the ISMTool service level node. All nodes up to Visual Packager are customer independent and OS independent. <platform> is customer independent and OS specific.

Configuring Options for a Packaging Server

To configure a packaging server, perform the following tasks:

1. Launch the OCC Client.
2. From the **Tools** menu, select **Options**.
3. In the Set Options window, select Packaging Servers in the object tree to display a list of managed servers that already have the IDK installed on them. Not Configured in the Default Server column indicates that the server has not been set up as a packaging server for a specific operating system

Figure 9-2: Sample Set Options Window.

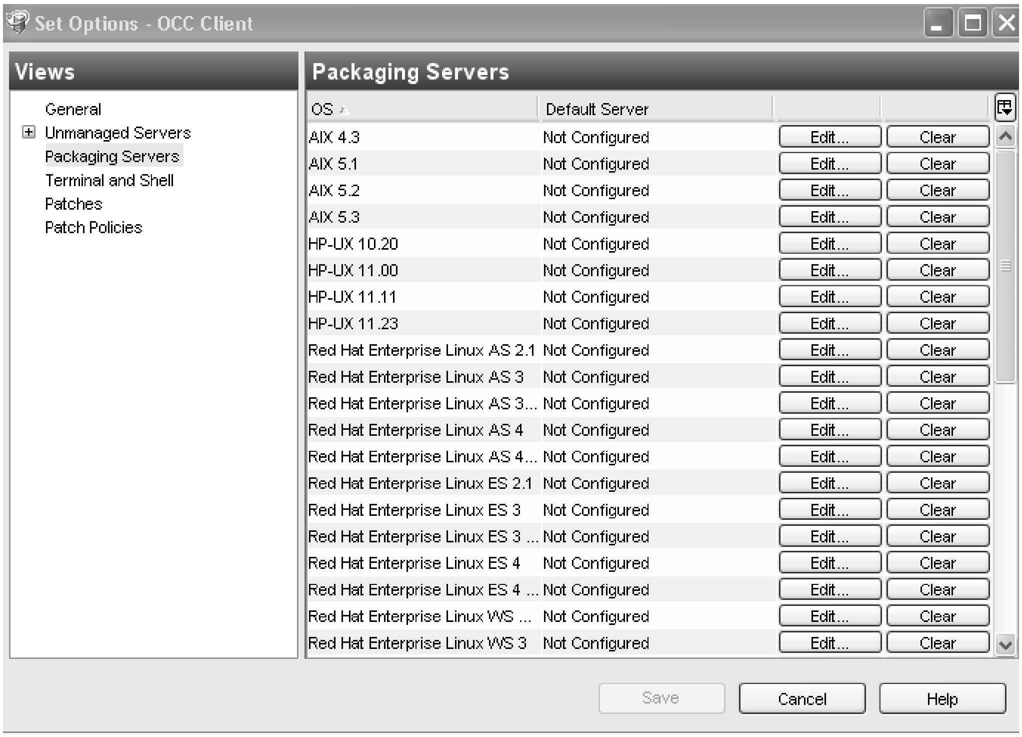

OS	Default Server		
AIX 4.3	Not Configured	Edit...	Clear
AIX 5.1	Not Configured	Edit...	Clear
AIX 5.2	Not Configured	Edit...	Clear
AIX 5.3	Not Configured	Edit...	Clear
HP-UX 10.20	Not Configured	Edit...	Clear
HP-UX 11.00	Not Configured	Edit...	Clear
HP-UX 11.11	Not Configured	Edit...	Clear
HP-UX 11.23	Not Configured	Edit...	Clear
Red Hat Enterprise Linux AS 2.1	Not Configured	Edit...	Clear
Red Hat Enterprise Linux AS 3	Not Configured	Edit...	Clear
Red Hat Enterprise Linux AS 3...	Not Configured	Edit...	Clear
Red Hat Enterprise Linux AS 4	Not Configured	Edit...	Clear
Red Hat Enterprise Linux AS 4...	Not Configured	Edit...	Clear
Red Hat Enterprise Linux ES 2.1	Not Configured	Edit...	Clear
Red Hat Enterprise Linux ES 3	Not Configured	Edit...	Clear
Red Hat Enterprise Linux ES 3 ...	Not Configured	Edit...	Clear
Red Hat Enterprise Linux ES 4	Not Configured	Edit...	Clear
Red Hat Enterprise Linux ES 4 ...	Not Configured	Edit...	Clear
Red Hat Enterprise Linux WS ...	Not Configured	Edit...	Clear
Red Hat Enterprise Linux WS 3	Not Configured	Edit...	Clear

4. In the content area, select an operating system and click **Edit** to display a list of managed servers for that operating system, or click **Clear** if you want to remove your edits.
5. In the Select Server window, select the managed server that you want to configure as the packaging server and then click **Select**.

The Select Server window will be empty if you do not have Read permission to the Service Level node. To obtain node stack permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

In the Set Options window, the Default Server column will now display the IP address of the server. Before you configured this as the packaging server, the Default Server column displayed Not Configured.

Use the search tool to dynamically filter by entering a server name, IP address, or operating system.

6 In the Set Options window, click **Save** to save your changes or click **Cancel** to close this window without saving your changes.

Overview of Packages

A package is installable software that includes server objects that contain applications, data, documentation, and configuration information for a managed server. These server objects can be files, directories, other packages and patches, Windows Registry, Windows Services, and so on. All packages are stored in the Software Repository.

Table 9-2 identifies the types of server objects you can include in a package, according to the operating system you plan to distribute the package to.

Table 9-2: Objects That Can Be Packaged

OBJECT TYPE	UNIX	WINDOWS
Files/Directories	Yes	Yes
Packages	Yes	Yes
Patches*	Yes	Yes
Windows Registry**	No	Yes
Windows Services	No	Yes
MTS/COM+	No	No
IIS Metabase	No	No

*Patches do not apply to Linux operating systems.

** You can package selected Windows registry keys, such as HKEY_CLASSES_ROOT, HKEY_CURRENT_CONFIG, HKEY_LOCAL_MACHINE, and HKEY_USERS.

Different types of server objects can be included in a package. For example, you can include multiple files, multiple patches, and multiple Windows services in one package. The IDK creates an installable package of these objects, in the native format of the operating system, such as Solaris Package for Unix, MSI for Windows, LPP for AIX, Depot for HPUX, and RPM for Linux.

In addition to selecting server objects that you want included in a package, you can also select packages and patches that already exist in Opsware SAS.

Packaging Process

The packaging process consists of several steps. The packaging process collects your package content, creates a background snapshot of the content, creates an executable wrapper for it, and then attaches it to an application node in the Software Tree. Figure 9-3 illustrates this workflow.

Figure 9-3: The Packaging Process

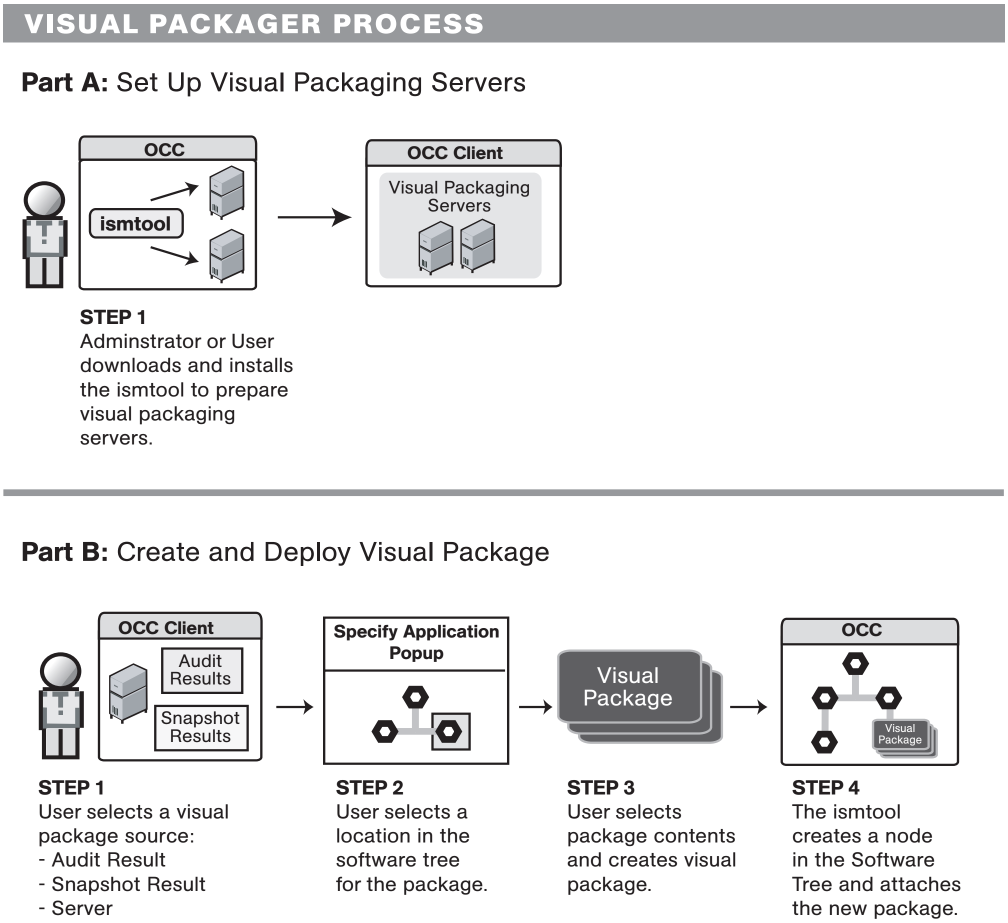

STEP 1
Adminstrator or User downloads and installs the ismtool to prepare visual packaging servers.

STEP 1
User selects a visual package source:
- Audit Result
- Snapshot Result
- Server

STEP 2
User selects a location in the software tree for the package.

STEP 3
User selects package contents and creates visual package.

STEP 4
The ismtool creates a node in the Software Tree and attaches the new package.

During the packaging process, a snapshot of your selected package content is transferred from the source server to the packaging server. This snapshot is used to create an Intelligent Software Module (ISM). The IDK (which must be installed on the packaging server) creates an executable wrapper for your package content. See the *Opsware® SAS ISM Development Kit Guide* for more information.

When you create a package, an application node is created in a user-defined, package location in the Software Tree. Your package is then attached to this node.

Ways to Create a Package

In Visual Packager, you can create a package in several different ways, depending on the content source. The content of a package is based on any of the following sources:

- An Opsware-managed server
- Audit results
- A server snapshot

Depending on the content source you select, you can launch the Visual Packager process from different windows in the OCC Client.

Creating a Package from a Managed Server

You can create a package from a managed server by using the following different windows in the OCC Client.

In the All Managed Servers window:

1. From the Navigation pane, select Servers and then select All Managed Servers.
2. From the Content pane, select one or more managed servers.
3. Select **Actions ➤ Create Package**.

In the Server Groups window:

1. From the Navigation pane, select Servers and then select Server Groups.
2. From the Content pane, select one or more servers in a server group.
3. Select **Actions ➤ Create Package**.

You *cannot* create a package for a server group. However, you can create a package for a server that is in a server group.

In the Server Explorer window:

1. From the Navigation pane, select Servers and then select All Managed Servers.

2. From the Content pane, select a managed server and open it.
3. From the Server Explorer, select **Actions ➤ Create Package**.

   Or

   Select a server object in the object tree, select an object in the content area, and then select **Actions ➤ Create Package**. If the object does not exist in the Software Repository, you will be prompted to provide the source.

Creating a Package from Audit Results

You can create a package from audit results by using the following different windows in the OCC Client.

In the Server Explorer window:

1. From the Navigation pane, select Servers and then select All Managed Servers.
2. From the Content pane, select a managed server and open it.
3. From the Server Explorer, select Compliance.
4. From the Content pane, select the Audit Results tab.
5. Select an audit and then select **Actions ➤ Create Package**.

In the Compliance window:

1. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Audit Results tab.
3. Select an audit and then select **Actions ➤ Create Package**.

In the Audit Results window:

1. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Audit Results tab.
3. Select an audit and then open it.
4. From the Audit Results window, select **Actions ➤ Create Package**.

Creating a Package from a Snapshot

You can create a package from a snapshot by using the following different windows in the OCC Client.

In the Server Explorer window:

1. From the Navigation pane, select Servers and then select All Managed Servers.
2. From the Content pane, select a managed server and open it.
3. From the Server Explorer, select Compliance.
4. From the Content pane, select the Snapshots tab.
5. Select a snapshot and then select **Actions ➤ Create Package**.

In the Snapshot Browser window:

1. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Snapshots tab.
3. Select a snapshot and then open it.
4. From the Snapshots Browser window, select **Actions ➤ Create Package**.

   Or

   In the object tree, select a server object and right-click. Only objects that were included in the selection criteria (in the snapshot template) are displayed.

In the Compliance window:

1. From the Navigation pane, select Software Library and then select Compliance.
2. From the Content pane, select the Snapshots tab.
3. Select a snapshot and then select **Actions ➤ Create Package**.

Creating a Package

You can create installable packages in native formats, such as Solaris Package and MSI. When you create an installable package, you must specify where you want it uploaded in the Software Tree. You can create a new application node or add the package to an existing node in the Software Tree. When you add a package to an existing node in the Software Tree, the contents of your new application will overwrite the contents of the existing application node. See the *Opsware® SAS Configuration Guide* for information about how to upload packages to and manage packages in the Software Tree.

Optionally, you can assign your package to a customer and specify the operating system versions that you want this package installed on.

To create a package, perform the following steps:

1. From one of the starting points described in “Ways to Create a Package”, select Create Package.

   ⚠ If you have not already set up your packaging server, a warning icon displays in the Create Package window. See “Packaging Server Setup” on page 364 in Chapter 9 for more information.

✓ You must have a set of permissions to manage selection criteria. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

2 In the Create Package window, select the Details tab.

Figure 9-4: Sample Create Package Window (Details Tab)

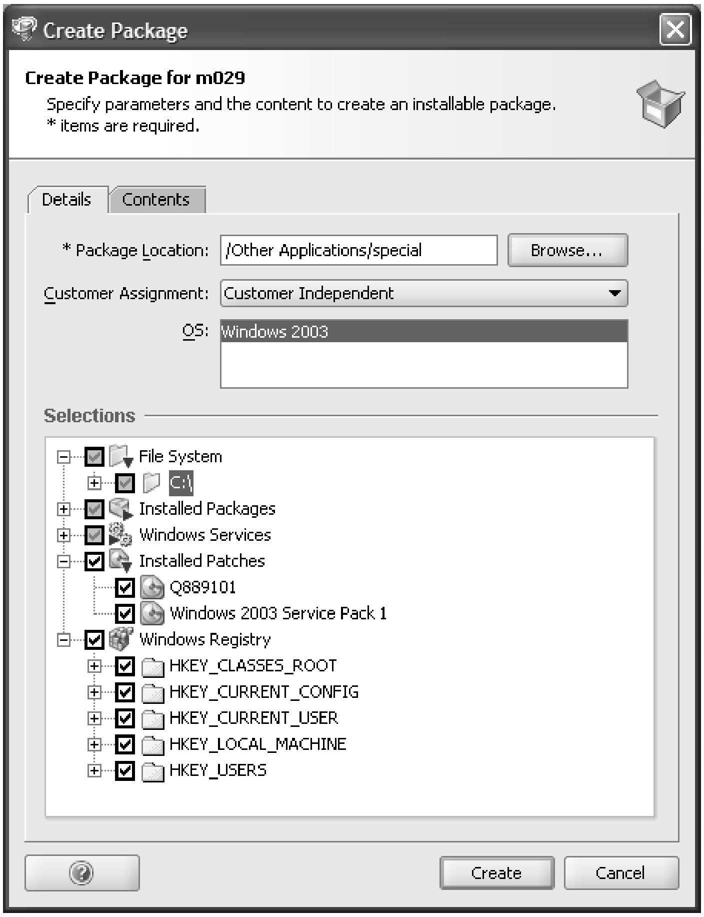

Create Package

Create Package for m029

Specify parameters and the content to create an installable package.
* items are required.

* Package Location: /Other Applications/special

Customer Assignment: Customer Independent

OS: Windows 2003

Selections

- ☑ File System
  - ☑ C:\
- ☑ Installed Packages
- ☑ Windows Services
- ☑ Installed Patches
  - ☑ Q889101
  - ☑ Windows 2003 Service Pack 1
- ☑ Windows Registry
  - ☑ HKEY_CLASSES_ROOT
  - ☑ HKEY_CURRENT_CONFIG
  - ☑ HKEY_CURRENT_USER
  - ☑ HKEY_LOCAL_MACHINE
  - ☑ HKEY_USERS

3 Next to the Package Location field, click **Browse.** This is a required field, as indicated by the asterisk (*). The Specify Application window appears.

4 In the Specify Application window, specify the location in the Software Tree where you want to create an application node by completing the following steps:

To navigate back up the hierarchy, modify the Look In field.

1. Enter a new application node Name. The backslash character (/) denotes a new hierarchy.
2. Click **OK** to add the new application node.

If you enter a new application node name that is identical to a node name that already exists in the Software Tree, the contents of your new application will overwrite the contents of the existing application node.

See the *Opsware® SAS Configuration Guide* for more information about the categories used in the top level of the Software Tree.

5 In the Customer Assignment field, select one of the following customer types to assign to your new package:

- **Customer Independent**: A global customer in Opsware SAS. Resources (applications, patches, and templates) that are associated with Customer Independent can be installed on any managed server, no matter what customer it is associated with.
- **Not Assigned**: The servers are not associated with a customer. You can install applications, patches, or templates that are Customer Independent on Not Assigned servers. However, you cannot install or use any resources associated with a customer on a server that is not assigned to a customer.
- **Other**: The name of this customer assignment varies, because it is created by the customer and is derived from the customer's environment.

If you modify the customer assignment and it does not correspond to the package location that is currently specified, Create Package will instruct you to also change the package location so that it corresponds to the new customer assignment.

6 In the OS field, select one or more operating system versions that you want to install this package on. The operating systems in this list represent the operating system family that you previously configured as a packaging server. You can only select from this list of operating system versions. See "Configuring Options for a Packaging Server" on page 367 in this chapter for more information.

7 In the Selections field, select the server objects that you want to include in your package, such as File System, Installed Packages, Installed Patches, Windows Services, and so on. If you select installed packages or installed patches that need to

be added to the Software Repository, Visual Packager will instruct you to select the Contents tab to perform this task. See “Adding New Package Content” on page 377 in this chapter for more information.

For Windows Services, you are selecting only the state of the service, such as Started, Stopped, Paused, and so on.

8. Click **Create** to create the package and save it in a new or existing application node in the Software Tree, or click **Cancel** to close this window without creating a package.

Adding New Package Content

After you have selected packages, patches, or non-package content, Create Package examines them to determine if additional information is required. At this point, Create Package also allows you to pick and choose which packages you want to include or exclude, and which non-package content objects will be packaged.

To add new package content, perform the following steps:

1. From one of the starting points described in “Ways to Create a Package”, select Create Package.

2 In the Create Package window, select the Contents tab.

Figure 9-5: Sample Create Package Window (Contents Tab)

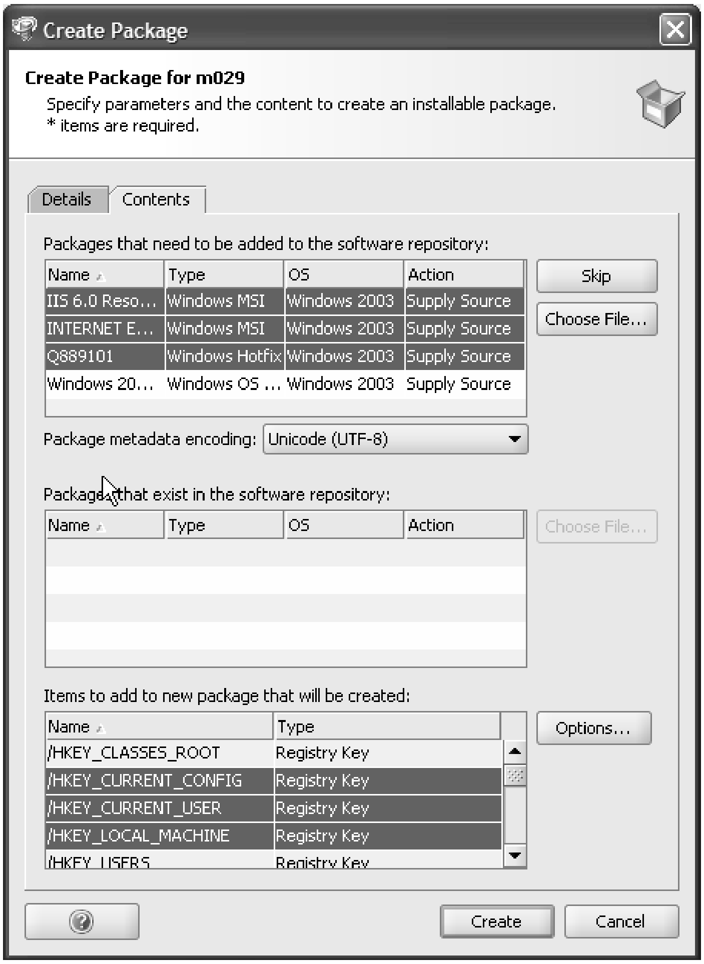

Create Package for m029

Specify parameters and the content to create an installable package.
* items are required.

Packages that need to be added to the software repository:

Name	Type	OS	Action
IIS 6.0 Reso...	Windows MSI	Windows 2003	Supply Source
INTERNET E...	Windows MSI	Windows 2003	Supply Source
Q889101	Windows Hotfix	Windows 2003	Supply Source
Windows 20...	Windows OS ...	Windows 2003	Supply Source

Package metadata encoding: Unicode (UTF-8)

Packages that exist in the software repository:

Name	Type	OS	Action

Items to add to new package that will be created:

Name	Type
/HKEY_CLASSES_ROOT	Registry Key
/HKEY_CURRENT_CONFIG	Registry Key
/HKEY_CURRENT_USER	Registry Key
/HKEY_LOCAL_MACHINE	Registry Key
/HKEY_USERS	Registry Key

3 In the "Packages that need to be added to the software repository" section, a list of packages or patches (previously selected in the Details tab) that the Software Repository does not have yet is displayed. Create Package requires that you specify the source location of these packages or patches so that they will be uploaded to the Software Repository. Without this source information, Create Package cannot create a package. The following information about these packages is displayed:

- **Name**: The name of the package.
- **Type**: The type of package, such as Sol_Pkg, MIS, RPM, and so on.
- **OS**: The operating system the package is intended for.

- **Action**: If the action status is Supply Source, you must specify the location of the package so that it will be included in the upload process.

You can pick and choose which packages and patches you want to provide source information for by performing the following steps:

1. To exclude particular packages from the upload process, select one or more packages and click **Skip** to exclude them from the upload process. These packages will *not* be attached to the application node.
2. To include a package in the upload process, select a package and click **Choose File** to locate it in your local file system or in the network file system of the packaging server. This package will be attached to the application node.

You can also specify the character encoding for the metadata of Unix packages that are displayed in the “Packages that need to be added to the software repository” section. Package metadata includes comments, READMEs, scripts, descriptions, and content lists. Before storing the package internally, Opsware SAS translates the metadata from the specified encoding into UTF-8. The default encoding is determined by the source of the package content, either a managed server, server snapshot, or audit results.

To specify character encoding for Unix package metadata, from the “Package metadata encoding” list, select a character encoding, such as Shift-JIS for Japanese. (For Windows packages, the encoding is set to UTF-8.)

4 In the “Packages that exist in the software repository” section, a list of packages or patches (previously selected in the Details tab) that currently exist in the Software Repository is displayed. If you want to overwrite these packages and patches with different versions, you must provide the source locations of them. Different package and patch versions are typically ones that you have previously used (and therefore know) are the ones you need to include in the package.

The following information about these packages is displayed:

- **Name**: The name of the package as it appears in the Software Tree.
- **Type**: The type of package as it exists in the Software Tree, such as Sol_Pkg, MIS, RPM, and so on.
- **OS**: The operating system the package is intended for, such as SunOS 5.8.
- **Action**: If the action status is Supply Source, you must specify the location of the package so that it will be included in the upload process.

To choose a package or patch that you want to overwrite the version in the Software Repository, select one and then click **Choose File** to point to its storage location in your local file system or in the network file system of the packaging server.

5 In the "Items to add to new package that will be created" section, a list of non-package content (previously selected in the Details tab) is displayed. By default, Create Package creates a new package that contains all of these items. You can specify a separate set of options that apply only to these non-package content items, such as reboot requirements and pre/post install and pre/post uninstall scripts. This non-package content is considered to be a package that is part of the same application node that the other packages (packages and patches) are attached to.

The following information about these packages is displayed:

- **Name**: The name of the object, such as a directory, file, Windows Registry, Windows Services, and so on.

- **Type**: The type of object.

Click **Options** to display the New Package Options window where you can specify additional settings for all items in the new package content. Visual Packager enables **Options** only if there is additional package content. See "Specifying Options for New Package Content" on page 380 in this chapter for more information.

6 Click **Create** to attach this new package content to the application node, or click **Cancel** to close this window without attaching any new package content to the application node.

Specifying Options for New Package Content

When you add new package content to a package that you have already created, you can specify additional settings that apply only to the new package content. These settings include which customer types this package applies to, whether to reboot after your package is installed or uninstalled, which pre/post install and pre/post uninstall scripts will be run, and so on.

To specify options for new package content, perform the following steps:

1 Follow the instructions in step 5 in "Adding New Package Content" on page 377 to display the New Package Options window.

2 In the New Package Options window, select or enter your preferred options.

Figure 9-6: Sample New Package Options Window

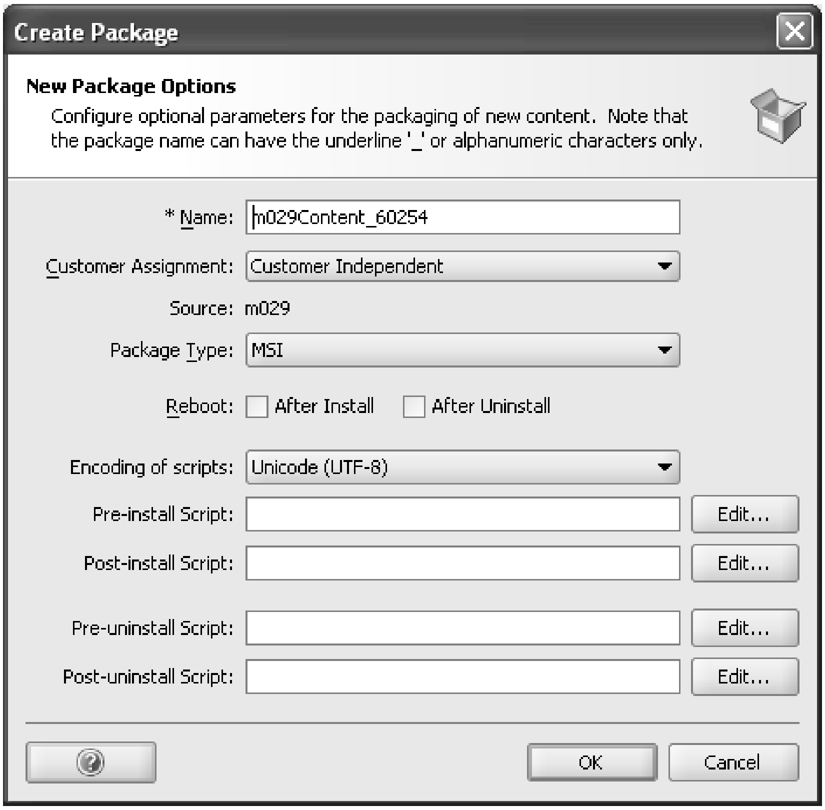

Create Package

New Package Options

Configure optional parameters for the packaging of new content. Note that the package name can have the underline '_' or alphanumeric characters only.

* Name: m029Content_60254

Customer Assignment: Customer Independent

Source: m029

Package Type: MSI

Reboot: □ After Install □ After Uninstall

Encoding of scripts: Unicode (UTF-8)

Pre-install Script: Edit...

Post-install Script: Edit...

Pre-uninstall Script: Edit...

Post-uninstall Script: Edit...

OK Cancel

3 In the Name field, Visual Packager presets the name of the package. However, if required, you can carefully modify the package name. The following rules apply to package names:

- For Solaris packages, the package name must not exceed five characters, such as 123_S.
- For all other types of packages, the package name must not exceed 64 characters, such as m123acmecomContent_89787. These types of package names are derived from the source name with the text Content and a 5-digit timestamp appended to it.
- For all package names, alphanumeric characters and the underscore (_) character are valid.

This is a required field, as indicated by the asterisk (*).

4 In the Customer Assignment field, select one of the following customer types you want to assign to your new package content:

- **Customer Independent**: A global customer in Opsware SAS. Resources (applications, patches, and templates) that are associated with Customer Independent can be installed on any managed server, no matter what customer it is associated with.
- **Not Assigned**: The servers are not associated with a customer. You can install applications, patches, or templates that are Customer Independent on Not Assigned servers. However, you cannot install or use any resources associated with a customer on a server that is not assigned to a customer.
- **Other**: The name of this customer assignment varies, because it is created by the customer and is derived from the customer's environment.

If you modify the customer assignment and it does not correspond to the package location that is currently specified, Create Package will instruct you to also change the package location so that it corresponds to the new customer assignment.

5 The Source field identifies the name of the source of the package. A package source can be a managed server, a snapshot, or an audit result. You cannot modify this field.

6 In the Package Type drop-down list, you can only select package formats for Solaris Package and LPP (for AIX). Create Package presets the package type for all MSI (for Windows), Depot (for HPUX), and RPM (for Linux) packages.

You must have RPM installed on your packaging server to enable the Visual Packager feature to create an RPM package for Solaris and AIX. The Visual Packager feature does not verify whether RPM is available on the packaging server.

7 In the Reboot field, perform any of the following steps:

- Check After Install to specify that the server will be rebooted after your package is installed. The default is an unchecked reboot option.
- Check After Uninstall to specify that the server will be rebooted after your package is uninstalled. The default is an unchecked reboot option.

Reboot settings are saved in the Model Repository.

8 In the “Encoding of scripts” list, select a character encoding, such as Shift-JIS for Japanese.

This encoding applies to the contents of the pre/post install and pre/post uninstall scripts when they are deployed on a managed server. The default encoding is determined by the source of the package content. Before installing the package on a managed server, Opsware SAS converts the scripts into the specified encoding. Internally, Opsware SAS stores the scripts in UTF-8.

For Linux, if a script contains non-ASCII characters, then include the shell execution command (such as `#! /bin/sh`) at the beginning of the script.

9 In pre/post install and pre/post uninstall scripts section, you can perform any of the following actions:

- To specify that a script is run before the package is installed, enter the Pre-install Script or click **Edit** to enter the script in the Specify Post-install Script window.
- To specify that a script is run after the package is installed, enter the Post-install Script or click **Edit** to enter the script in the Specify Post-install Script window.
- To specify that a script is run before the package is uninstalled, enter the Pre-uninstall Script or click **Edit** to enter your script in the Specify Pre-uninstall Script window.
- To specify that a script is run after the package is uninstalled, enter the Post-uninstall Script or click **Edit** to enter your script in the Specify Post-uninstall Script window.

Pre/post install and pre/post uninstall scripts are stored in the Model Repository.

Create Package cannot verify that a script will successfully execute. Opsware, Inc. recommends that you test all scripts before you add them as package options.

10 Click **OK** to save your new package options, or click **Cancel** to close this window without saving your new package options.

Viewing Package Details

You can use the Job Manager in the OCC Client to review the following information about your package process:

- Job ID, start and end time, the name of the packaging server, the status of the packaging process, and so on.
- Log details about the package creation process, such as error descriptions.

Chapter 10: Global Shell

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Global Shell
- Remote Terminal
- Global Shell Session
- Opsware Global File System (OGFS)
- Directories in the OGFS
- Server Filtering in the OGFS
- Remote Opsware Shell (rosh)
- Character Encoding for OGFS
- Error Messages

Overview of Global Shell

The Opsware Global Shell feature enables you to remotely perform the following tasks:

- Complete routine maintenance tasks on managed servers.
- Troubleshoot, identify, and remediate problems on managed servers.

The Global Shell consists of a file system and a command-line interface to that file system for managing servers in Opsware SAS. The file system is known as the Opsware Global File System (OGFS). All object types in the OGFS (such as servers, customers, and facilities) are represented as directory structures in this file system. Global Shell enables you to manage user permissions for accessing the file system, Windows Registry, and Windows Services objects on managed servers, and it provides an audit trail that records the events performed on managed servers. Opsware SAS also provides recommended scripts that you can run from the Global Shell interface; for more information on these scripts see

Remote Terminal

The Remote Terminal feature opens a terminal window for each selected Unix server or an RDP client window for each selected Windows server.

You must have login permissions on the managed server to use this feature. See the *Opsware® SAS Configuration Guide* for more information.

In the OCC Client, the Remote Terminal feature is accessible from the Server list and Server Explorer windows. In the Server list window, select a managed server and select **Remote Terminal**. In the Server Explorer, to access the File System, Windows Services, and Windows Registry nodes of a managed server, you must first select a node, and then right-click and select a valid user group, such as root, Administrator, or LocalSystem.

You can specify your terminal and RDP client preferences for a Remote Terminal or Global Shell session, such as `telnet` for a Windows terminal, `xterm` for a Unix terminal, `rdesktop` for an RDP client, and so on, in the Set Preferences window. See "Terminal and Shell" on page 110 in Chapter 3 for more information.

Opening a Remote Terminal

The type of remote terminal you open for a managed server depends on the operating system of the server. For a Unix server, a terminal window is opened. For a Windows server, an RDP session is opened. (You cannot login to a remote terminal with a user name or password that contains multi-byte characters.)

For a Unix Server:

You can open a terminal window for a Unix server from the Server list window and from the Server Explorer.

In the Server list window, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and then select Managed Servers.
2. Select a managed Unix server.
3. From the **Actions** menu, select **Remote Terminal**.

In the Server Explorer (for a Unix server), perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and then select Managed Servers.
2. Select a managed Unix server. Open it. The Server Explorer window opens.
3. Select File System in the object tree, right-click, and then select a login name (such as root or Administrator). See the *Opsware® SAS Configuration Guide* for more information.
4. Select **Actions ➤ Remote Terminal**.

For a Windows Server:

You can open an RDP session for a Windows server from the Server list window and from the Server Explorer.

In the Server list window, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and then select Managed Servers.
2. Select a managed Windows server.
3. Select a Windows server, and then select **Actions ➤ Remote Terminal**.

In the Server Explorer (for a Windows server), perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and then select Managed Servers.
2. Select a managed Windows server. Open it. The Server Explorer window opens.
3. Select File System in the object tree, right-click, and then select a login name (such as LocalSystem). See the *Opsware® SAS Configuration Guide* for more information.
4. Select **Actions ➤ Remote Terminal**.

Global Shell Session

The Global Shell feature enables you to open a terminal window for the OGFS in Opsware SAS.

Opening a Global Shell Session

You can open a Global Shell session with an `ssh` client or from within the OCC Client. When you open a session, the working directory is `/home/`*`user-name`*.

To open a Global Shell session with an ssh client, perform the following steps:

1. On a host that is not an Opsware core server or a managed server, open a terminal window.

2. In the terminal window, enter an ssh command with the following syntax:

   ```
   ssh -p 2222 user-name@ogfs-host
   ```

   To use this command, port 2222 must be open on the firewall that protects the OGFS server. The *user-name* is your Opsware user and the *ogfs-host* is the host name (or IP address) of the core server running the OGFS. After you enter the ssh command, the OGFS prompts for the password of the Opsware user.

To open a Global Shell session from within the OCC Client, from the **Actions** menu, select **Global Shell**.

Copying Files To and From the OGFS

You can securely copy files between the OGFS and a server that is not part of Opsware SAS. To copy the files, perform the following steps:

1. On a host that is not an Opsware core server or a managed server, open a terminal window.

2. In the terminal window, enter either the scp, sftp, or rsync command and specify port 2222, your Opsware user name, and the host running the OGFS.

The following three scp examples perform the same operation: They copy the file myscript.sh from the local machine to the file /home/jdoe/myscript.sh in the OGFS. The Opsware user is jdoe and the host running the OGFS is 192.168.166.178.

```
scp -P 2222 myscript.sh jdoe@192.168.166.178:myscript.sh
scp -P 2222 myscript.sh jdoe@192.168.166.178:/home/jdoe
scp -P 2222 myscript.sh jdoe@192.168.166.178:
```

The following example copies myscript.sh from the home directory of jdoe in the OGFS to the local machine:

```
scp -P 2222 jdoe@192.168.166.178:myscript.sh myscript.sh
```

The following sftp example copies myscript.sh from the local machine to the OGFS:

```
sftp -oPort=2222 jdoe@192.168.166.178
Connecting to 192.168.166.178...

Opsware Global Shell
```

```
jdoe@opsware's password:
sftp> put myscript.sh

....
```

The following `rsync` example transfers files from `/path` on the local machine to `/other/path` in the OGFS:

```
rsync -av -e "ssh -p 2222" /path \
jdoe@192.168.166.178:/other/path
```

Opsware Global File System (OGFS)

The OGFS is a tree hierarchy of directories and files (objects) on all managed servers in Opsware SAS. The structures of these directories range from a list of objects to a hierarchical structure of objects, which can contain lists of related objects. To optimally navigate through the OGFS to get to a destination of your choice it is important to understand these directory structures.

The following example illustrates how to search a directory that contains a hierarchical structure of different types of related objects:

To find all servers that are associated with facilities and customers, specify the following path to search the `Server` directory:

```
$ ls /opsw/Server/@/
```

With regard to a destination of your choice, the Opsware Global Shell feature recognizes that you may not always be searching for a server. Therefore, hierarchical directory structures allow you to search by a variety of object types.

When you add or edit server custom attributes in the Opsware Global File System (using Global Shell), Opsware SAS preserves leading and trailing whitespace characters in custom attribute values.

Directories in the OGFS

The following directories in the OGFS are best explained by structure diagrams and supporting text:

- root Directory
- opsw Directory
- Facility Directory
- Group Directory
- Server Directory
- Script Directory
- Application Directory
- Customer Directory
- OS Directory
- Realm Directory
- ServiceLevel Directory

Diagram Conventions

The structure diagrams use the following typographical and formatting conventions, as shown in Table 10-1.

Table 10-1: Opsware Global File System (OGFS) Diagram Conventions

NOTATION	DESCRIPTION
Times	The name of the object type as it appears in the OGFS, such as bin, Group, Facility, Server, Customer, and so on.
`Courier`	An object in the OGFS that is user-defined, such as `server-1-name`, `group-1-name`, `facility-1-name`, `custom-attribute-1-name`, and so on.
`Courier Italics`	A brief description of the corresponding object and its relationship to other objects in the OGFS. Used only in the Server directory diagram.

root Directory

At the root level of the OGFS is a directory for each Opsware user. Each user directory and all files and directories under it are visible only to processes in an authenticated session. Figure 10-1 illustrates the root directory structure.

Figure 10-1: root Directory

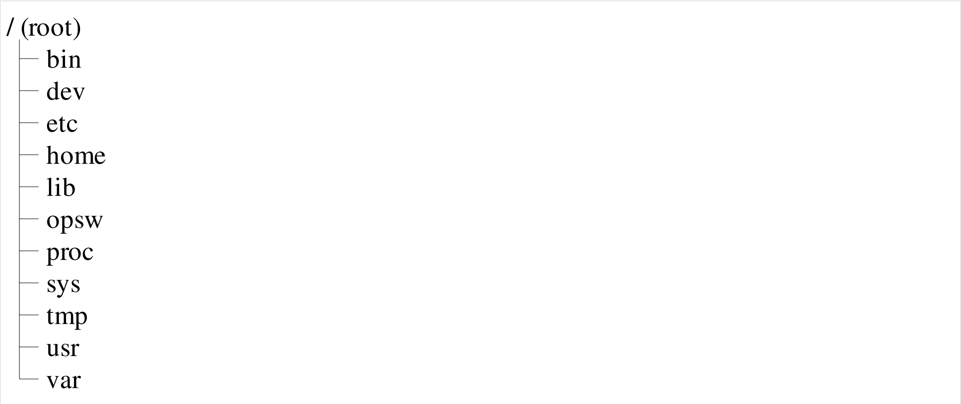

```
/ (root)
 ├─ bin
 ├─ dev
 ├─ etc
 ├─ home
 ├─ lib
 ├─ opsw
 ├─ proc
 ├─ sys
 ├─ tmp
 ├─ usr
 └─ var
```

Each Opsware user has the following private directories:

- A home directory which is at /home/*user-name* (Opsware user name)
- Temporary directories which are at /tmp, /var/tmp and /usr/tmp

Each user's home directory contains a public directory (/home/*user-name*/public) which is readable by all other Opsware users and can be used to share files with other Opsware users.

opsw Directory

The opsw directory is also known as the Opsware model space. The Application, Customer, Facility, Group, OS, Realm, Script, Server, and ServiceLevel directories represent objects in the Opsware Model Space. Applications can navigate through this model space. The bin directory contains only Opsware SAS programs.

Figure 10-2 illustrates the structure of the `opsw` directory.

Figure 10-2: opsw Directory

```
/opsw
 ├─ bin
 ├─ Application
 ├─ Customer
 ├─ Facility
 ├─ Group
 ├─ OS
 ├─ Realm
 ├─ Script
 ├─ Server
 └─ ServiceLevel
```

Facility Directory

The `Facility` directory contains a list of facilities and facility details. Figure 10-3 illustrates the structure of the `Facility` directory.

Figure 10-3: Facility Directory

```
/opsw/Facility
 ├─ facility-1-name
 │   └─ @
 │       ├─ .id
 │       ├─ info
 │       └─ CustAttr
 │           ├─ custom-attribute-1-name
 │           ├─ custom-attribute-2-name
 │           ┆
 ├─ facility-2-name
 ┆
```

Group Directory

The `Group` directory contains `Public` and `Private` user groups, user group details (where @ represents the group itself), child groups, and grandchild groups. This structure essentially contains groups within groups. Figure 10-4 illustrates the structure of the `Group` directory.

Figure 10-4: Group Directory

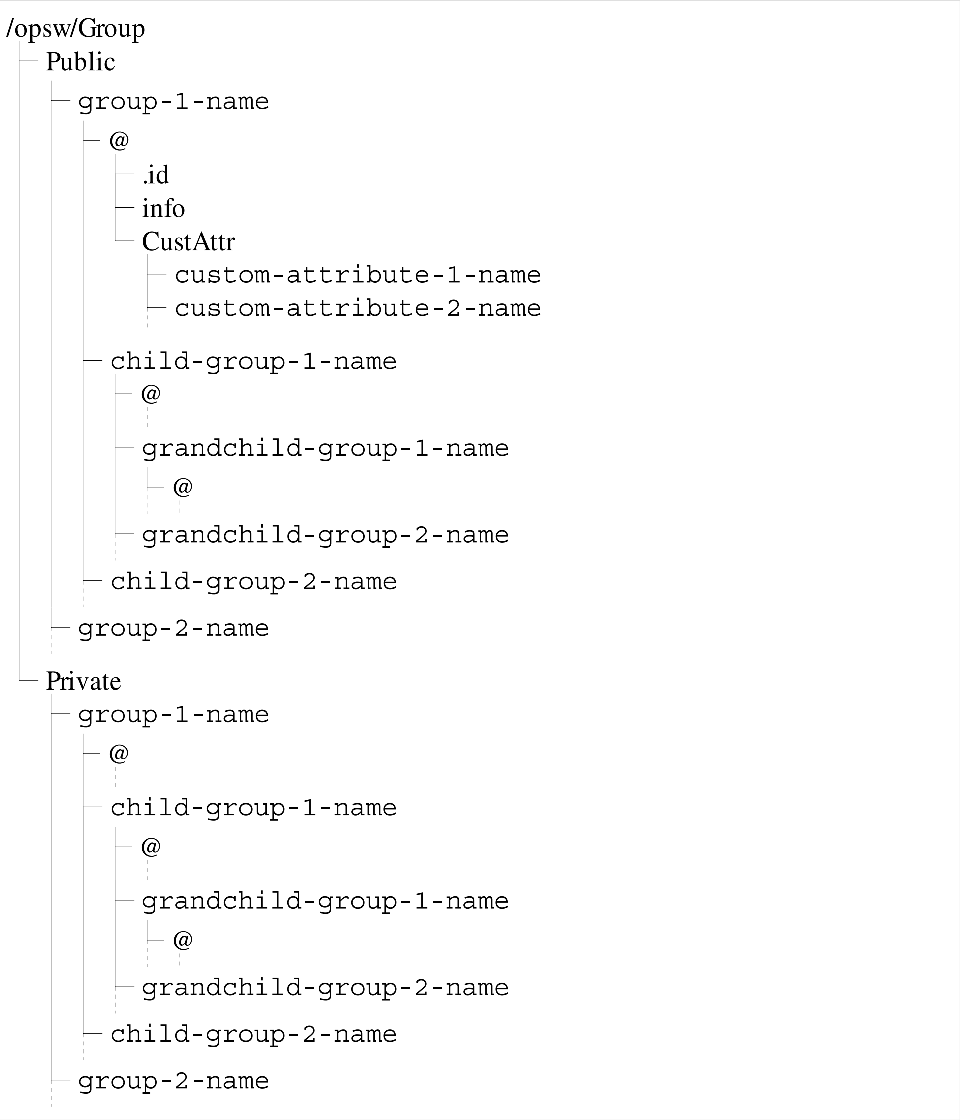

```
/opsw/Group
 ├─ Public
 │   ├─ group-1-name
 │   │   ├─ @
 │   │   │   ├─ .id
 │   │   │   ├─ info
 │   │   │   └─ CustAttr
 │   │   │       ├─ custom-attribute-1-name
 │   │   │       ├─ custom-attribute-2-name
 │   │   │       ┆
 │   │   ├─ child-group-1-name
 │   │   │   ├─ @
 │   │   │   ┆
 │   │   │   ├─ grandchild-group-1-name
 │   │   │   ┆   ├─ @
 │   │   │   ┆   ┆
 │   │   │   ├─ grandchild-group-2-name
 │   │   │   ┆
 │   │   ├─ child-group-2-name
 │   │   ┆
 │   ├─ group-2-name
 │   ┆
 └─ Private
     ├─ group-1-name
     │   ├─ @
     │   ┆
     │   ├─ child-group-1-name
     │   │   ├─ @
     │   │   ┆
     │   │   ├─ grandchild-group-1-name
     │   │   ┆   ├─ @
     │   │   ┆   ┆
     │   │   ├─ grandchild-group-2-name
     │   │   ┆
     │   ├─ child-group-2-name
     │   ┆
     ├─ group-2-name
     ┆
```

Server Directory

The `Server` directory contains all managed servers and different types of related objects. To optimally navigate this directory and find servers by group, by facility, by customer, and so on, you can use a server filter. See "Server Filtering in the OGFS" on page 400 in this chapter for more information.

Figure 10-5 illustrates the structure of the `Server` directory.

There are lower server directories (such as *`server-1-name`*, *`server-2-name`*, and so on) within the top-level `Server` directory. When you modify files in the lower directories, you do so as a specific user (such as *`login-1-name`*). Modifying files includes the adding, deleting, and editing actions.

Text in parentheses in *`(Courier Italics)`* denotes a brief description of the corresponding object and its relationship to other objects in the OGFS.

Figure 10-5: Server Directory

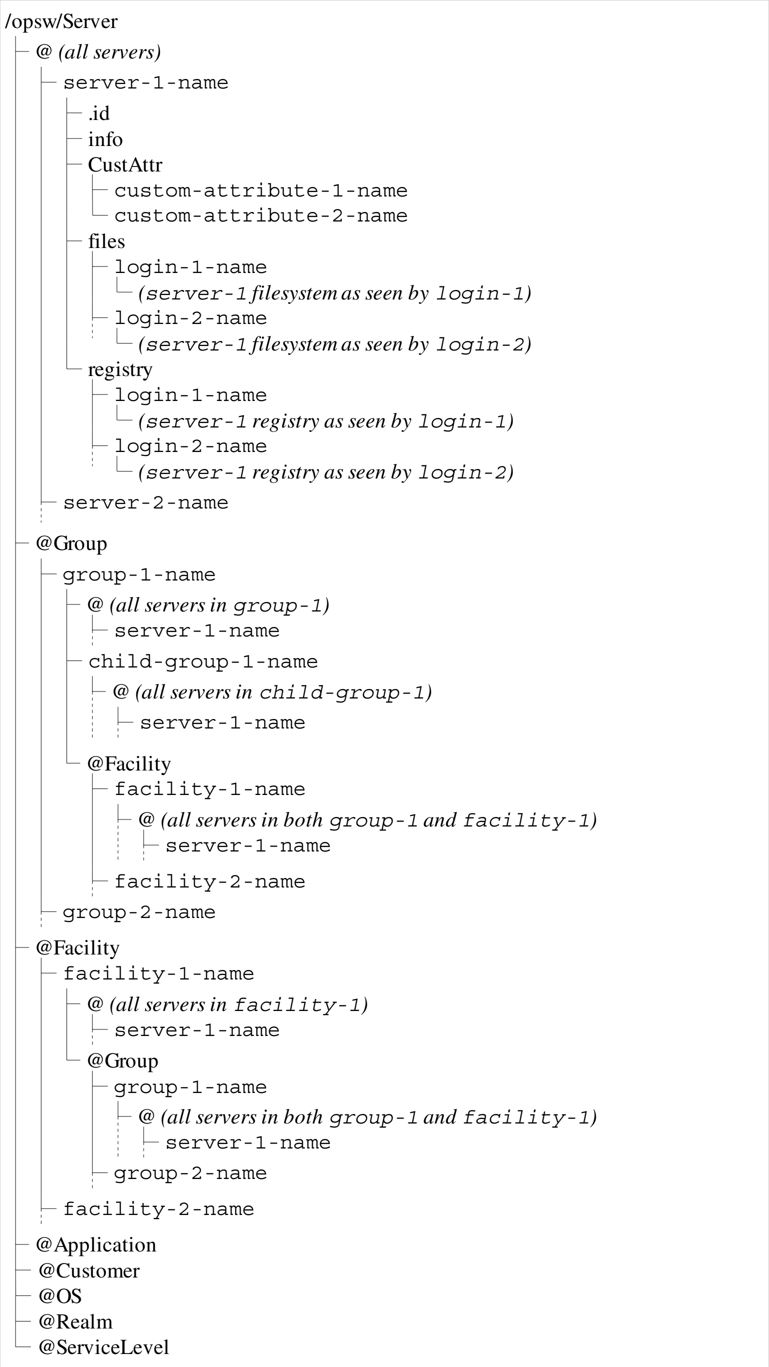

```
/opsw/Server
├─ @ (all servers)
│  ├─ server-1-name
│  │  ├─ .id
│  │  ├─ info
│  │  ├─ CustAttr
│  │  │  ├─ custom-attribute-1-name
│  │  │  └─ custom-attribute-2-name
│  │  ├─ files
│  │  │  ├─ login-1-name
│  │  │  │  └─ (server-1 filesystem as seen by login-1)
│  │  │  ├─ login-2-name
│  │  │  ┊  └─ (server-1 filesystem as seen by login-2)
│  │  └─ registry
│  │     ├─ login-1-name
│  │     │  └─ (server-1 registry as seen by login-1)
│  │     ├─ login-2-name
│  │     ┊  └─ (server-1 registry as seen by login-2)
│  ├─ server-2-name
│  ┊
├─ @Group
│  ├─ group-1-name
│  │  ├─ @ (all servers in group-1)
│  │  │  ├─ server-1-name
│  │  ├─ child-group-1-name
│  │  │  ├─ @ (all servers in child-group-1)
│  │  │  ┊  ├─ server-1-name
│  │  │
│  │  └─ @Facility
│  │     ├─ facility-1-name
│  │     │  ├─ @ (all servers in both group-1 and facility-1)
│  │     │  ┊  ├─ server-1-name
│  │     │
│  │     ├─ facility-2-name
│  ├─ group-2-name
│  ┊
├─ @Facility
│  ├─ facility-1-name
│  │  ├─ @ (all servers in facility-1)
│  │  │  ├─ server-1-name
│  │  └─ @Group
│  │     ├─ group-1-name
│  │     │  ├─ @ (all servers in both group-1 and facility-1)
│  │     │  ┊  ├─ server-1-name
│  │     ├─ group-2-name
│  │
│  ├─ facility-2-name
│  ┊
├─ @Application
├─ @Customer
├─ @OS
├─ @Realm
└─ @ServiceLevel
```

Script Directory

The `Script` directory contains a list of shared scripts and script details. Figure 10-6 illustrates the structure of the `Script` directory.

Figure 10-6: Script Directory

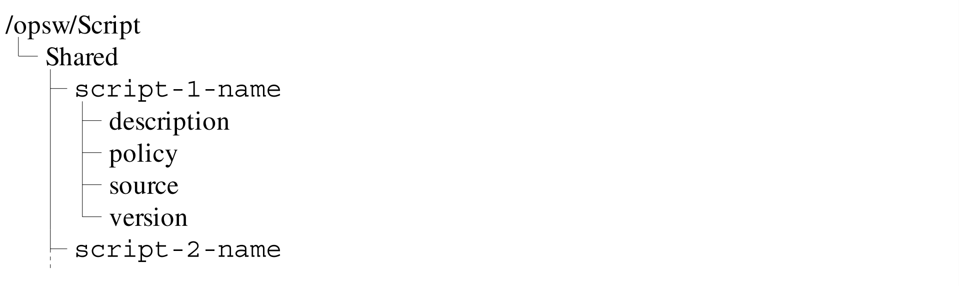

```
/opsw/Script
 └─ Shared
     ├─ script-1-name
     │   ├─ description
     │   ├─ policy
     │   ├─ source
     │   └─ version
     ├─ script-2-name
     ┆
```

Application Directory

The `Application` directory contains applications, application details, child applications, and grandchild applications. This directory structure is similar to the application node hierarchy in the Software Tree. Figure 10-7 illustrates the structure of the `Application` directory.

Figure 10-7: Application Directory

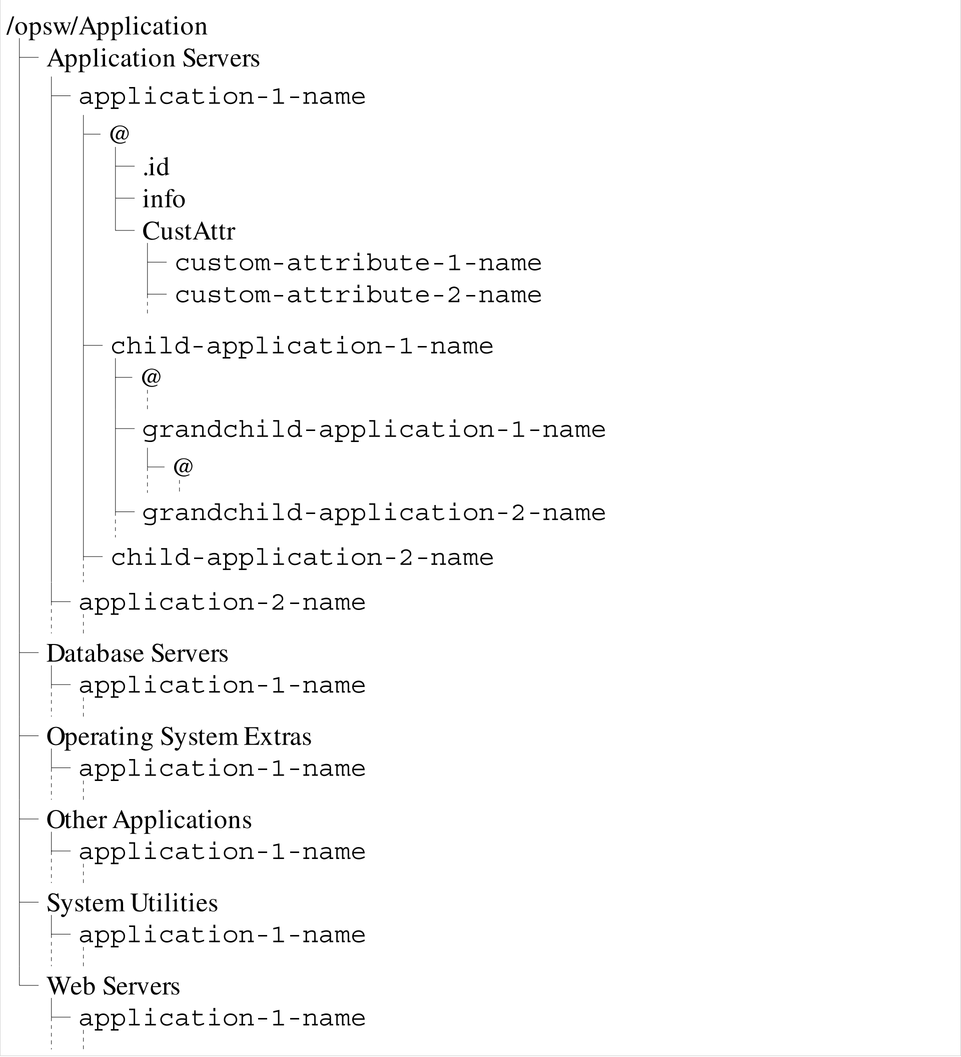

Customer Directory

The `Customer` directory contains a list of customers and customer details. Figure 10-8 illustrates the structure of the `Customer` directory.

Figure 10-8: Customer Directory

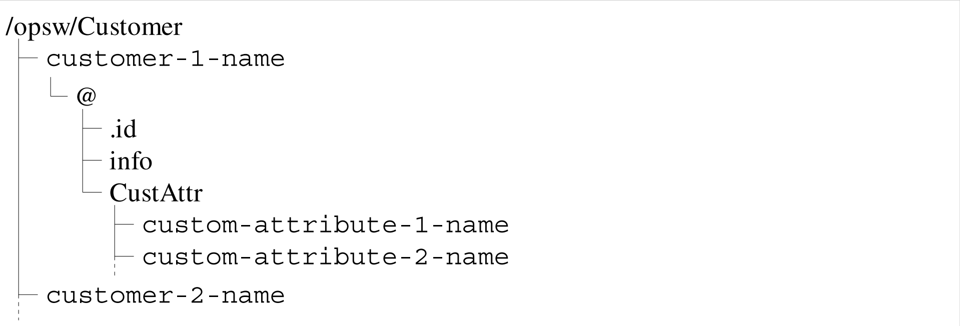

```
/opsw/Customer
 ├─ customer-1-name
 │   └─ @
 │       ├─ .id
 │       ├─ info
 │       └─ CustAttr
 │           ├─ custom-attribute-1-name
 │           ├─ custom-attribute-2-name
 │           ⋮
 ├─ customer-2-name
 ⋮
```

OS Directory

The `OS` directory contains a list of operating systems and operating system details. Figure 10-9 illustrates the structure of the `OS` directory.

Figure 10-9: OS Directory.

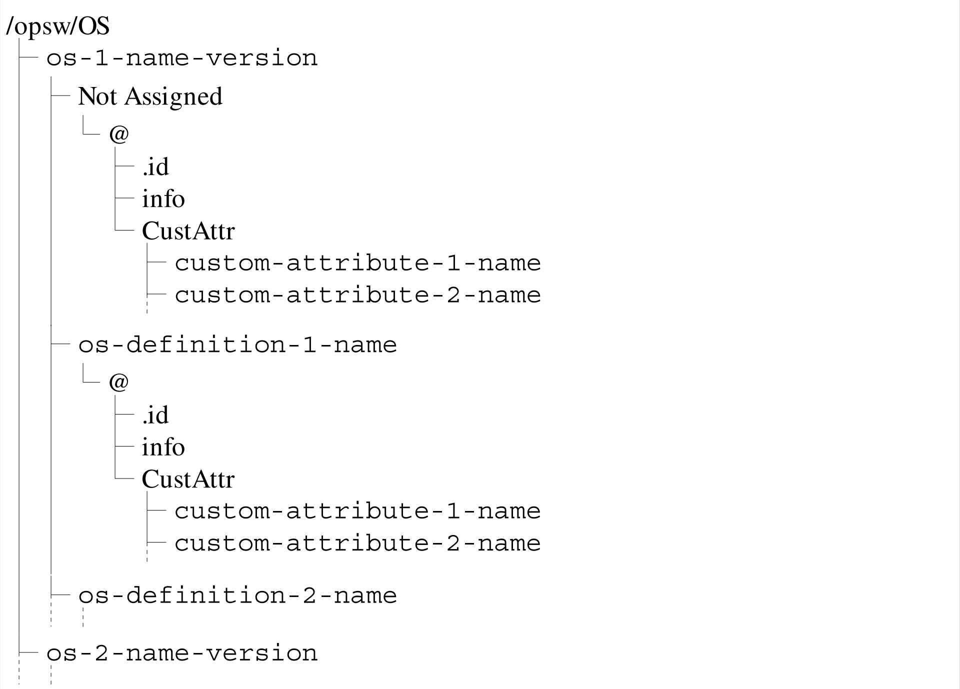

```
/opsw/OS
 ├─ os-1-name-version
 │   ├─ Not Assigned
 │   │   └─ @
 │   │       ├─ .id
 │   │       ├─ info
 │   │       └─ CustAttr
 │   │           ├─ custom-attribute-1-name
 │   │           ├─ custom-attribute-2-name
 │   │           ⋮
 │   ├─ os-definition-1-name
 │   │   └─ @
 │   │       ├─ .id
 │   │       ├─ info
 │   │       └─ CustAttr
 │   │           ├─ custom-attribute-1-name
 │   │           ├─ custom-attribute-2-name
 │   │           ⋮
 │   ├─ os-definition-2-name
 │   ⋮   ⋮
 ├─ os-2-name-version
 ⋮   ⋮
```

Realm Directory

The `Realm` directory contains a list of realms and realm details. Figure 10-10 illustrates the structure of the `Realm` directory.

Figure 10-10: Realm Directory

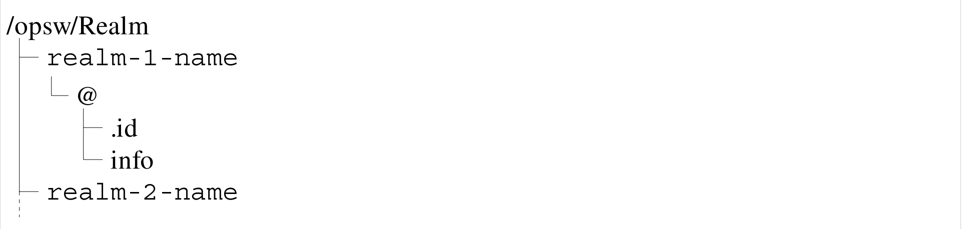

ServiceLevel Directory

The `ServiceLevel` directory contains service levels, service level details, service level children, and service level grandchildren. Figure 10-11 illustrates the structure of the `Service Level` directory.

Figure 10-11: ServiceLevel Directory

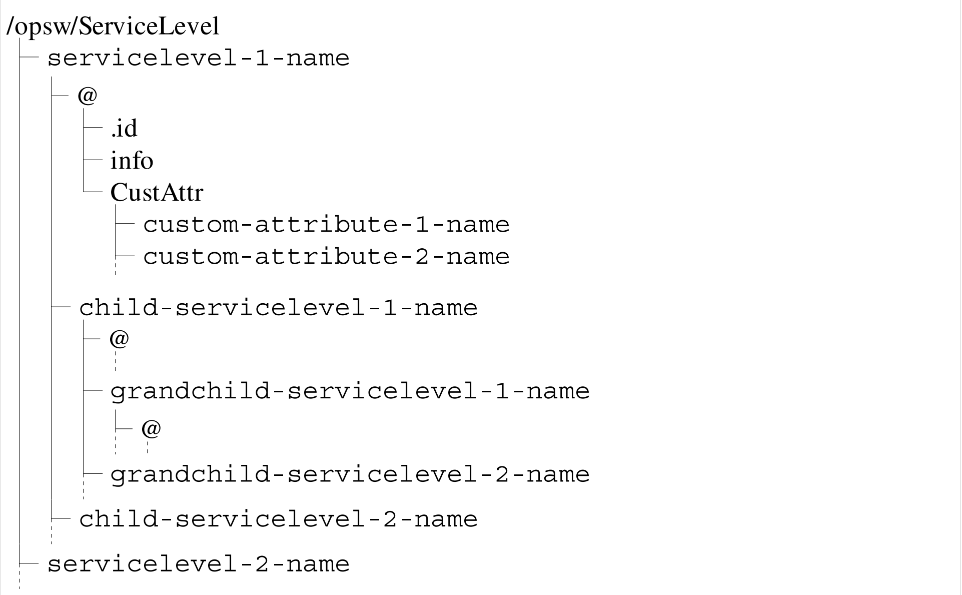

Server Filtering in the OGFS

As you navigate down the OGFS tree, the path grows longer and more specific—where fewer servers are visible in the `Server` directory. In the OGFS, the `/opsw` directory contains subdirectories for several types of objects in the Opsware model space, such as `Server`, `Group`, `Facility`, `OS`, `Application`, `Customer`, and so on.

In the Global Shell interface, you can filter your view of these object types in the `Server` directory by specifying an axis (`@`) in the path. A path in the Opsware model space can be a list of filtering criteria that selects objects of a given type. This path begins with the desired object type, such as `/Server`, and each filtering criteria begins with an `@`, such as `@Customer`. An ending `@` denotes the end of the filtering criteria.

Figure 10-12 is graphical representation of related objects (customers and facilities) in a hierarchical `Server` directory. The small boxes represent managed servers. Examples of ways that you can filter this directory immediately follow the diagram.

Figure 10-12: Filtering in the Server Directory

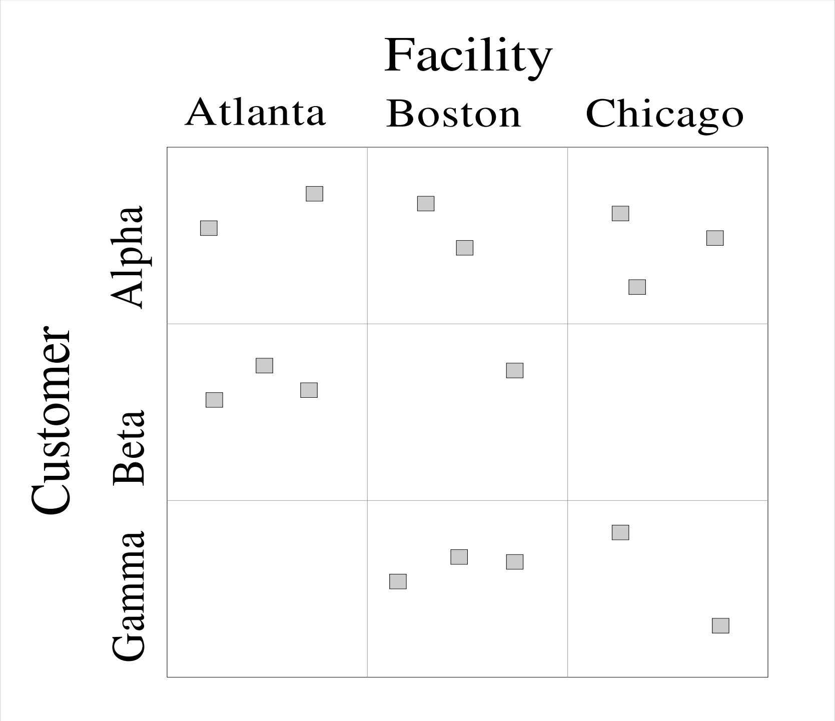

Based on Figure 10-12, the following examples illustrate ways to narrow your search for servers:

- To find all 16 servers, specify the following path:

```
$ ls /opsw/Server/@
```

- To find servers in the Atlanta facility, specify the following path:

```
$ ls /opsw/Server/@Facility/Atlanta/@
```

- To find servers that belong to customer Alpha, specify the following path:

```
$ ls /opsw/Server/@Customer/Alpha/@
```

- To find servers in the Atlanta facility that belong to customer Alpha, specify either of the following paths:

```
$ ls /opsw/Server/@Facility/Atlanta/@Customer/Alpha/@
$ ls /opsw/Server/@Customer/Alpha/@Facility/Atlanta/@
```

The following paths are filtered away by the OGFS, because they would yield a dead-end. There are no servers belonging to customer Gamma in the Atlanta facility.

```
$ ls /opsw/Server/@Facility/Atlanta/@Customer/Gamma/@
$ ls /opsw/Server/@Customer/Gamma/@Facility/Atlanta/@
```

This same filtering logic can be applied to `@Realm`, `@Group`, and `@Application`.

Remote Opsware Shell (rosh)

The Remote Opsware Shell (`rosh`) command makes a client connection that enables you to remotely run programs on managed servers. You invoke the rosh command from within a Global Shell session. The `rosh` command has the following syntax:

```
rosh (-n server-name | -i server-id)[-d dir] [-l login-name]
[-s] [-t | -T] [command [arg ...]]
```

Table 10-2 describes the `rosh` options.

Table 10-2: rosh Options and Commands

OPTION	DESCRIPTION
`-d` *`dir`*	Sets the working directory (path) on the remote server. The default is the remote user's home directory.
`-i` *`server-id`*	Specifies the server by its ID, which must already exist in the `/opsw/.Server.ID` directory.

Table 10-2: rosh Options and Commands (continued)

OPTION	DESCRIPTION
`-l` *`login-name`*	Specifies the login name of the remote user who performs operations on a remote server, which must already exist in the `/opsw/Server` directory.
`-n` *`server-name`*	Specifies the server by its name, which must already exist in the `/opsw/Server` directory.
`-r`	Relays RDP data to a managed server (on Windows).
`-s` *`script-name`*	Treats a command as the name of a saved script that will be sent to and run on the remote server.
`-t`	Forces the remote session to run in a pseudo terminal (for Unix servers only).
`-T`	Forces the remote session to run without a pseudo terminal (for Unix servers only).
`command [args . . .]`	Runs a program or saved script.

The following usage rules apply to the `rosh` program:

- Specify either the `-n` or `-i` option to log into or run programs on a managed server. These options are mutually exclusive, but if both are specified the `-i` option has precedence.
- If neither the `-n`, `-i,` and `id` options are specified, the managed server can be inferred if your working directory is at or below:

  `/opsw/Server/.../`*`server-name`*`/`

  Or

  `/opsw/.Server.ID/`*`server-id`*`/`
- If `-r` is specified, no other option (excluding `-n` or `-i`) can be specified.
- If `-l` is not specified, the login-name can be inferred if your working directory is at or below:

  `/opsw/Server/.../`*`server-name`*`/files/`*`login-name`*`/`

  Or

/opsw/.Server.ID/*server-id*/files/*login-name*/

- If `-s` is specified and *`command`* is a saved shared script with a `setuid` policy, the `login-name` specified by the `-l` option will be overridden. In this case, the `-l` option may be omitted. These scripts are stored in `/opsw/Script/Shared`.
- If your working directory is not below *`server`*`/files/`*`login-name`* and `-d` is not specified, the `cwdpath` defaults to the home directory for l*`ogin-name`*. To default to the home directory, you must specify `-l`.

rosh Operations

The `rosh` command establishes a client connection which enables you to remotely run programs on managed servers. The Opsware Global Shell feature provides the following modes of operation for `rosh`:

- **jump**: This operation starts a shell session in a pseudo-terminal on a managed server. This mode operates when you do not use the `-s` option and when you do not specify a command or a script. You must have the `loginToServer` permission on the managed server to jump.
- **reach**: This is a remote execution of commands that are native to the operating system of the managed server. This mode operates when you specify a command. You must have the `runCommandOnServer` permission on the managed server to reach.
- **push**: This is a remote execution of a script on a managed server. The script is stored in the OGFS and is sent to the managed server by `rosh`. Depending on the type of script, the following permissions may be required:
  - For shared saved scripts, which are stored in `/opsw/Script/Shared`, you can require the `runTrustedOnServer` permission on the managed server or other permissions that are specified by the policy of the script.
  - For any other script file, you must have the `runCommandOnServer` permission on the managed server.

The following examples illustrate what these operations look like for an Opsware user named `psi` at the path:

```
/opsw/Server/@/salish.snv1.corp.opsware.com/files/root/etc
```

```
[psi@m168 etc](538) $ uname -n; id; pwd
m168.dev.opsware.com
uid=59796(psi) gid=59796(psi) groups=59796(psi)
/opsw/Server/@/salish.snv1.corp.opsware.com/files/root/etc
```

The rosh jump command would display the following information about the managed server:

```
[psi@m168 etc](539) $ rosh
[root@salish etc]# uname -n; id; pwd
salish.snv1.corp.opsware.com
uid=0(root) gid=0(root)
groups=0(root),1(bin),2(daemon),3(sys),4(adm),6(disk),10(wheel)
,12(mail),7(lp),4(adm),9(kmem),6(disk),5(tty),3(sys),2(daemon),
8(mem)
/etc
[root@salish etc]# logout
```

The rosh reach command displays the following information about the managed server:

```
[psi@m168 etc](541) $ rosh "uname -n; id; pwd"
salish.snv1.corp.opsware.com
uid=0(root) gid=0(root)
groups=0(root),1(bin),2(daemon),3(sys),4(adm),6(disk),10(wheel)
,12(mail),7(lp),4(adm),9(kmem),6(disk),5(tty),3(sys),2(daemon),
8(mem)
/etc
```

The rosh push command displays the following information about the managed server:

```
[psi@m168 etc](544) $ cat /tmp/who.sh
#!/bin/sh
uname -n
id
pwd

[psi@m168 etc](543) $ rosh -s /tmp/who.sh
salish.snv1.corp.opsware.com
uid=0(root) gid=0(root)
groups=0(root),1(bin),2(daemon),3(sys),4(adm),6(disk),10(wheel)
,12(mail),7(lp),4(adm),9(kmem),6(disk),5(tty),3(sys),2(daemon),
8(mem)
/etc
```

Character Encoding for OGFS

To support international environments, the OGFS can display information in different character encodings such as Shift-JIS (Japanese) and EUC-KR (Korean). You can control the encoding of Global Sessions in the following ways:

- To specify the encoding of your Global Shell sessions, in the Terminal and Shell Preferences of the OCC Client, select an item from the Encoding drop-down list.

- To change the encoding of an active Global Shell session, run the `swenc` command with the `-e` option.

If you change the encoding of an active session, you must also change the encoding of the terminal application for that session. This procedure varies according to the terminal application.

Terminal Application Configuration

The terminal application that is hosting a Global Shell or Remote Terminal session must be configured to use the same encoding expected by the session. If the encodings do not match, the data might be displayed incorrectly.

When the OCC Client launches the terminal application, it composes the command specified by the Terminal Client field in the Terminal and Shell Preferences. If the Terminal Client field includes the `%e` substitution variable, the OCC Client replaces `%e` with an encoding name. For Global Shell sessions, this encoding name is specified by the Encoding field of the Terminal and Shell Preferences. For Remote Terminal sessions, this encoding name is the value of the Encoding field in the Properties section of the Server Explorer. If the terminal application does not support the encoding that replaces the `%e` variable, or if the `%e` variable is not specified, you must change the encoding manually in the terminal application after it starts.

Data that Cannot Be Displayed

Data that cannot be displayed might be from a managed server (such as the contents of files) or it might be the name of an object in the Opsware model. If the session's encoding does not support the data to be displayed, the data often appears as question marks. (However, it might appear as other characters such as exclamation points.) The session attempts to display this data with the current encoding. Usually, this data cannot be accessed. To access this data, select a compatible encoding for the session.

Objects in the Opsware model, such as Facility and Server, appear as file names in the OGFS. If these file names contain characters that cannot be represented by the encoding of the session, they are displayed as question marks, appended by the Opsware ID of the object. In the following example, the IDs are 10002 and 11002:

```
New York
Paris
Montr?al~10002
??~11002
```

If the model object does not have an ID, such as a custom attribute, then the session attempts to display the name with the current encoding.

swenc Command (Switching Character Encoding)

The `swenc` command enables you to change the character encoding within a Global Shell session. The `swenc` command has the following syntax:

```
swenc [-e encoding] [-T {on | off}] [-E] [-x] [-c command]
```

Table 10-3 describes the `swenc` options.

Table 10-3: swenc Options

OPTION	DESCRIPTION
`-c` *`command`*	Executes *`command`* and exits, reverting the session encoding to its previous state.
`-E`	Lists the valid character encodings.
`-e` *`encoding`*	Changes the character encoding of the current session.
`-T {on | off}`	Turns on or off the transcoding of data from the Unix managed server. (Data from Windows servers does not need to be transcoded.) See "Transcoded Data in a Managed Server" on page 408 in this chapter for more information.
`-x`	Prevents the launching of a sub-shell.

The following usage rules apply to the `swenc` utility:

- If you specify no options, `swenc` displays the character encoding and transcoding mode of the current session.
- Unless you specify the `-x` option, `swenc` starts a new sub-shell, which uses the encoding specified with the `-e` option. To leave the sub-shell and revert to the previous encoding, enter `exit`.
- Changing the encoding with `swenc` affects all processes in the current session, including background processes. If you change the encoding while background processes are running, the background processes might encounter errors.

- The swenc command affects only the current Global Shell session. For example, if you run the rosh command after the swenc -e command, the rosh command does not inherit the encoding that you changed with the swenc command.
- The swenc command does not change the working directory of the session, unless the working directory contains path names that cannot be represented in the new encoding. In this case, the working directory is the user's home directory.
- If you change the character encoding, make sure that the encoding of the terminal application that hosts the Global Shell and Remote Terminal sessions is set properly. To view or change the terminal client, go to the Terminal and Shell Preferences of the OCC Client.

LANG and LC_CTYPE Environment Variables

Many Unix commands (such as ls) rely on the character encoding, which is determined by the LANG or LC_CTYPE environment variables. In a Global Session, if the encoding is changed with the swenc command, the system attempts to reset these variables.

Opsware SAS determines the new value of the LANG variable with the following process:

1. The value of LANG is generated by combining the language of the user's profile in the Opsware Command Center with the current session encoding. For example, if the language is English and the session encoding is UTF-8, LANG is set to en_ US.utf8.
2. The value determined by the preceding step is compared with the set of valid locales on the OGFS server (according to the output of locale -a). If the value is a valid locale, LANG is set to this value.
3. If the value is not a valid locale, the system attempts to find a valid locale that specifies the user's language without the encoding. For example, if the user's language is English and the session encoding is EUC-JP, and this combination does not form a valid locale, LANG is set to en_US. If no matching locale is found, LANG is left unspecified.

The new value of the LC_CTYPE variable is determined in the following order:

1. Opsware SAS attempts to find a valid locale that matches the session encoding, regardless of the language.
2. If a valid locale is found, LC_CTYPE is set to this locale. For example, if the session encoding is EUC-JP, the LC_CTYPE variable is set to ja_JP.eucjp.

3. If no matching locale is found, `LC_CTYPE` is left unspecified.
4. If Opsware SAS cannot set the `LANG` or `LC_CTYPE` variables with the preceding process, you should set them manually.

Transcoded Data in a Managed Server

Transcoding is the conversion of data from one character encoding to another. Opsware SAS automatically transcodes some of the data between Global Shell sessions and other sources of data. For example, the file names of managed servers are transcoded, but the contents of the files are not. To see which data is transcoded, see Table 10-4. To display the transcoding mode of the current Global Shell session, enter the `swenc` command with no options.

Table 10-4: Data Transcoded for Global Shell Sessions

DATA	TRANSCODING
Objects in the Opsware model space, such as `Facility`, `Customer`, and `Server`. These objects are stored in the Opsware Model Repository (database) in UTF-8.	Between UTF-8 and the session encoding.
File and directory names of managed servers.	Between the managed server encoding and the session encoding.
Meta-data of managed servers, such as user names and registry key names.	Between the managed server encoding and the session encoding.
File contents of managed servers.	None
Contents of Windows registries, services, COM objects, and IIS metabases.	None
`rosh`: Contents of saved scripts executed on managed servers (a `rosh` push operation).	None
`rosh`: Ad-hoc scripts executed on managed servers.	None
`rosh`: Command-line arguments to programs executed on managed servers.	None

Table 10-4: Data Transcoded for Global Shell Sessions

DATA	TRANSCODING
rosh: Data streams (such as stdin and stdout) of programs executed on managed servers.	None
rosh: Data streams of rosh jump operations.	None
OGFS home, tmp, and var/tmp directories.	None

Disabling the Transcoding of Managed Server Data

On Unix servers, file and directory names can contain characters in arbitrary encodings. When accessed through the OGFS, file and directory names are transcoded by Opsware SAS. If the encoding of the names does not match the default encoding of the managed server, the transcoded data might be unusable.

You can disable transcoding with the swenc command. To turn transcoding on or off, use the -T option:

```
swenc -T {on | off}
```

If transcoding is disabled, file and directory names are passed unmodified from the managed servers. Therefore, you must manually configure the encoding of the terminal application to display the names correctly.

Windows servers store their file system data internally in the UTF-16 encoding. Because the encoding is known, transcoding is performed correctly and does not need to be disabled. Therefore, the -T option of the swenc command has no effect on Windows servers.

Error Messages

The Global Shell feature provides the file system error messages that are described in Table 10-5

Table 10-5: Global Shell Errors

ERROR	DESCRIPTION	ACTION
Input/output error	Your session has exceeded the time-out limit or the Agent is not running.	Start a new session or check the status of the Agent.
Operation not permitted.	No password was found.	Verify that you have a valid password.
Permission denied.	You are not allowed to view a directory. This does not mean that the directory does not exist on a given server. See the *Opsware® SAS Configuration Guide* for more information.	Verify that you have `readFileSystem` permissions.
RFS Specific error	You do not have permissions on the managed server. For example, this error will occur if you are trying to perform an operation on a managed server and you do not belong to the Administrators group that has the required permissions assigned to it.	You must have a set of permissions to perform operations on managed servers. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

For more detailed error messages, you can use the `rosh` program to help troubleshoot the server that is problematic. See "Remote Opsware Shell (rosh)" on page 401 in this chapter for more information.

Chapter 11: Distributed Script Execution

IN THIS CHAPTER

This section discusses the following topics:

- Distributed Script Execution
- Ways to Initiate Script Operations
- Script Management Tasks
- Script Execution
- Script Execution Results
- Error Resolution for Script Execution
- Opsware SAS Custom Extensions
- Control Scripts and Intelligent Software Modules

Distributed Script Execution

This section provides information about the Distributed Script Execution feature within Opsware SAS and contains the following topics:

- Overview of Distributed Script Execution
- Scripts Types – My Scripts, Shared Scripts, and Ad-Hoc Scripts
- Script Execution Functionality
- Run Distributed Script Link

Overview of Distributed Script Execution

The Script Execution feature provides tools for automating the management and execution of server scripts. Previously, a user created a script and then manually executed the script on individual servers, one server after another. With the Script Execution feature, a user performs all script tasks from one location – the Opsware Command Center.

From the Opsware Command Center, you can create or upload a script, set it up to run simultaneously across multiple Unix or Windows servers, and monitor it as it executes on each server. After a script runs, job- and server-specific execution results are available for review. You can modify, delete, or rerun the script again at a later date. See Figure 11-1

Figure 11-1: Script Execution Process

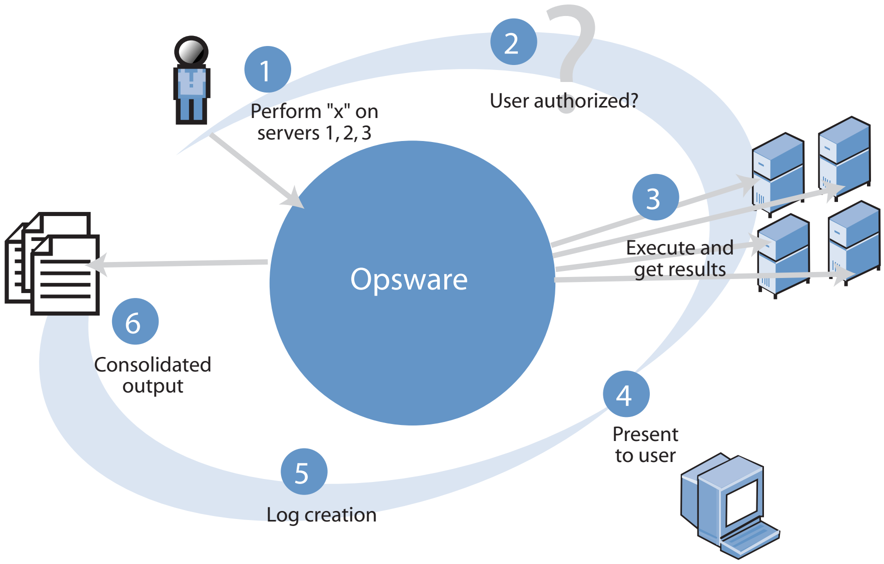

You need a specific set of feature permissions to manage and execute scripts. You'll also need permissions to access the servers associated with customers, facilities, or server groups. To obtain these permissions, contact your Opsware administrator. For more information, see the Permissions Reference Appendix in the *Opsware® SAS Configuration Guide*.

Scripts Types – My Scripts, Shared Scripts, and Ad-Hoc Scripts

The Script Execution Subsystem supports the three major types of scripts for Unix and Windows operating systems: Unix/Linux shell, Windows batch (.BAT), and Windows Visual Basic (VBScript).

After you create or upload Unix or Windows scripts in Opsware SAS, the scripts are stored in Opsware SAS in one of two ways:

- As private scripts, accessible only to the user who created them. In Opsware SAS, private scripts are called My Scripts.

  My Scripts can only be edited, deleted, or executed by the user who created the script. My Scripts are intended for personal use.

- As public scripts, accessible to all Script Execution users. In Opsware SAS, public scripts are called Shared Scripts.

A third type of Opsware SAS script is created (or uploaded) and then immediately executed by a user. The script is intended for one-time use and is not stored in Opsware SAS. In Opsware SAS, this type of script is referred to as an Ad-Hoc Script. During the Ad-Hoc Script creation and execution process, only one user has access to the script.

After you create a script and store it as a specific type of script in Opsware SAS, you cannot convert the script to the other type of script. My Scripts cannot be converted to Shared Scripts (and vice versa).

Script Execution Functionality

The Script Execution feature provides three basic functions:

- Script management
- Script execution
- Viewing script execution results

Script Management Tasks

Script Management tasks include:

- Viewing contents of stored My Scripts or Shared Scripts
- Creating (or uploading) My Scripts or Shared Scripts for storage in Opsware SAS
- Editing or deleting stored My Scripts or Shared Scripts
- Viewing version history of stored My Scripts or Shared Scripts

Script management functionality is handled by two Opsware SAS components – the Command Engine and the Opsware Command Center. The Command Engine handles the entry of scripts into the Opsware Model Repository (the script storage location in Opsware SAS) and the versioning of stored scripts. The Opsware Command Center provides the user interface for script management activities. It provides tools that allow users to create or upload scripts for storage. It also allows users to create Ad-Hoc (one-time-use, not stored) scripts for immediate execution.

Script Execution Tasks

Script execution tasks include:

- Executing a My Script or Shared Script, stored in Opsware SAS, on one or more servers
- Creating (or uploading) an Ad-Hoc Script and then immediately executing it on one or more servers

Script Execution Functionality

Script execution functionality is handled by three Opsware SAS components – the Opsware Command Center, the Command Engine, and the Opsware Agent. The Opsware Command Center provides the user interface for script execution activities. Script execution tasks are automated by a wizard that leads users through the following script execution steps:

1. Select scripts.
2. Select servers.
3. Specify execution options.
4. Confirm settings.
5. Execute scripts across one or more servers.

During script execution on the servers, the Command Engine runs a script that issues an execution command to the Opsware Agent on each server. Each Opsware Agent handles script execution and sends execution results to the Command Engine.

Execution of a script on the managed servers cannot be rolled back.

Script Execution Results Functionality

Execution results display immediately after a script runs and can be viewed any time after a script is executed. The functionality that displays execution results is handled by two Opsware SAS components – the Command Engine and the Opsware Command Center. The Command Engine enters the execution results data into the Model Repository. The

Opsware Command Center retrieves and displays execution results data from the Model Repository and provides tools for the user to download execution results data (output and error files) as a zip file.

Ways to Initiate Script Operations

In the Opsware Command Center, you can initiate a Script Execution operation in the following three ways:

- Scripts link
- Run Distributed Script link
- By selecting **Server ➤ Run Script** from these initiating scripts from the Scripts link and from the Run Distributed Script Wizard on the Home page.

Scripts Link

The Scripts link displays tools for managing scripts. Select this link to create or upload scripts that you want to save or run, and to view, edit, or delete scripts that have already been stored in Opsware SAS. The Scripts link is located on the navigation panel under Software. See Figure 11-2.

Figure 11-2: The Scripts Link Is Used to Create, Run, Upload, Edit, Delete, or View Scripts

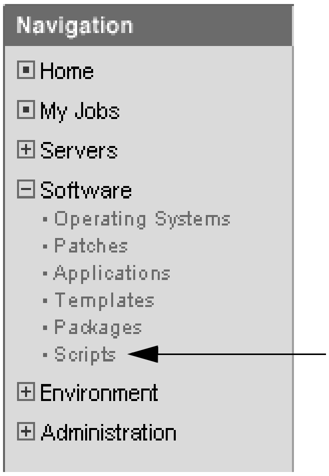

Run Distributed Script Link

The quickest way to execute a script that is stored in Opsware SAS, or to create and immediately execute an Ad-Hoc Script, is to select the Run Distributed Script link in the Tasks panel of the Opsware Command Center. See Figure 11-3.

Figure 11-3: Click the Run Distributed Scripts Link to Launch the Script Execution Wizard

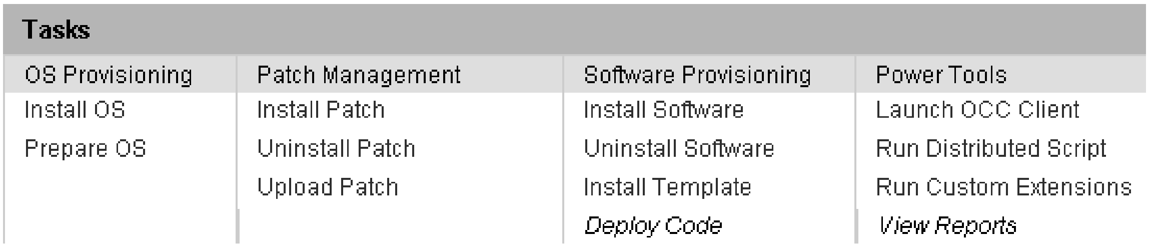

Tasks

OS Provisioning	Patch Management	Software Provisioning	Power Tools
Install OS	Install Patch	Install Software	Launch OCC Client
Prepare OS	Uninstall Patch	Uninstall Software	Run Distributed Script
	Upload Patch	Install Template	Run Custom Extensions
		Deploy Code	*View Reports*

You can view the status of scripts in the My Jobs area of the Opsware Command Center Home page, which displays information about jobs that have run, are currently running, or are scheduled to run, including script execution jobs. Through My Jobs, you can display script execution results. The name of the job is also a hyperlink to a pop-up window that allows you to change scheduling information for that job. In addition to the name, start time, and status, the number of servers affected by the job also appears.

The hyperlink to script execution jobs is called Run Script. A particular execution event can be identified by the start time of the execution event. Click any Run Script name to view the results of that particular script execution if it has completed. Otherwise, information about when the job is scheduled to run appears instead.

You can display a complete list of My Jobs by clicking the See All link in the My Jobs area of the home page, or by clicking the My Jobs link from the navigation panel.

Script Management Tasks

Supported script management tasks include:

- Script Creation Guidelines
- Viewing the Scripts List
- Editing, Deleting, and Downloading a Script
- Script Version History

Script Creation Guidelines

Opsware SAS supports the three major types of scripts for Unix and Windows operating systems: Unix/Linux shell, Windows batch (.BAT), and Windows Visual Basic (VBScript).

When creating scripts, adhere to the following guidelines:

- 4 MB is the maximum size allowable for a script.
- When you create a Unix shell script with a language other than the Bourne (sh) shell, use the sh-bang (#!) format at the top of the script to specify the correct command interpreter. The command interpreter needs to be present on the Opsware-managed server.

  For example, if you are using Perl, the beginning of the script would contain the following line:

  ```
  #!/usr/bin/perl
  ```

  The following example shows a short Perl script (it displays “hello world”):

  ```
  #!/usr/bin/perl
  print “hello world\n”
  ```
- VBScripts are executed by the VBScript interpreter on the Windows server.
- To access command line parameters with Unix shell commands, use the following convention: `$1 $2...`
- To access command line parameters with Windows .BAT, use: `%1 %2...`
- Script lines do not need to be terminated in a specific way. But with Windows scripts, Opsware SAS converts all `\n` to `\r\n`. With Unix scripts, all `\r\n` are converted to `\n`.
- Scripts should be written to send error output to standard error (file descriptor #2).
- Scripts should use the standard convention of returning a zero code to indicate success. For other return codes, there is no standard code system to follow. Create unique non-zero return codes to handle each type of error.

Creating or Uploading a Script

You can create scripts in the Opsware Command Center or by uploading a script into Opsware SAS. To create or upload a My Script or a Shared Script, perform the following steps:

1. From the navigation panel, select Software ➤ Scripts.
2. The Scripts page has two tabs: My Scripts and Shared Scripts.

- To create or upload My Scripts, select the My Scripts tab, and then New Script. The New Script page appears.
- To create or upload a Shared Script, select the Shared Scripts tab, and then New Script. The New Script page appears, as Figure 11-4 shows.

Figure 11-4: Scripts: New Script Page

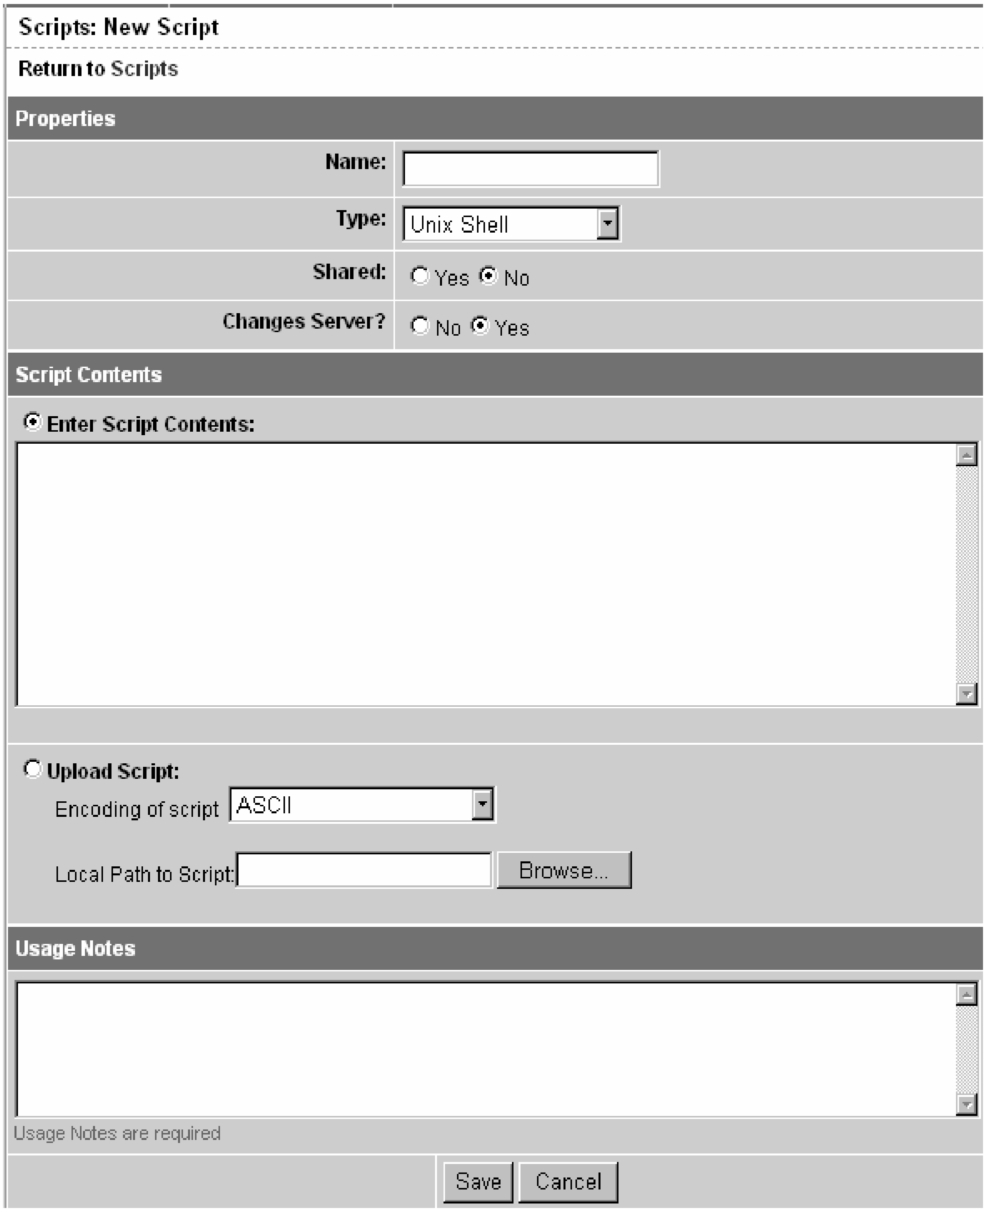

Scripts: New Script

Return to Scripts

Properties

Name:

Type: Unix Shell

Shared: ○ Yes ● No

Changes Server? ○ No ● Yes

Script Contents

● Enter Script Contents:

○ Upload Script:

Encoding of script ASCII

Local Path to Script:

Usage Notes

Usage Notes are required

3 On the Scripts: New Script page, enter the following data under Properties:

- Enter the name of the script. The name must be a unique shared-use or personal-use name.
- Select the script type: Unix shell, Windows .BAT, or Windows VBScript.
- If you are creating a Shared Script, select Yes next to Shared. If you are creating a My Script, select No.
- In the Changes Server field, select No if the script does not modify the server, and Yes if it does. If you select Yes, locked servers cannot be selected to run that script.

4 Under Script Contents, enter or upload a script by performing one of the following tasks:

- To upload a script, click **Upload Script**. In the Local Path to Script box, either manually enter the path to the script, or click **Browse** to locate the script.

When you upload a script, you must select the character encoding of the script's contents from the list. (Before storing the script, Opsware SAS converts the script contents from the encoding that you select to the UTF-8 encoding.)

- To create a script, click **Enter Script Contents** and manually enter the script in the text box.

The script editor does not recognize tabs and its functionality is browser-dependent.

5 In the Usage Notes section of the page, enter script details or other descriptive information.

6 Click **Save** to store the script. The script is saved in the Model Repository.

The Scripts page appears and confirms that the script is now stored. The script is included in the list of available Shared Scripts or My Scripts.

Viewing the Scripts List

After you save a script in Opsware SAS, you can view it on the list of stored My Scripts or Shared Scripts. My Scripts can be viewed only by the user who created them, while Shared Scripts can be viewed by all users.

To view the list of stored My Scripts or Shared Scripts, perform the following steps:

1. From the navigation panel, select Software ➤ Scripts.
2. The Scripts page has two tabs: My Scripts and Shared Scripts.
   - To view the list of My Scripts, select the My Scripts tab.
   - To view the list of Shared Scripts, select the Shared Scripts tab.

   A list of My Scripts or Shared Scripts displays. Each script name is also a link.
3. To view a script, click its name.

Editing, Deleting, and Downloading a Script

After you save a script in Opsware SAS, you can edit, delete, or download the script. Before you edit or delete a script, you might want to view the script properties, contents, usage notes, and change log.

My Scripts are accessible only to the user who created them and can only be edited or deleted by this user.

Editing a Script

To edit a stored script, perform the following steps:

1. From the navigation panel, select Software ➤ Scripts.
2. The Scripts page has two tabs: My Scripts and Shared Scripts.
   - To locate a My Script, select the My Scripts tab.
   - To locate a Shared Script, select the Shared Scripts tab.

   A list of scripts displays. Each script name is also a link.

3 To view the script, click the name. The View Script page appears, displaying script properties, contents, and usage notes, as Figure 11-5 shows.

Figure 11-5: The View Script Page

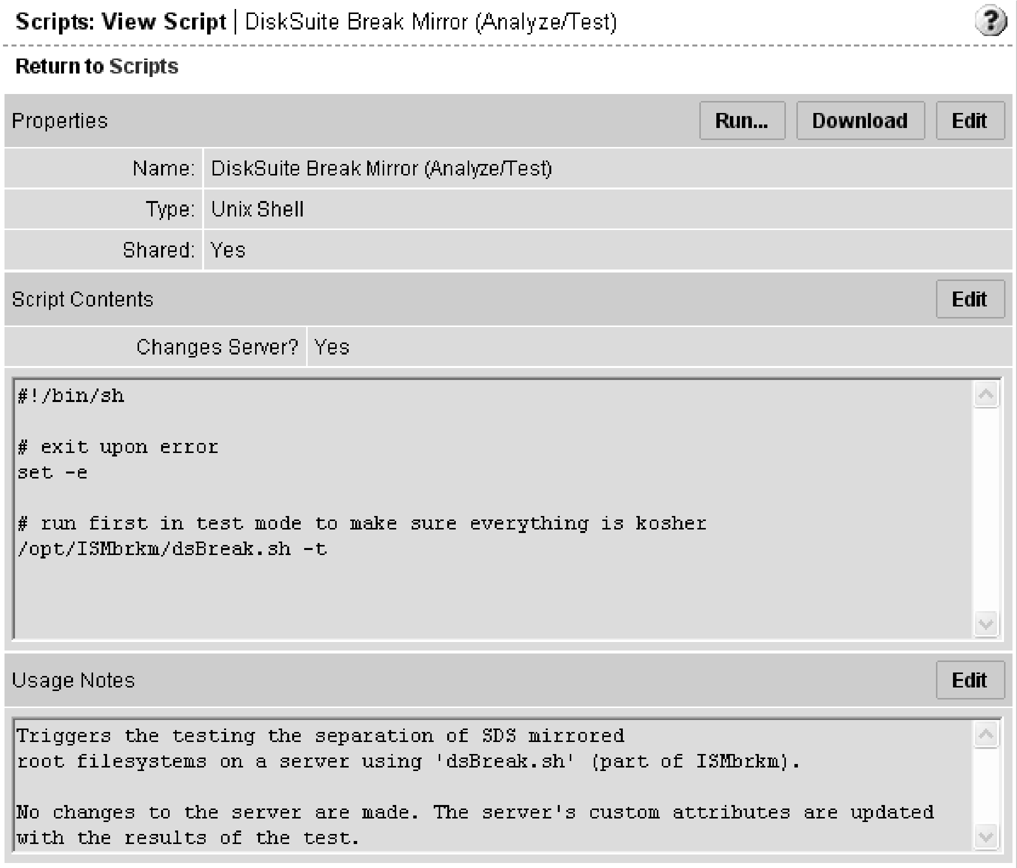

Scripts: View Script | DiskSuite Break Mirror (Analyze/Test)

Return to Scripts

Properties

Name:	DiskSuite Break Mirror (Analyze/Test)
Type:	Unix Shell
Shared:	Yes

Script Contents

Changes Server? Yes

```
#!/bin/sh

# exit upon error
set -e

# run first in test mode to make sure everything is kosher
/opt/ISMbrkm/dsBreak.sh -t
```

Usage Notes

```
Triggers the testing the separation of SDS mirrored
root filesystems on a server using 'dsBreak.sh' (part of ISMbrkm).

No changes to the server are made. The server's custom attributes are updated
with the results of the test.
```

4 Click **Edit** in the Properties panel for details about script properties.

5 On the Scripts: Edit Script page, Change Log information is now available. Select the Change Log tab to view the current script's version history.

6 Use the tabs on the Scripts: View Script page to edit script contents or the script name, and enter usage notes. See Figure 11-6.

Figure 11-6: View Scripts Page That Shows Properties, Contents, Usage Notes, and Change Log Tabs

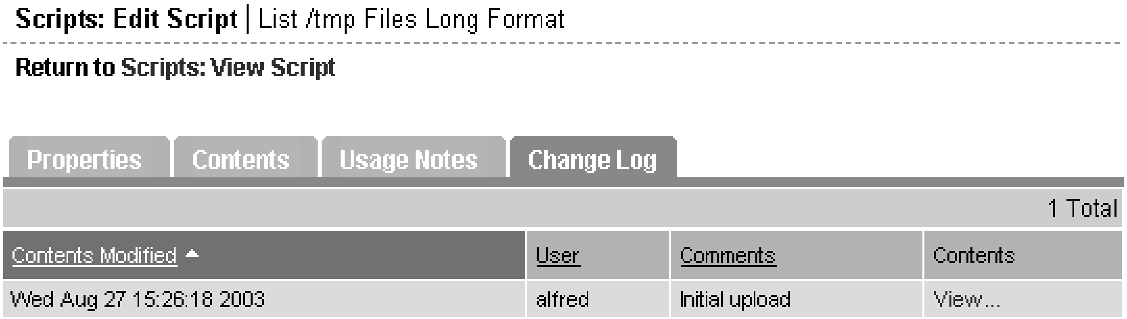

Scripts: Edit Script | List /tmp Files Long Format

Return to Scripts: View Script

Contents Modified ▲	User	Comments	Contents
Wed Aug 27 15:26:18 2003	alfred	Initial upload	View...

- To edit script contents, select the Contents tab. In the Edit Contents panel, select the Edit Script Contents radio button and edit the contents of the script.

  Instead of manually editing script contents, you can also upload new script contents. The uploaded script overwrites current script contents. To upload, select the Upload & Overwrite Script radio button and enter the location of the script you want to upload.

  After you edit script contents or upload new content, enter change log comments in the text box below the Edit Contents panel. Change log comments are required when you edit a script. When you are finished, click **Save**.

- To edit the name of the script, select the Properties tab. Edit the name that currently displays in the Name box and click **Save**.

- To enter usage notes, select the Usage Notes tab and enter information on the Edit Usage Notes panel. When finished, click **Save**.

Deleting a Script

To delete a stored script, perform the following steps:

1 From the navigation panel, select Software ➤ Scripts.

2 The Scripts page has two tabs: My Scripts and Shared Scripts.

- To locate a My Script, select the My Scripts tab.

- To locate a Shared Script, select the Shared Scripts tab.

A list of scripts displays. Each script name is also a link.

3. To review a script before you delete it, click the name. To return to the list of scripts, click the Scripts link in Return to Scripts (located at the top of the Scripts: View Script page).

4. On the Scripts page, select the scripts that you want to delete by clicking the box located to the left of the script name. You can delete more than one script at a time.

5. Click **Delete**.

6. The Delete Scripts confirmation window appears.

   - To review the list of scripts you selected for deletion, click **View Details**.

   - To cancel the entire operation, click **Cancel**.

   - To delete the selected scripts, click **Delete**.

After you delete a script, you can still view the results of executions performed with that script.

Downloading a Script

To download a stored script, perform the following steps:

1. From the navigation panel, select Software ➤ Scripts.

2. The Scripts page has two tabs: My Scripts and Shared Scripts.

   - To locate a My Script, select the My Scripts tab.

   - To locate a Shared Script, select the Shared Scripts tab.

   A list of scripts displays. Each script name is also a link.

3. Click the name of the script to download.

4. On the Scripts page, click **Download**.

5. Select the character encoding of the target server from the list. Opsware SAS converts the script contents from the UTF-8 encoding to the encoding that you select. Internally, Opsware SAS stores the script in the UTF-8 encoding.

6. in the pop-up window for the encoding, click **Download**.

Script Version History

Version history for a script is maintained in a change log, which is stored with other My Script or Shared Script management information. Each time a script is modified, new script version information is created and stored.

Viewing *Script Version History*

To view script version history, perform the following steps:

1. From the navigation panel, select Software ➤ Scripts.
2. The Scripts page has two tabs: My Scripts and Shared Scripts.
   - To locate a My Script, select the My Scripts tab.
   - To locate a Shared Script, select the Shared Scripts tab.

   A list of scripts displays. Each script name is also a link.
3. To view the script, click the name. The Scripts: View Script page appears and displays script properties, content, and usage notes.
4. Click **Edit** in the Properties panel.
5. On the Scripts: Edit Script page, select the Change Log tab to access the change log and view version history for the current script.

   The change log provides the following script version information:
   - Date and time the script is modified
   - Users who modified the script
   - Comments associated with each script modification
   - Script contents for that modification

Script Execution

This section provides information about script execution tasks, tips, and procedures and contains the following topics:

- Overview of Distributed Script Execution
- Script Execution Guidelines
- Executing My Script or Shared Scripts
- Creating and Executing an Ad-Hoc Script

Overview of Script Execution

A Script Execution Wizard automates script set up and execution processes and steps the user through the following execution stages (in the order shown):

1. **Select Script**: You can select only one My Script or Shared Script for each execution run. You can select scripts in the Wizard or from the Scripts page by selecting Software ➤ Scripts in the navigation panel.

2. **Select Servers**: You can select one or more servers or groups from the displayed list of available servers. Only servers running an operating system applicable to the selected script are shown (for example, only servers running Unix are shown for a Unix shell script).

3. **Specify Options**: Runtime data, information, and execution options are entered at this stage.

4. **Confirm Settings**: This stage allows the user to review settings prior to execution. Click **Next**.

5. On the Schedule and Notify page, you have the following options:

   - **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.

   - **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

     When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

     - **Option 1**: Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.

     - **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

6 **Run Script**: While the script executes, progress information displays about the total run and each server. When execution ends, summary information displays. After execution, script contents, output data, and error data can be immediately reviewed and downloaded.

The default amount of output data that Opsware SAS stores is 10 KB. This amount can be modified by the Opsware administrator.

Script Execution Guidelines

When executing a script, adhere to the following guidelines:

- No specific assumptions should be made about the execution environment. No particular environment variables are set.
- When executing processes on Unix servers, make sure long running processes (Web servers, databases, and so forth) are started as daemons. Also, make sure server processes properly daemonize themselves.
- For Windows servers, do not start anything that causes a window to open and wait for input.

Executing My Script or Shared Scripts

To execute My Script or Shared Scripts, perform the following steps:

1 Click the Run Distributed Script link in the Tasks area of the Opsware Command Center home page to launch the Run Distributed Script Wizard.

2 On the Overview page, make sure that you select Saved Script.

3 Click **Start**.

Alternatively, you can initiate the script execution process by clicking Software ➤ Scripts on the navigation panel, and then selecting a script and clicking **Run**.

4 To list Shared Scripts, select the Shared Script tab. Or, to list My Scripts, select the My Scripts tab.

5 From the script list, locate the script and select the script's radio button located to the left of the script name.

6 Click **Next** to continue (or click **Previous** to return to the previous step).

7 At the Select Servers page, browse or search for a server or servers to use.

- To browse, select the Browse tab to obtain a list of servers available to you. At the top of the page, use the Customers and Server Status to narrow your selection.
- To search, select the Search tab. On the Search page, indicate search criteria by selecting the appropriate check box.

Only servers are listed that use the operating system appropriate for the type of script that you want to run. For information about a server, click the name of the server. To sort the server list by server name, IP address, OS version, facility, or customer, click the heading of the column that you want to sort. For example, to sort by customer, click the heading titled Customer.

8 From the server list, specify the servers that you want to use. Select one or more servers or groups for script execution.

9 Click **Next** to continue (or click **Previous** to return to the previous step).

10 On the Specify Options page, enter the runtime information, as Figure 11-7 shows.

Figure 11-7: The Specify Options Page

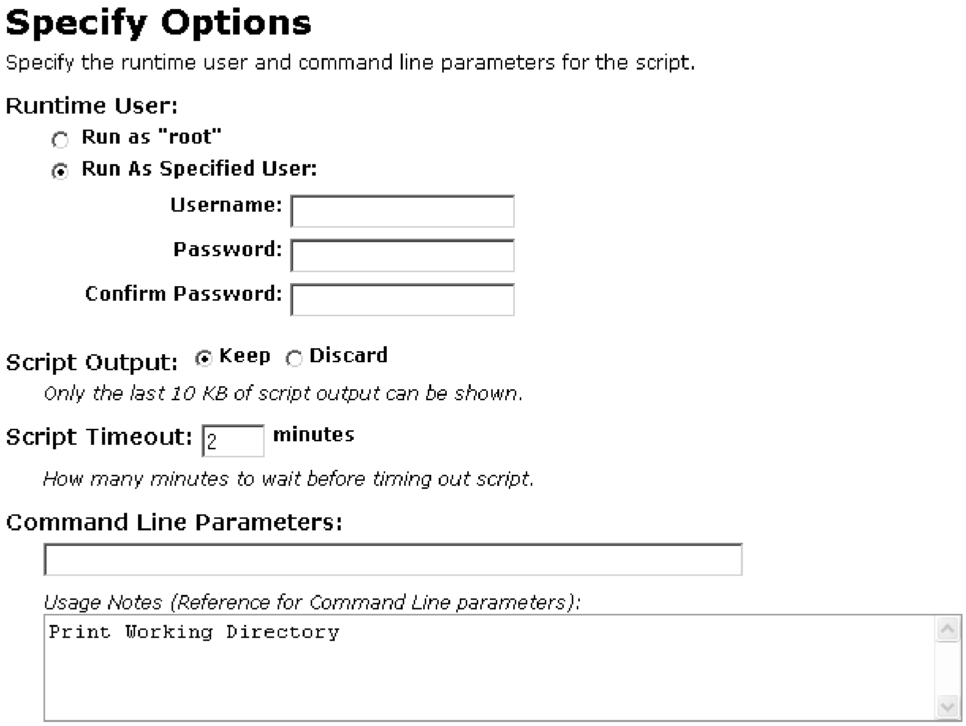

Specify Options

Specify the runtime user and command line parameters for the script.

Runtime User:

○ Run as "root"

● Run As Specified User:

Username:

Password:

Confirm Password:

Script Output: ● Keep ○ Discard

Only the last 10 KB of script output can be shown.

Script Timeout: 2 minutes

How many minutes to wait before timing out script.

Command Line Parameters:

Usage Notes (Reference for Command Line parameters):

```
Print Working Directory
```

- Specify the runtime user. If you have the appropriate permission, you are able to execute the script as root or local system without entering a password. Otherwise, enter a user name and password for the servers you intend to execute on. The user name and password you use must be the same across all servers.
- Specify if you want to keep or discard script output. Only a portion of script output is saved (the amount of script output that is saved is configurable by the Opsware administrator, but the default is 10 KB).
- Enter a script time-out value in minutes.

  This value is the amount of time a script has to complete execution activities on a server. If the script is not finished when the time-out value is reached, the script is stopped by Opsware SAS and a script error occurs.

  Select a time-out value that is greater than the time required for execution to complete.

- Enter command line parameters.

  Use the same syntax as when entering a script on a command line. For Unix scripts, use Bourne (sh) shell syntax. For Windows scripts, use `cmd.exe` parameter syntax.

- Add usage notes.

11 Click **Next** to continue (or click **Previous** to return to the previous step).

12 On the Confirm Settings page, review your script execution settings before you proceed. The Script to be run panel provides information about your script and the Servers panel displays the servers selected for script execution. The Confirm Settings page displays execution settings, as Figure 11-8 shows.

Figure 11-8: The Confirm Settings Page

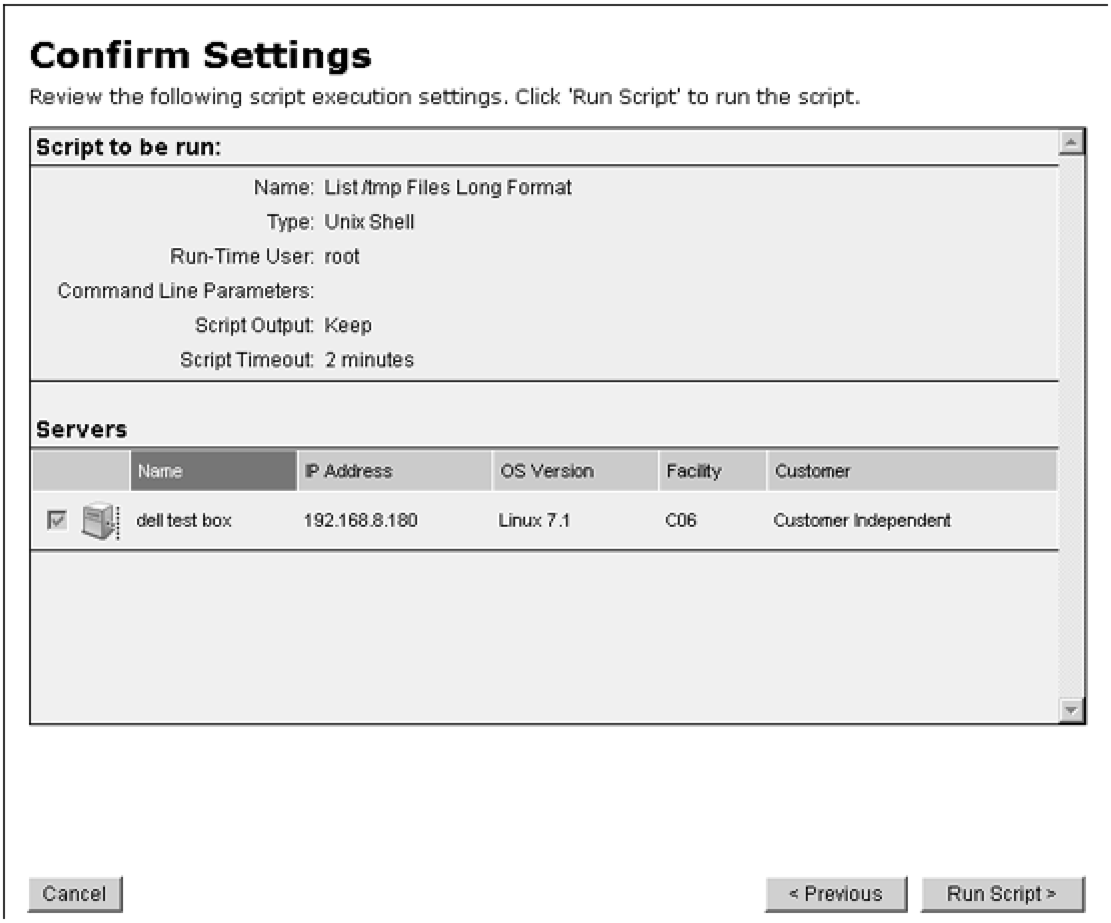

Confirm Settings

Review the following script execution settings. Click 'Run Script' to run the script.

Script to be run:

Name:	List /tmp Files Long Format
Type:	Unix Shell
Run-Time User:	root
Command Line Parameters:	
Script Output:	Keep
Script Timeout:	2 minutes

Servers

	Name	IP Address	OS Version	Facility	Customer
☑	dell test box	192.168.8.180	Linux 7.1	C06	Customer Independent

If you discover that script or server changes are needed, click **Previous** as many times as you need to return to the Script or Select Servers page.

13 On the Schedule and Notify page, you have the following options:

- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.
- **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

  When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have "Allow Run Refresh Jobs" permissions, you will see additional options. Select either of the following options:

  - **Option 1**: Membership is determined based on the "Time of Confirm Selection." Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
  - **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

14 Click **Run Script** to execute the script. While the script executes, progress information displays about the total run and each server. When execution ends, summary information displays. After execution, script contents, output data, and error data can be immediately reviewed and downloaded.

If you select to run the job at that time, a progress bar appears that shows the progress of the script execution.

15 After script execution, summary details on each server and the total job display. An ID number at the top of the page identifies the job.

- To see script output, summary details and information about errors that might have occurred, click **View Details**.
- For each server, you can click **Download** to obtain a zip file (`results.zip`) that contains script execution output and error data.

  The data is in two files called `stdout` (output data) and `stderr` (error data), in a directory named <servername>. For example, for a server named m0043, the output and error files would be located as follows: `m0043/stdout` (output file for server named m0043) and `m0043/stderr` (error file for server named m0043).

16 Click **Close** to exit the wizard.

Creating and Executing an Ad-Hoc Script

An Ad-Hoc script is written (or uploaded) and then executed without being stored in Opsware SAS.

Ad-Hoc script setup and execution are handled by the same wizard that steps the user through the processes used for stored scripts. However, the initial steps differ because Ad-Hoc script creation is integrated with execution activities.

Perform the following steps to create or upload an Ad-Hoc script, and then execute the script:

1 Click the Run Distributed Script link in the Tasks area of the Opsware Command Center home page to launch the Run Distributed Script Wizard.

2. On the Overview page, click Define Ad-Hoc Script ➤ Start. The Define Script page appears, as Figure 11-9 shows.

Figure 11-9: The Define Script Page

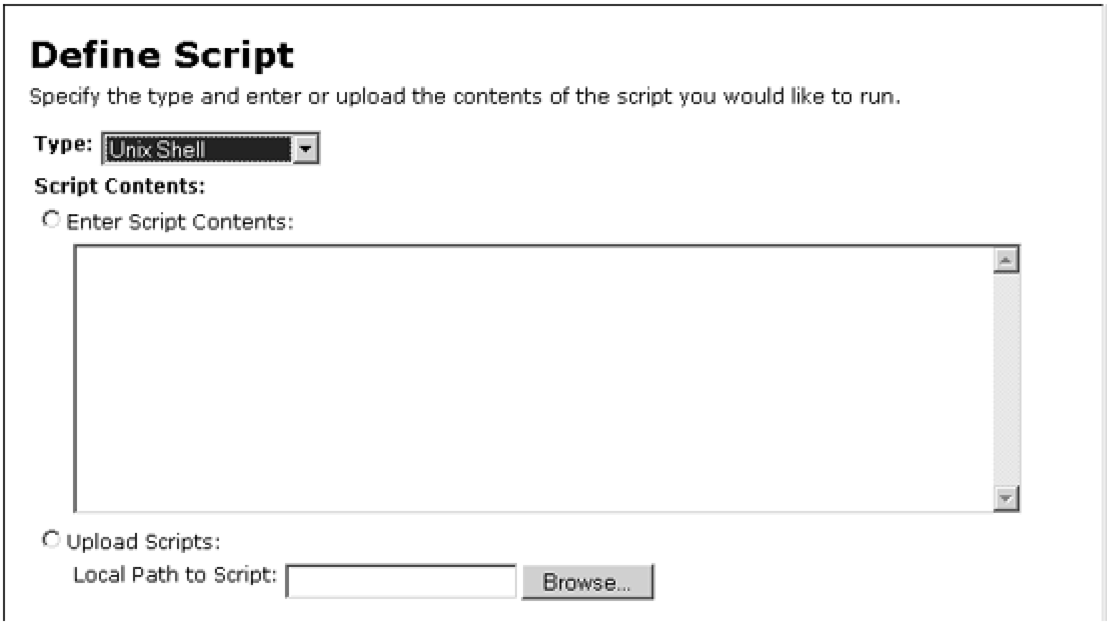

Define Script

Specify the type and enter or upload the contents of the script you would like to run.

Type: Unix Shell

Script Contents:

○ Enter Script Contents:

○ Upload Scripts:

Local Path to Script: Browse...

3. Select the type of script that you are creating (Unix shell, Windows VBScript, or Windows .BAT).

4. For Script Contents, indicate if you are entering script contents or uploading the script:

   - To upload a script, click **Upload Script**. In the Local Path to Script box, either manually enter the path to the script or click **Browse** to locate the path.

   - To create a script, click **Enter Script Contents** and enter the script in the text box.

The Scripts editor does not recognize tabs and its functionality is browser-dependent.

5. Click **Next** to continue (or click **Previous** to return to the previous step).

6. On the Select Servers page, browse or search for the server or servers to use.

- To browse, select the Browse tab to obtain a listing of servers available to you. At the top of the page, use the Customers and Server Status filters to narrow your selection.
- To search, select the Search tab. At the Search page, indicate search criteria by selecting the appropriate check box.

Only servers are listed that use the operating system appropriate for the type of script that you want to run. For information about a server, click the name of the server.

7 From the list of servers, select one or more servers or groups for script execution.

8 Click **Next** to continue (or click **Previous** to return to the previous step).

9 On the Specify Options page, enter the following runtime information:

- Specify the runtime user. If you have the appropriate permission, you can execute the script as root or local system without entering a password. Otherwise, enter an appropriate user name and password.
- Specify if you want to keep or discard script output. Only a portion of script output is saved (the amount of saved script is configured by the Opsware administrator).
- Enter a script time-out value in minutes.

  This is the amount of time a script has to complete execution activities on a server. If the script is not finished when the time-out value is reached, the script is stopped by Opsware SAS and a script error occurs.

  Select a time-out value that is greater than the time required for execution to complete.
- Enter command line parameters.

  Use the same syntax as when you enter a script on a command line. For Unix scripts, use Bourne (sh) shell syntax. For Windows scripts, use `cmd.exe` parameter syntax.

10 Click **Next** to continue (or click **Previous** to return to the previous step).

11 On the Confirm Settings page, review your script execution settings before you proceed to run the script. The Script to be run panel provides information about your script and the Servers panel displays the servers selected for script execution.

If you discover that script or server changes are needed, click **Previous** as many times as you need to return to the Script or Select Servers page.

12 On the Schedule and Notify page, you have the following options:

- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.
- **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

  When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have "Allow Run Refresh Jobs" permissions, you will see additional options. Select either of the following options:

  - **Option 1**: Membership is determined based on the "Time of Confirm Selection." Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
  - **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

13 Click **Run Script** to execute the script. While the script executes, progress information displays about the total run and each server. When execution ends, summary information displays. After execution, script contents, output data, and error data can be immediately reviewed and downloaded.

If you select to run the job at that time, a progress bar appears that shows the progress of the script execution.

14 After script execution, summary details display for the total job and each server. An ID number at the top of the page identifies the job.

- To see script output, summary details and information about errors that might have occurred, click **View Details**.

- For each server, you can click **Download** to obtain a zip file (`results.zip`) that contains script execution output and error data.

  The data is in two files called `stdout` (output data) and `stderr` (error data), in a directory named <servername>. For example, for a server named m0043, the output and error files would be located as follows: `m0043/stdout` (output file for server named m0043) and `m0043/stderr` (error file for server named m0043).

15 Click **Close** to exit the wizard.

Script Execution Results

This section provides information about script execution results and contains the following topics:

- Viewing Execution Results Immediately After Script Execution
- Viewing Execution Results Stored in Opsware SAS

During execution, the current status of script execution events at each server displays on the Run Script page. When script execution activities finish, the page immediately displays a summary of execution results for each server and for the entire run. See Figure 11-10.

Figure 11-10: Script Execution Progress Page

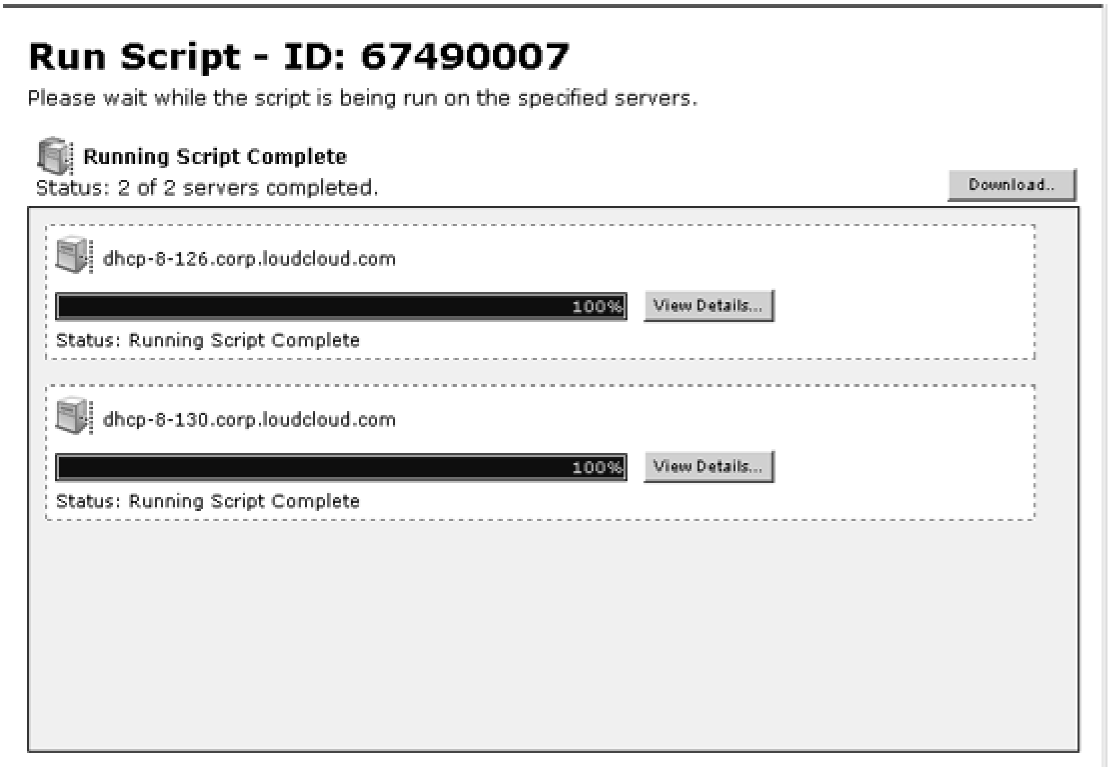

Run Script - ID: 67490007

Please wait while the script is being run on the specified servers.

Running Script Complete

Status: 2 of 2 servers completed.

dhcp-8-126.corp.loudcloud.com

100%

Status: Running Script Complete

dhcp-8-130.corp.loudcloud.com

100%

Status: Running Script Complete

Script execution data and information are stored in the Model Repository and are later accessible through the My Jobs feature for the user who performed the execution.

Viewing Execution Results Immediately After Script Execution

After a script runs, execution data and information for a specific server are available by clicking **View Details** for that server. Information that is displayed includes script execution output, errors (if there are any) and summary data (such as the script name and contents, and the date of the run). See Figure 11-11.

Figure 11-11: Execution Results After Script Execution

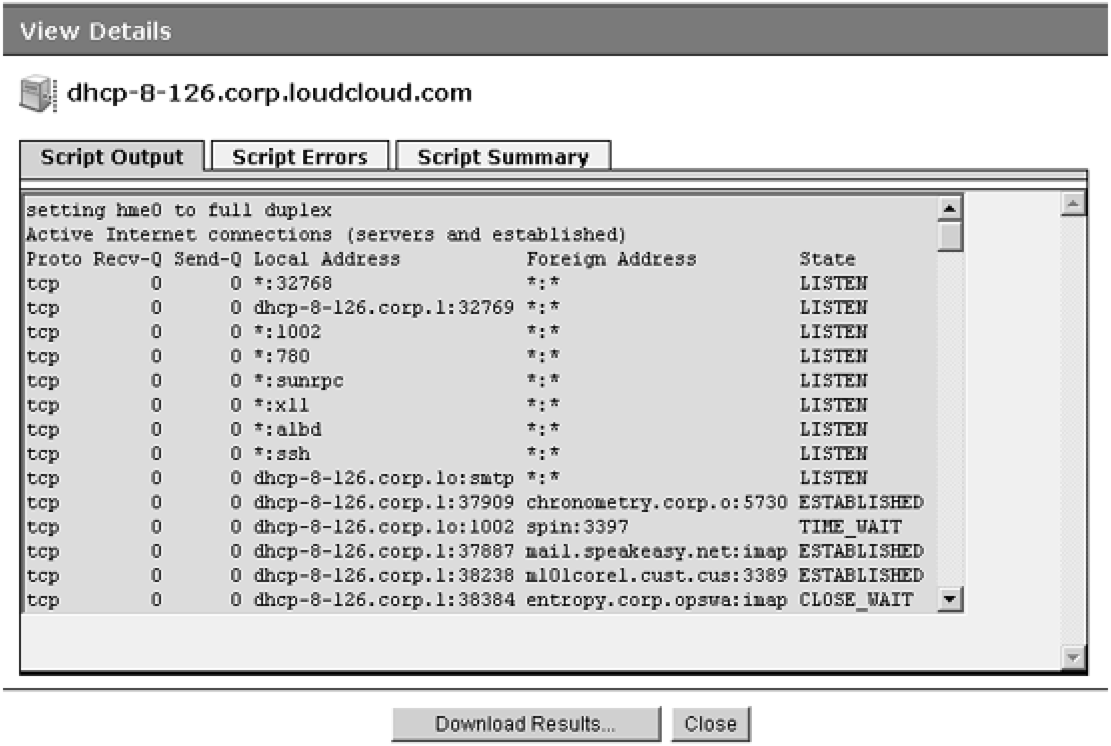

```
setting hme0 to full duplex
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address          Foreign Address        State
tcp        0      0 *:32768                *:*                    LISTEN
tcp        0      0 dhcp-8-126.corp.l:32769 *:*                   LISTEN
tcp        0      0 *:1002                 *:*                    LISTEN
tcp        0      0 *:780                  *:*                    LISTEN
tcp        0      0 *:sunrpc               *:*                    LISTEN
tcp        0      0 *:x11                  *:*                    LISTEN
tcp        0      0 *:albd                 *:*                    LISTEN
tcp        0      0 *:ssh                  *:*                    LISTEN
tcp        0      0 dhcp-8-126.corp.lo:smtp *:*                   LISTEN
tcp        0      0 dhcp-8-126.corp.l:37909 chronometry.corp.o:5730 ESTABLISHED
tcp        0      0 dhcp-8-126.corp.lo:1002 spin:3397             TIME_WAIT
tcp        0      0 dhcp-8-126.corp.l:37887 mail.speakeasy.net:imap ESTABLISHED
tcp        0      0 dhcp-8-126.corp.l:38238 ml01corel.cust.cus:3389 ESTABLISHED
tcp        0      0 dhcp-8-126.corp.l:38384 entropy.corp.opswa:imap CLOSE_WAIT
```

A zip file (`results.zip`) is available that contains standard script execution output and error data in two files (`stdout` and `stderr`), which are located inside a directory name <*`servername`*>. You can download the zip file by clicking **Download**.

Viewing Execution Results Stored in Opsware SAS

After script execution, script data results are available through the My Jobs feature.

The My Jobs list is available either at the My Jobs panel (Opsware Command Center home page) or the My Jobs page (opened by clicking the My Jobs link at Opsware Command Center's navigation panel). In the list, executed scripts are identified from other types of Opsware SAS jobs by the name Run Script.

You can view Information about a specific executed script (started on a particular day and time) by clicking the Run Script link for the specified day and time. See Figure 11-12.

Figure 11-12: Execution Results for a Particular Script Execution Event

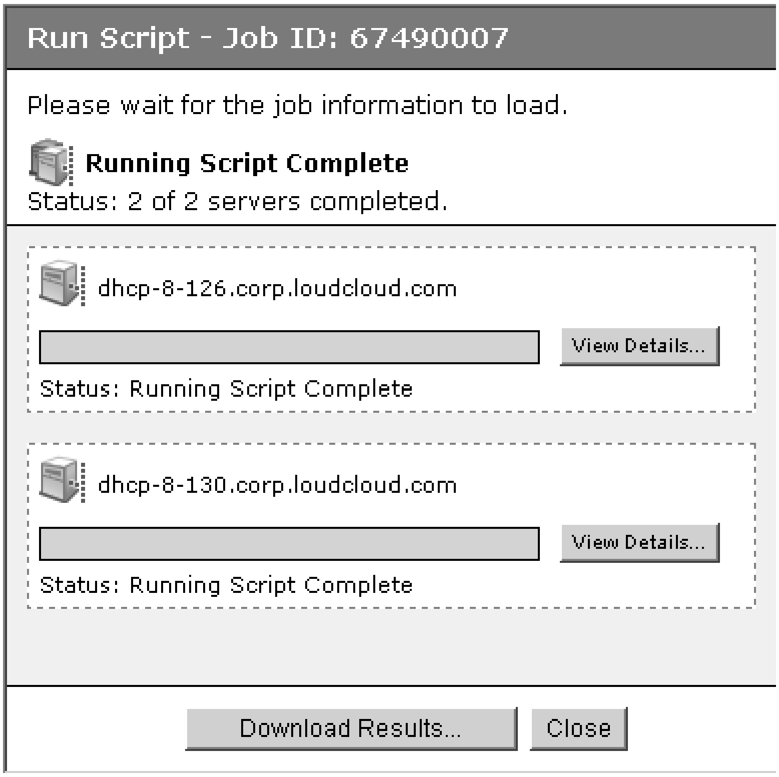

Run Script - Job ID: 67490007

Please wait for the job information to load.

Running Script Complete

Status: 2 of 2 servers completed.

dhcp-8-126.corp.loudcloud.com

Status: Running Script Complete

dhcp-8-130.corp.loudcloud.com

Status: Running Script Complete

Script results data are held in the Model Repository for a limited amount of time (30 days, for example). The Opsware administrator can configure this time period.

Error Resolution for Script Execution

The following section describes the events that might occur during Script Execution feature activities and provide suggestions for resolving the problems. This section contains the following topics:

- Resolving Script Uploading Errors
- Resolving Script Time-out Events
- Investigating Script Non-Zero Return Codes
- Investigating Server Authentication Errors
- Investigating Partial Executions

Resolving Script Uploading Errors

If errors occur when a user uploads a script, the Opsware Command Center error pages display information about the event.

If errors occur, perform the following steps:

- The maximum size for a script is 4 MB. Verify that the script being uploaded does not exceed this size.
- If communication with the Opsware Command Center is lost during the script upload process, upload the script again.

Resolving Script Time-out Events

When a script runs on an Opsware-managed server, if the script hangs or is not finished before the script time-out value is reached, an Opsware SAS error occurs and the script is stopped on the server.

The script time-out value is entered on the Specify Options page of the Run Distributed Script Wizard. The time-out value is the amount of time before a script times out.

Check that the length of time it takes for script execution does not exceed the current time-out value. Creation of debug output can also be added into the script for troubleshooting purposes. If the time-out event is due to the script hanging, then further examination of the script and server should occur.

Investigating Script Non-Zero Return Codes

If execution of a script on an Opsware-managed server returns a non-zero code, an error is reported. (A zero return code indicates successful – for example, normal – operation.) Information about the error is available to the user immediately following the error event and script execution. Information is also available after execution through the My Jobs feature.

Depending on how a script is written, a non-zero return code might indicate a fatal error.

Investigating Server Authentication Errors

If an authentication error occurs at an Opsware-managed server, verify that the user correctly enters the password. Also, verify that the correct user name and password are being used.

If authentication fails during execution of a script, an Opsware SAS error is raised for the server running the script. Information about the error is displayed in the Opsware Command Center for that server during and immediately following script execution. Script execution information is also available through the My Jobs feature.

Investigating Partial Executions

If a script runs on only some of the selected servers, this might be due to an Opsware Command Engine failure or because someone else has locked the managed servers by, for example, running another Opsware SAS task that uses the servers. Or, the servers might be unreachable.

Opsware SAS Custom Extensions

Opsware, Inc. Professional Services can extend functionality for customers by creating custom extensions to Opsware SAS. Opsware Custom Extensions (which are custom Command Engine scripts) extend Opsware SAS functionality to cover specific customer needs.

In Opsware SAS, the Command Engine is a system for running distributed programs across many servers (utilizing the Opsware Agents running on the servers). Opsware SAS features, such as the Code Deployment feature, use Command Engine scripts to implement part of their functionality

The Custom Extension feature is accessed through the Run Custom Extension Wizard. This Wizard allows a user to choose a custom extension to run, specify necessary input data for the extension, validate the data required to run the extension, run the extension, and view or download the results from the job. When a user runs a custom extension, the job shows up in My Jobs.

When a custom extension is added to one facility, it is automatically propagated to the other facilities in the multimaster mesh of Opsware SAS facilities.

To access the Run Custom Extension Wizard, users must be assigned to a user group that has the permission Wizard: Opsware Extension. When users have this permission, they can run any custom extensions on the servers they have access to in the Opsware Command Center.

Running a Custom Extension

To run a custom extension, perform the following steps:

1. From the Opsware Command Center home page, click the Run Custom Extensions link in the Tasks panel.

   Or

   From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers list appears. Select the servers on which you want to run a custom extension and choose **Run Extension** from the **Tasks** menu.

   The Run Custom Extension Wizard appears.

2. Select the custom extension that you want to run and click **Next**.

   If you have already selected servers from the Manage Servers list, you must ensure that the custom extension that you select can run on the operating systems of the selected servers. Otherwise, an error message appears in this page and you cannot proceed.

   Some custom extensions do not require that you select servers from the Server list. For example, the extension might prompt you in the Specify Settings step to enter server host names in a text box. If you already selected servers from the Servers list, an error message appears in the page.

   If you have not already selected servers from the Manage Servers list, the Select Servers page appears.

3. If prompted, select the servers or groups on which you want to run the custom extension and click **Next**. You can find the servers that you want to run a custom extension on by browsing the list or by searching.

The Specify Settings page appears. See Figure 11-13.

Figure 11-13: The Specify Settings Page of the Run Custom Extension Wizard

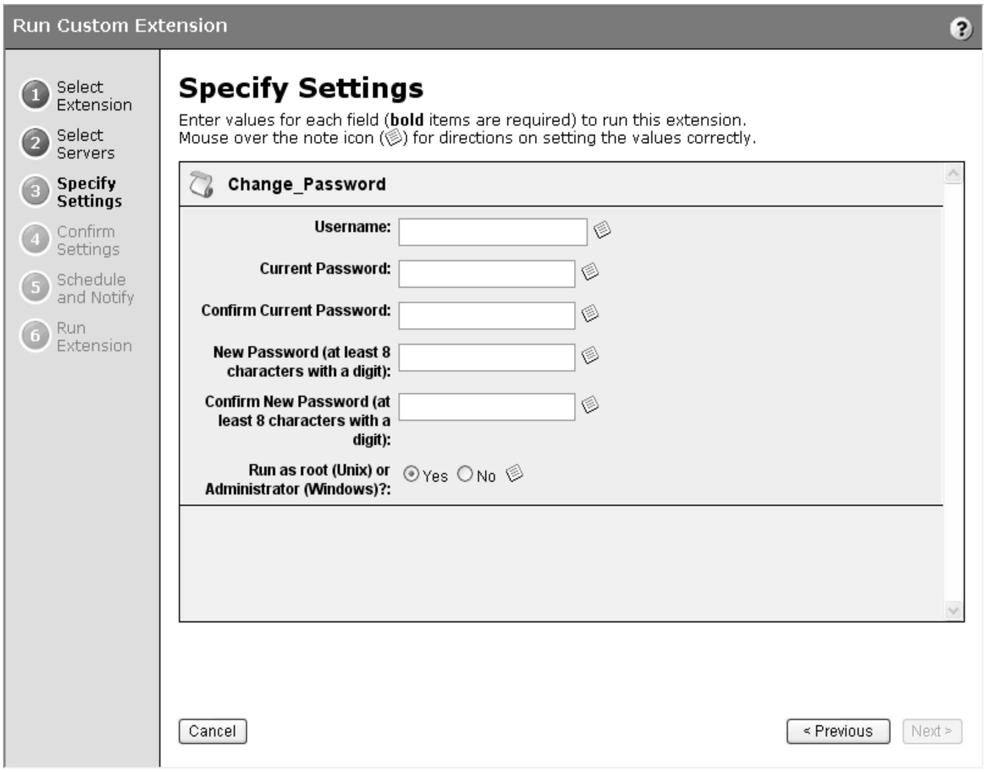

Specify Settings

Enter values for each field (**bold** items are required) to run this extension.
Mouse over the note icon (▨) for directions on setting the values correctly.

Change_Password

Username:

Current Password:

Confirm Current Password:

New Password (at least 8 characters with a digit):

Confirm New Password (at least 8 characters with a digit):

Run as root (Unix) or Administrator (Windows)?: ● Yes ○ No

4 Specify the settings for the custom extension and click **Next**.

The settings that appear in the page are unique to the custom extension that is being run. For information about what data to enter in a field, move your mouse pointer over a Note icon.

The Confirm Settings page appears. See Figure 11-14.

Figure 11-14: The Confirm Settings Page of the Run Custom Extension Wizard

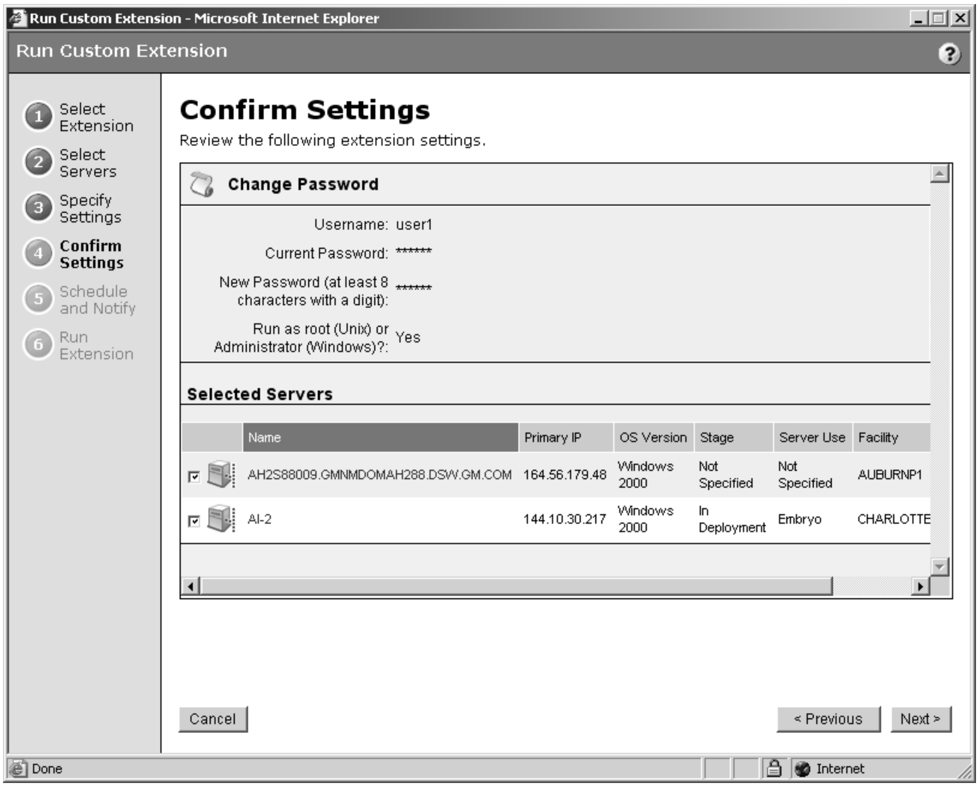

Confirm Settings

Review the following extension settings.

Change Password

Username: user1

Current Password: ******

New Password (at least 8 characters with a digit): ******

Run as root (Unix) or Administrator (Windows)?: Yes

Selected Servers

	Name	Primary IP	OS Version	Stage	Server Use	Facility
☑	AH2S88009.GMNMDOMAH288.DSW.GM.COM	164.56.179.48	Windows 2000	Not Specified	Not Specified	AUBURNP1
☑	AI-2	144.10.30.217	Windows 2000	In Deployment	Embryo	CHARLOTTE

5 Review the values that you entered and the servers that you selected on which to run the custom extension. (You can remove servers from the list by deselecting their check boxes. The list displays the first nine servers on which the custom extension will run. Click the "Show remaining servers" link to display the complete list of selected servers.)

6 Click **Next**.

7 On the Schedule and Notify page, you have the following options:

- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.

- **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

- **Option 1**: Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
- **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

8 Click **Run** to start the Custom Extension Wizard.

If you selected to run the job at that time, a progress bar appears for the servers on which the extension is running that shows the progress of the job. Depending upon how the custom extension was written, you will see the progress displayed in one of two ways:

- If the custom extension was written to show progress and status for individual servers, you will see individual progress bars for each server. When the custom extension has finished running, you will see the **View Details** button. You can click **View Details** to display detailed error information.
- If the custom extension was written to show progress for all servers, then you will see a single progress bars for all servers. When the custom extension has finished running for all servers, you will see **View Details**. You can click **View Details** to display detailed error information.

If an error occurs while the custom extension is running on servers, the progress bar moves to 100% and an error message appears below the bar.

For any servers where the custom extension has failed to run, you will have to either re-launch "Run Custom Extension" from the home page and pick the servers you want to try again or choose those servers from a Manage Server ➤ My Servers ➤ Server Search list and choose the menu item **Run Custom Extension**.

9 (Optional) When the custom extension finishes running, you can click **View Details** to see the results.

The Custom Extension Results window appears. The tabs in the window can vary depending on how the custom extension was implemented. To download the results to a file, click the download link. When you are done viewing the results, click **Close** to close the window.

10 Click **Close** to end the wizard.

Closing the wizard does not stop the custom extension if it is still running. After you close the wizard, you can view the progress of the running custom extension by viewing My Jobs (accessible from the Home page or the navigation panel). Each custom extension job is identified with the name Run Custom Extension. Click the name link to identify which extension was run.

See "Server Management Scheduling and Notification" on page 276 in Chapter 7 for information about the My Jobs feature.

Control Scripts and Intelligent Software Modules

An Intelligent Software Module (ISM) is an installable software package created with the Opsware ISM Development Kit (IDK). (For more information, see the *Opsware® SAS ISM Development Kit Guide*.) After the IDK uploads an ISM into Opsware SAS, the ISM appears in the Opsware Command Center as an application in the Software Tree. A wrench icon next to the application name identifies it as an ISM. Using the Opsware Command Center, you install the ISM onto managed servers.

An ISM can contain control scripts, which perform day-to-day, application-specific tasks such as starting software servers. For example, an ISM for Apache might contain control scripts that start and stop the HTTP server. You run control scripts with the Control window of the Opsware Command Center. (See Figure 11-15 on page 448.)

To run control scripts, you must belong to a user group with the ISM Controls permission. You'll also need permissions to access the servers associated with customers, facilities, or server groups. To obtain these permissions, contact your Opsware administrator. For more information, see the Permissions Reference Appendix in the *Opsware® SAS Configuration Guide*.

Figure 11-15: Control Window Showing Parameters, Schedule, and Notification

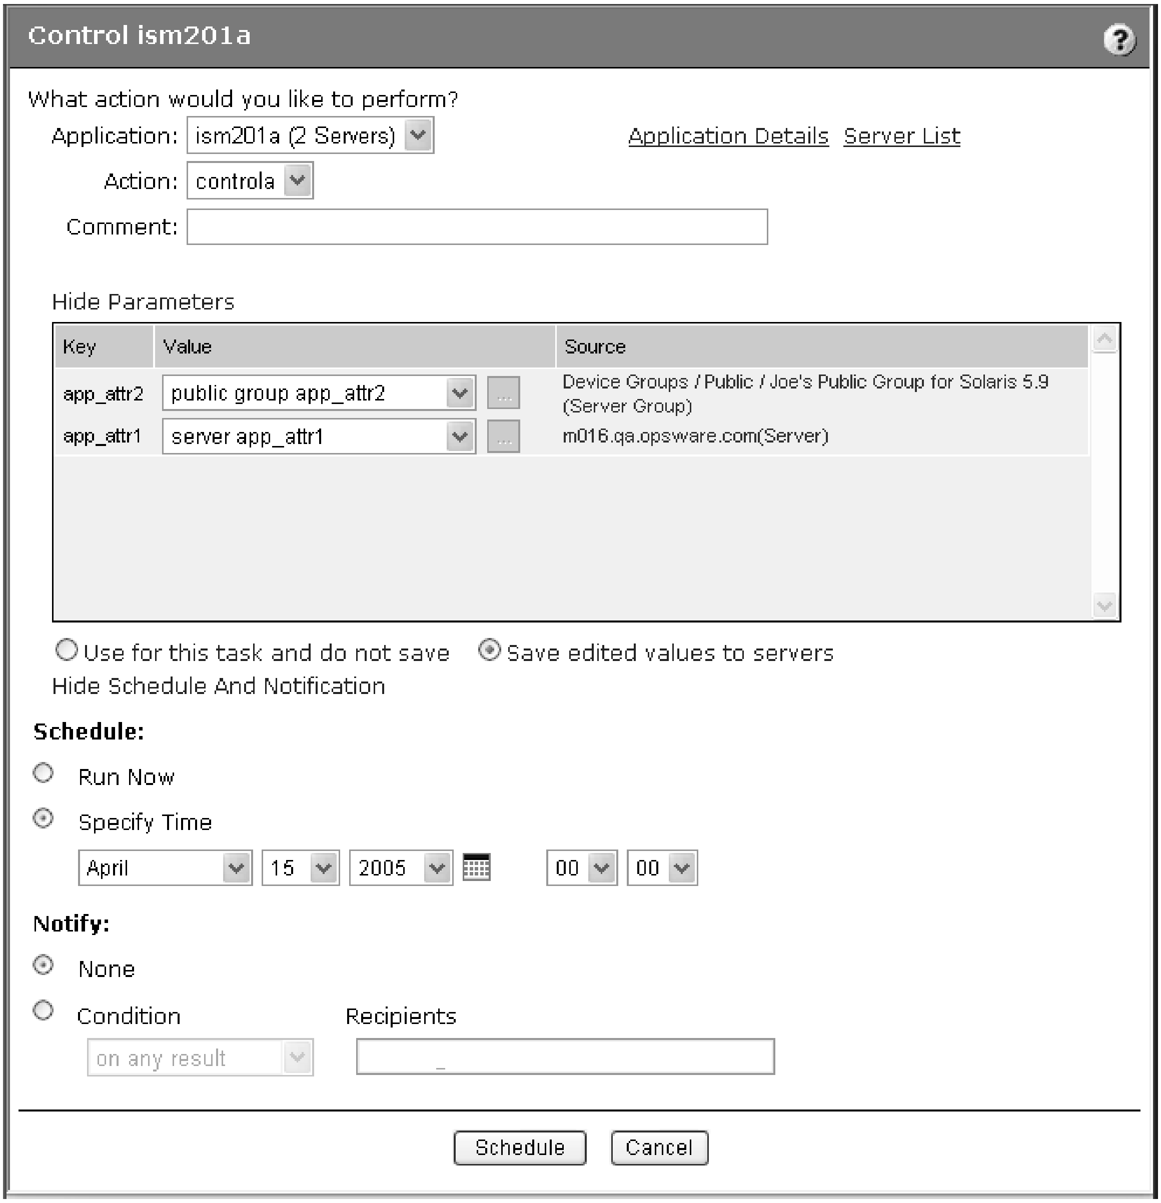

Control ism201a

What action would you like to perform?

Application: ism201a (2 Servers)

Action: controla

Comment:

Key	Value	Source
app_attr2	public group app_attr2	Device Groups / Public / Joe's Public Group for Solaris 5.9 (Server Group)
app_attr1	server app_attr1	m016.qa.opsware.com(Server)

○ Use for this task and do not save ● Save edited values to servers

Schedule:

○ Run Now

● Specify Time

April 15 2005 00 00

Notify:

● None

○ Condition on any result

Recipients

Parameters and Custom Attributes

A control script can have parameters corresponding to custom attributes. The key of a parameter matches the name of its correpsonding custom attribute. The value of a custom attribute determines the value of the parameter. The source of a custom attribute is an Opsware SAS object, such as a facility, customer, server, or server group.

Custom attributes with the same name (but with different values) can be specified on different Opsware SAS objects. If a server is associated with objects that have identically named custom attributes, Opsware SAS uses a predefined search order to determine the custom attribute that provides the parameter value. In the Control window of the Opsware Command Center, you can view the search order and the source of the custom attribute with the Value drop-down list. You can also override existing parameter values.

Control Scripts

An advanced user creates a control script with a text editor, packages the script into an ISM, and uploads the ISM into Opsware SAS. See the *Opsware® SAS ISM Development Kit Guide* for more information.

The control script name defined with the IDK might appear differently in the Control window. The Action field of the Control window displays the name of the control script, but without the leading `ism_` or the file type extension. For example, a control script named `ism_start.cmd` appears in Action field as `start`. The Action field displays only the first 25 characters of a control script name.

Opening the Control Window

The Control window of the OCC enables you to run or schedule a control script of an ISM. You can open the Control window when viewing either a server or an application.

To open the Control window from the My Servers (Summary View), Manage Servers (Summary View), or Search page, perform the following steps:

1. From the Navigation panel, expand the Servers link.

2. From the My Servers (Summary View), Manage Servers (Summary View), or Search page, locate the servers (or server group) with the ISM.

3. Select the checkboxes for the servers.

4. From the **Tasks** menu, select **Run ➤ Control**. The Control window appears. (If the **Control** menu item is disabled, then the server might not be reachable.)

To open the Control window from the My Servers (Software View) and Manage Servers (Software View) pages, perform the following steps:

1. From the Navigation panel, expand the Servers link.
2. From the My Servers (Summary View) or Manage Servers (Summary View) page, locate the servers (or server group) with the ISM.
3. From the **View** menu, select **Software**. The Software View page appears.
4. In the Software column, click the wrench icon of the ISM. The Control window appears.

To open the Control window from an application, perform the following steps:

1. From the Navigation panel, select Software ➤ Applications.
2. Locate the ISM.
3. Select the Members tab of the ISM.
4. Select one or more checkboxes for the servers.
5. From the **Tasks** menu, select **Run ➤ Control**. The Control window appears.

Running a Control Script on an ISM

To run a control script, perform the following steps:

1. Open the Control window. (See "Opening the Control Window" on page 449 in this chapter for more information.)
2. If the server you selected has more than one control script, select an entry from the Application (ISM) list.
3. Select an entry from the Action (control script) list.
4. In the Comment field, you can enter a brief description.
5. To view the Software Details window for this ISM, click the Application Details link. The Software Details window includes information such as the customer, full path (location) of the ISM in the tree, and the packages contained in the ISM.
6. To see the servers affected by running the control script, click the Server List link. The Server List window includes information such as the IP address, OS, customer, and facility.
7. To see the parameters of the control script, select the View Parameters link.

The Key, Value, and Source fields for each parameter appear in a table. The Key field identifies the name of a parameter. The Value field is the data assigned to the parameter. The Source identifies the object (such as a customer or server) with the custom attribute that corresponds to the Value field.

If custom attributes with the same name have been assigned at different levels, the Value field has a drop-down list. For a given execution of a control script, a parameter can have just one value. The drop-down list displays the Value data for the custom attributes in order of precedence (highest at the top). See "Parameters and Custom Attributes" on page 449 in this chapter for more information.

8 You can select a different entry from the Value drop-down list. Note that the Source field changes when you select a different entry.

9 If you want to override the data in the Value field, select "Set new value" from the drop-down list, click the ellipsis (...) button, and type the data in the pop-up window.

10 If you changed the Value field in either of the two preceeding steps, you can choose the duration of the change. For example, if you want the change only for the current execution of the script, select "Use for this task and do not save" (the default). If you want to save your change for subsequent executions, select "Save edited values to servers."

11 Click the View Schedule and Notification link.

12 If you want to run the script at this time, click **Run Now**. Then, **Run** appears at the bottom of the window.

13 If you want to run the script at a later time, select Specify Time. Then, **Schedule** appears at the bottom of the window. Choose the date and time.

14 If you want an e-mail notification sent when the script finishes running, select Condition and choose an entry from the drop-down list. In the Recipients field, enter the e-mail addresses separated by commas.

15 To run the script now, click **Run** at the bottom of the window.

16 To run the script later, click **Schedule** at the bottom of the window.

Viewing the Results of a Control Script

A control script runs as a job on one or more managed servers. To view the results of a job, perform the following steps:

1. From the Navigation panel, click the My Jobs link.
2. On the My Jobs page, enter the search criteria.
3. In the drop-down list of job types, select the Control item.
4. Click **Update**.

See "Viewing Execution Results Stored in Opsware SAS" on page 438 in this chapter for more information.

Chapter 12: Patch Management for Windows

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Patch Management for Windows
- Patch Management Roles for Windows
- Patch Management Process
- Patch Details
- Policy Management
- Patch Compliance
- Patch Administration for Windows
- Patch Installation
- Patch Uninstallation

Overview of Patch Management for Windows

The Patch Management for Windows feature enables you to identify, install, and remove Microsoft® Windows patches and maintain a high level of security across managed servers in your organization. With the OCC Client user interface, you can identify and install patches that protect against security vulnerabilities for the Windows 2000, Windows 2003, and Windows NT4.0 operating systems. These patches include Service Packs, Update Rollups, and hotfixes.

The OCC Client interface organizes Microsoft patches by operating systems and displays detailed vendor security information about each patch, such as Microsoft Security Bulletins. You can browse patches by the date Microsoft released the patch, by the severity level, by the Security Bulletin ID, QNumber, and so on. You can also browse all patches that are installed on a server, and view and edit patch metadata.

In Patch Management, you can separately schedule when you want patches imported from Microsoft (either automatically or on demand) into Opsware SAS and when you want these patches downloaded to managed servers. As a best practice, patch installations are typically scheduled for a time that causes minimal disruption to an organization's business operation. If you are installing one patch on one server, the installation operation will start only after the download operation has completed.

Patch Management also allows you to set up email notifications that alert you whether the download and installation operations completed, succeeded, or failed. When you schedule a patch installation, you can also specify your reboot preferences to adopt, override, postpone, or suppress the vendor's reboot options.

To provide flexibility in how you identify and distribute patches on managed servers or server groups, Patch Management allows you to create patch policies that define groups of patches you need to install. By creating a patch policy and attaching it to a server or server group, you can effectively manage which patches get installed where in your organization. In case you need to include or exclude a patch from a patch installation, Patch Management allows you to deviate from a patch policy by specifying that individual patch in a patch policy exception. An additional patch is one that is not already specified in the patch policy and is one that you want to include in (add to) the patch installation. A patch that you want to exclude from a patch installation is one that is already specified in a patch policy and is identified in the patch policy exception as one you do not want installed. In cases where it is already known that a certain Windows patch may cause a server or application to malfunction, you should create a patch policy exception to exclude it from being installed on that server or on all servers that have that application.

While Patch Management allows you to react quickly to newly discovered security vulnerabilities, it also provides support for strict testing and standardization of patch installation. After you have identified patches you need to install, Patch Management allows you to simulate (preview) the installation before you actually install a patch. This preview process tells you whether the servers you selected for the patch installation already have that patch installed. In some cases, a server could already have a patch installed if a system administrator had manually installed it. After this type of patch install, if a compliance scan has not been run or the installed patch has not been registered, Opsware SAS does not know about it. The preview process provides an up-to-date report of the patch state of servers. The preview process also reports on patch dependency and supersedence information, such as patches that require certain Windows products, and patches that obsolete other patches or are obsoleted by other patches.

Patch Management also provides a solution for remediating servers that are not operating properly due to installed patches. If installed patches cause problems, even after being tested and approved, Patch Management allows you to uninstall the patches in a safe and standardized way. Patch Management allows you to specify uninstall options that control server reboots and the execution of uninstall commands, and pre-uninstall and post-uninstall scripts. Similar to previewing a patch install, you can also preview a patch uninstall.

To help you track the patch state of servers or server groups, Patch Management allows you to export this information. This information can be exported in a comma-separated value (CSV) file and includes details about when a patch was last detected as being installed, when a patch was installed by Opsware SAS, the patch compliance level, what patch policy exceptions exist, and so on. You can then import this information into a spreadsheet or database to perform a variety of patch analysis tasks.

Features of Patch Management for Windows

The Patch Management for Windows feature is fully integrated with Opsware SAS. This feature leverages from the Opsware SAS server automation features. For example, Opsware SAS maintains a central database (the Model Repository) that includes detailed information about every server under management, the patches and software (applications) installed on managed servers, and the patches and software available for installation. You can use this data to determine the severity of a server’s exposure to a newly discovered vulnerability, and to evaluate the benefits of rolling out a patch versus the costs that might be incurred during downtime and testing activities.

By automating the patching procedure, the Patch Management can reduce the amount of downtime required for patching. Opsware SAS also allows you to separately schedule patch downloads and patch installations so that patching occurs during off-peak hours.

Opsware SAS automates patch management by providing the following features:

- A central repository where patches are stored and organized in their formats
- A database that includes information on every patch that has been applied
- Customized scripts that can be run before and after a patch is installed
- Advanced search abilities identify servers that require patching
- Auditing abilities so that security personnel can track the deployment of important patches

Software Library

The Software Library provides flexibility in searching for and displaying Microsoft patches by operating system, severity level, release date, bulletin ID, and so on. See Figure 12-1. The number in parenthesis is the total number of patches (for that operating system version) that were uploaded from the Microsoft web site. In the Content pane, a dimmed patch icon indicates that the patch has not yet been uploaded to the Software Library. Use the column selector to control which columns of patch metadata data to display, depending on what you find useful for any given patch.

Since the Software Library is integrated with Microsoft patch metadata, you can review vendor information (in real-time) in the Preview pane.

Figure 12-1: OCC Client Software Library

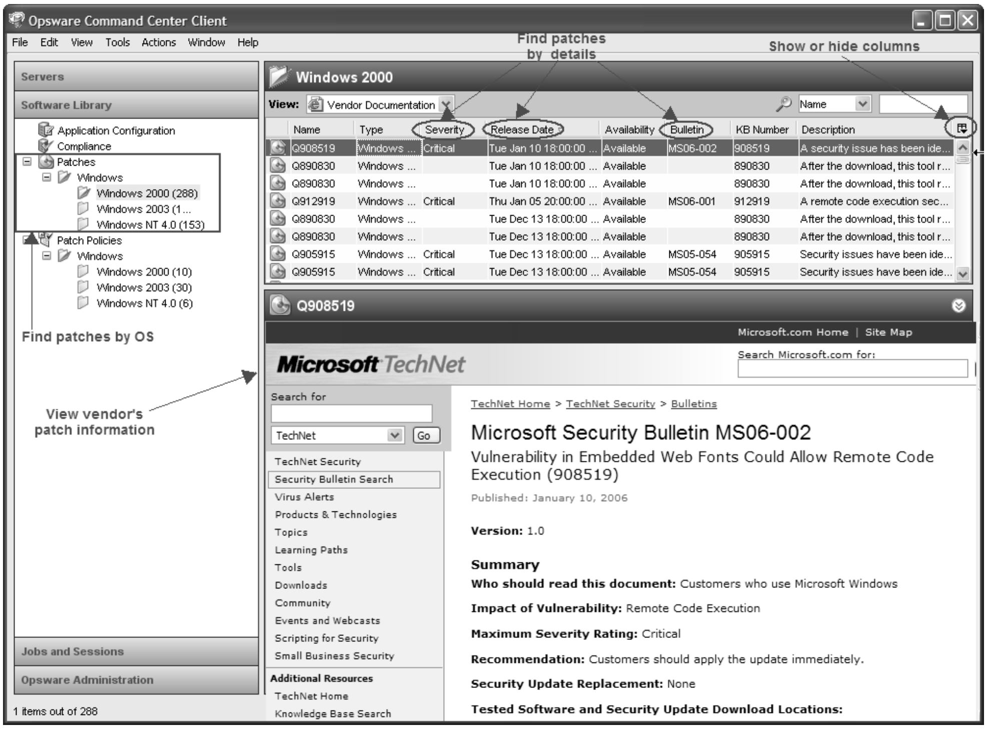

Name	Type	Severity	Release Date	Availability	Bulletin	KB Number	Description
Q908519	Windows ...	Critical	Tue Jan 10 18:00:00 ...	Available	MS06-002	908519	A security issue has been ide...
Q890830	Windows ...		Tue Jan 10 18:00:00 ...	Available		890830	After the download, this tool r...
Q890830	Windows ...		Tue Jan 10 18:00:00 ...	Available		890830	After the download, this tool r...
Q912919	Windows ...	Critical	Thu Jan 05 20:00:00 ...	Available	MS06-001	912919	A remote code execution sec...
Q890830	Windows ...		Tue Dec 13 18:00:00 ...	Available		890830	After the download, this tool r...
Q890830	Windows ...		Tue Dec 13 18:00:00 ...	Available		890830	After the download, this tool r...
Q905915	Windows ...	Critical	Tue Dec 13 18:00:00 ...	Available	MS05-054	905915	Security issues have been ide...
Q905915	Windows ...	Critical	Tue Dec 13 18:00:00 ...	Available	MS05-054	905915	Security issues have been ide...

TechNet Home > TechNet Security > Bulletins

Microsoft Security Bulletin MS06-002

Vulnerability in Embedded Web Fonts Could Allow Remote Code Execution (908519)

Published: January 10, 2006

Version: 1.0

Summary

Who should read this document: Customers who use Microsoft Windows

Impact of Vulnerability: Remote Code Execution

Maximum Severity Rating: Critical

Recommendation: Customers should apply the update immediately.

Security Update Replacement: None

Tested Software and Security Update Download Locations:

Patch Management for Windows Prerequisites

You must have Internet Explorer 6.0 or later installed on a managed server to support Patch Management. This version of Internet Explorer supports the Microsoft XML parser and related DLL files that are required for its native Microsoft Baseline Security Analyzer

(MBSA) tool (mbsacli.exe). Opsware SAS uses version 1.2.1 and 2.0 of the MBSA tool for patch management. Vendor-recommended patches that are installed during the patch reconcile process are based on MBSA 2.0.

Microsoft Patch Database

The Microsoft patch database contains information about released patches and how they should be applied. Patch Management compares all Windows servers to this database to enable the policy setter to determine which patches must be applied.

Microsoft posts patches on its web site on the second Tuesday of each month, unless a special circumstance requires an immediate release. Windows patches released on *patch Tuesday* are available immediately to import into Opsware SAS. Before Patch Management can install a patch on a managed server, the patch must be downloaded from the Microsoft web site and imported into the Software Repository. You can download and import patches with either the OCC Client or with a script. For information about automatically importing the Microsoft patch database, see the *Opsware® SAS Administration Guide*.

Once every twenty-four hours, the Opsware Agent on a Windows server compares the server's current state against the Microsoft patch database (based on the latest version of the MBSA) that has been imported into Opsware SAS by the patch administrator. The Opsware Agent reports the results of that comparison, and the data is stored in the Model Repository. When a user requests a patch compliance scan of a Windows server, the data is retrieved from the Model Repository and displayed in the OCC Client. By storing the data in the Model Repository, rather than performing an actual comparison on the server itself when a user requests an analysis, the data can be quickly retrieved and displayed.

If you perform a patch analysis of a Windows server immediately after importing a new version of the Microsoft patch database, the analysis does not yet include the data from the new patch database. Instead, Opsware SAS reports the data from the last time that the Opsware Agent recorded the results of its comparison. For example, the Opsware 5.5 Agent on a Windows server uses Microsoft's latest detection engine (MBSA 2.0) to identify installed patches. If you used a previous version of the Opsware Agent to create a package of installed patches (from a server snapshot), a previous version of Microsoft's detection engine (MBSA 1.2.1) was used. Because different versions of MBSA were used to identify patches installed on a Windows server, you should expect to see a difference between the list of installed patches that the OCC Client displays and the installed patches in the package that was created from a snapshot.

While MBSA 2.0 can include programs that are not patches in the Microsoft patch database, such as Malicious Software Removal Tool entries, these programs are excluded from Patch Management.

Opsware SAS Integration

When a server is brought under management by Opsware SAS, the Opsware Agent installed on the server registers the server's configuration, including installed patches, with Opsware SAS. (The Opsware Agent repeats this registration every 24 hours.) This information, which includes data about the exact operating system version, hardware type, installed software and patches, is immediately recorded in the Model Repository. Also, when you first provision a server with Opsware SAS, the same data is immediately recorded.

When a new patch is issued, you can use the OCC Client to immediately identify which servers require patching. Opsware SAS provides a Software Repository where you upload patches and other software. Users access this software from the OCC Client to install patches on the appropriate servers.

After a server is brought under management, you should install all Windows patches by using the Patch Management feature. If you install a patch manually, Opsware SAS does not have data about that patch until the next software registration. If you install a patch manually, it can take as long as 24 hours until the data about that server in the Model Repository is up-to-date. However, whenever you install patches with Opsware SAS, the Opsware Agent immediately updates the information about the server in the Model Repository.

You cannot use Opsware SAS to uninstall a patch that was not installed by using the Patch Management feature.

Support for Patch Testing and Installation Standardization

Opsware SAS offers features to minimize the risk of rolling out patches. When a patch is initially imported into Opsware SAS, its status is marked as untested (Limited) and only administrators with the required permissions can install it.

The patch administrator then defines patch installation and uninstallation options and tests the patch. Only after the patch is tested and the patch administrator marks it as available for use (Available) can other administrators install the patch.

The Patch Management feature allows you to standardize the way that patches are installed and uninstalled, thereby preventing ad-hoc installation procedures. Patch administrators standardize patch installation by providing pre-install and post-install scripts, install and uninstall flags, instructions on when to reboot, and how to handle error codes from the pre-install and post-install scripts.

Supported Windows Patch Types

The following table lists the Windows patch types that Patch Management supports.

Table 12-1: Windows Patch Types

OS VERSIONS	PATCH TYPES
Windows NT 4.0	Windows Hotfix Windows OS Service Pack
Windows 2000	Windows Hotfix Windows OS Service Pack
Windows 2003	Windows Hotfix Windows OS Service Pack

Supporting Technologies for Patch Management

Patch Management uses patching utilities and technologies for each supported Windows operating system. Opsware SAS uses these tools behind the scenes, which allows you to perform patch management through a single interface, without having to worry about invoking a number of different patching utilities.

The following patch management and installation tools are used for the supported Windows operating systems:

mbsacli.exe: Lists and verifies patches that are installed on a managed server. Detects which application files are already installed on a managed server and, subsequently, recommends the correct patch to install if multiple patches have the same QNumber.

msiexec.exe: Installs and uninstalls MSI packages.

qchain.exe: Enables a single reboot when you are installing more than one hotfix.

unzip.exe: Extracts info-zip compatible zip archives.

Windows Update Agent: Microsoft framework for patch update.

Windows Hotfixes

After a Microsoft Windows hotfix is imported into Opsware SAS, you can specify options to reboot the server when a hotfix is installed or uninstalled. A Windows hotfix typically requires a reboot if it updates system files. This reboot enables Opsware SAS to use the newly updated system files.

When a hotfix is installed along with other hotfixes, this process is called hotfix chaining. If one or more hotfixes normally require that the server is rebooted, the reboot can sometimes be postponed until all hotfixes have been installed. The user performing the install must first run Qchain.exe before doing the reboot, in order to be sure that the Pending File Rename Queue is correctly ordered.

Postponing reboots is not always possible, due to a defect in Qchain.exe that was resolved in December 2002. All Windows hotfixes created after May 2001 included the Pending File Rename Queue manipulation logic in Qchain.exe. Therefore, all hotfixes created between May 2001 and December 2002 are vulnerable to the same Qchain.exe defect. See the Microsoft Article Q815062 for more information.

If a Windows Service Pack or Security Rollup Package is being installed in the same hotfix chaining process, they will require a reboot and cannot be postponed. Before the reboot that is associated with this package occurs, QChain.exe must be run.

When multiple hotfixes are chained by Opsware SAS, the setting that specifies that a reboot on install is required for each hotfix is honored. Opsware SAS analyzes the set of hotfixes being installed to determine whether one or more reboots can be postponed until the end of the chaining operation.

If you are installing a Windows hotfix that does not support the -z flag, remember to use the -z option to prevent the Patch Management feature from passing in the -z flag.

Opsware SAS examines the date each hotfix was created, to determine whether any associated reboot could be safely postponed until the end of the chained installation.

Opsware SAS will *not* reorder the install order of the chained hotfixes (as an attempt to further reduce the number of reboots), whether or not Service Pack or Security Rollup Packages are being installed in the chained operation.

When Opsware SAS installs a hotfix in isolation (not as part of a chained installation operation), Opsware SAS honors the value of the reboot on the install operation.

Opsware SAS runs Qchain.exe on the managed server after the install of each Windows hotfix and before any associated reboot to guard against problems associated with an incorrectly ordered Pending File Rename Queue. This problem could occur if another hotfix was installed on the managed server outside of Opsware SAS.

Patch Management Roles for Windows

Opsware SAS provides support for rigorous change management by assigning the functions of patch management to several types of users in an organization: a policy setter, a patch administrator, and a system administrator.

Policy Setter

The policy setter is a member of a security standards group that reviews patch releases from operating system vendors and determines which vendor patches will be included in the organization's patch policies. A policy setter is responsible for reviewing the latest security threats and the patches that vendors have released to address these problems. A policy setter is generally known as an expert in the operating systems and applications that they manage, and are able to assess the necessity of applying patches issued by vendors. A policy setter is able to diagnose common problems that arise after patches are installed, allowing for a thorough test of the patch application process.

Patch Administrator

The patch administrator has the authority to import, test, and edit patch options. The patch administrator is often referred to as the security administrator in an organization. A patch administrator is granted specific permissions to import patches into Opsware SAS, test the patches, and then mark them as available for use. Basic users can import patches, but they cannot install them or mark them as available. Patch administrators are also able to edit patch options (such as installation scripts) through patch management. Other types of users are not allowed to import or edit patches. Typically, a patch administrator imports the Microsoft patch database and tests patches on non-production reference hardware. After testing the patches and determining that the patches are safe to apply to production systems, a patch administrator marks the patches available in the Software Library and then advises the system administrators that they must apply the approved patches.

System Administrator

The system administrator installs patches (that have been approved for use) uniformly and automatically, according to the options that the patch administrator specifies. The system administrator is an Opsware user who is responsible for the day-to-day maintenance of the servers in a deployment. These users are not required to have the same level of expertise in low-level system details as the policy setter and patch administrator. Because the patch administrator has set up the patch installation, the system administrators can attach policies to servers, set an exception for a patch, and install patches on a large number of managed servers. They are responsible for searching for the servers that require the approved patch, installing the patches, and verifying that the patches were successfully installed. The system administrator can import patches but cannot install a patch until the patch administrator has marked it as available. The system administrator can also uninstall patches.

These responsibilities are enforced by assigning permissions for managing patches in Opsware SAS. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

Patch Management Process

The Windows patching process consists of several key phases: setup, policy management, patch compliance, and deployment. Setup steps include getting the Microsoft database (patches and metadata) into Opsware SAS, identifying products you want to track patches for, and configuring patch compliance. Policy management steps include investigating released patches, creating and updating patch policies or exceptions, marking patches available to use, and attaching policies or exceptions to servers or server groups. Patch compliance steps include running compliance scans to determine whether a server is out of compliance, reconciling patches in policies, setting up installation options, and installing applicable patches. To deploy patches on demand, you can import the required patches, test them, update policies or create new policies, mark them available to use, specify install options, and install the required patches. Figure 12-2 and Figure 12-3 illustrate these phases and steps.

Figure 12-2: Windows Patching Process: Part A and Part B

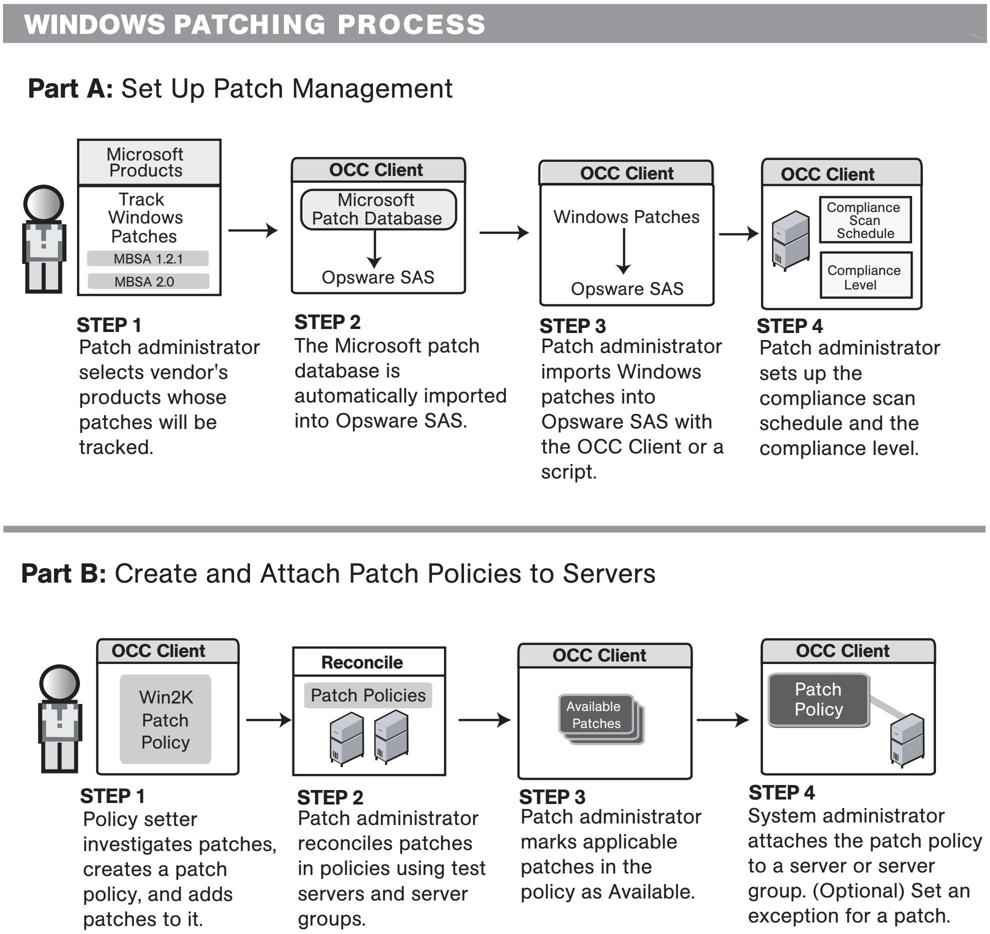

WINDOWS PATCHING PROCESS

Part A: Set Up Patch Management

STEP 1
Patch administrator selects vendor's products whose patches will be tracked.

STEP 2
The Microsoft patch database is automatically imported into Opsware SAS.

STEP 3
Patch administrator imports Windows patches into Opsware SAS with the OCC Client or a script.

STEP 4
Patch administrator sets up the compliance scan schedule and the compliance level.

Part B: Create and Attach Patch Policies to Servers

STEP 1
Policy setter investigates patches, creates a patch policy, and adds patches to it.

STEP 2
Patch administrator reconciles patches in policies using test servers and server groups.

STEP 3
Patch administrator marks applicable patches in the policy as Available.

STEP 4
System administrator attaches the patch policy to a server or server group. (Optional) Set an exception for a patch.

Figure 12-3: Windows Patching Process: Part C and Part D

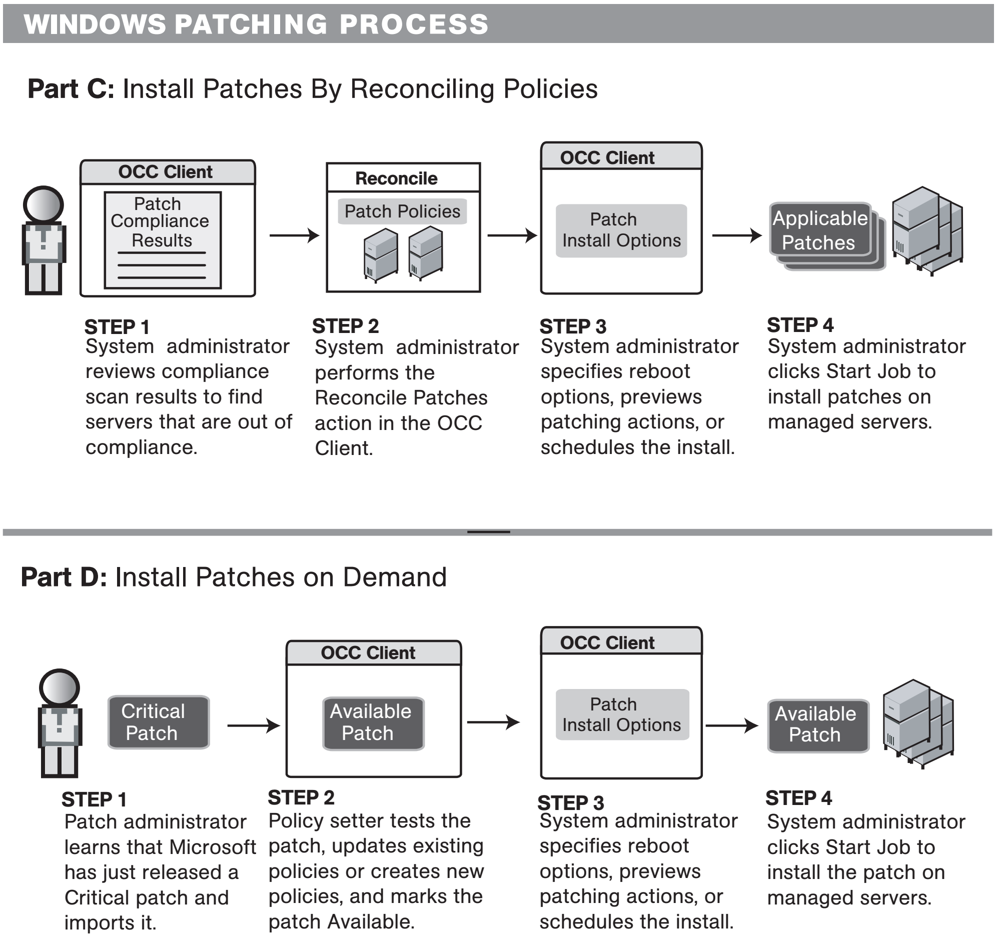

Part C: Install Patches By Reconciling Policies

STEP 1
System administrator reviews compliance scan results to find servers that are out of compliance.

STEP 2
System administrator performs the Reconcile Patches action in the OCC Client.

STEP 3
System administrator specifies reboot options, previews patching actions, or schedules the install.

STEP 4
System administrator clicks Start Job to install patches on managed servers.

Part D: Install Patches on Demand

STEP 1
Patch administrator learns that Microsoft has just released a Critical patch and imports it.

STEP 2
Policy setter tests the patch, updates existing policies or creates new policies, and marks the patch Available.

STEP 3
System administrator specifies reboot options, previews patching actions, or schedules the install.

STEP 4
System administrator clicks Start Job to install the patch on managed servers.

Patch Details

A patch is a piece of object code (binaries) that is inserted into (patched into) an executable program to temporarily fix a known defect. Patch Management displays detailed information (metadata) about a patch.

Figure 12-4: Patch Metadata

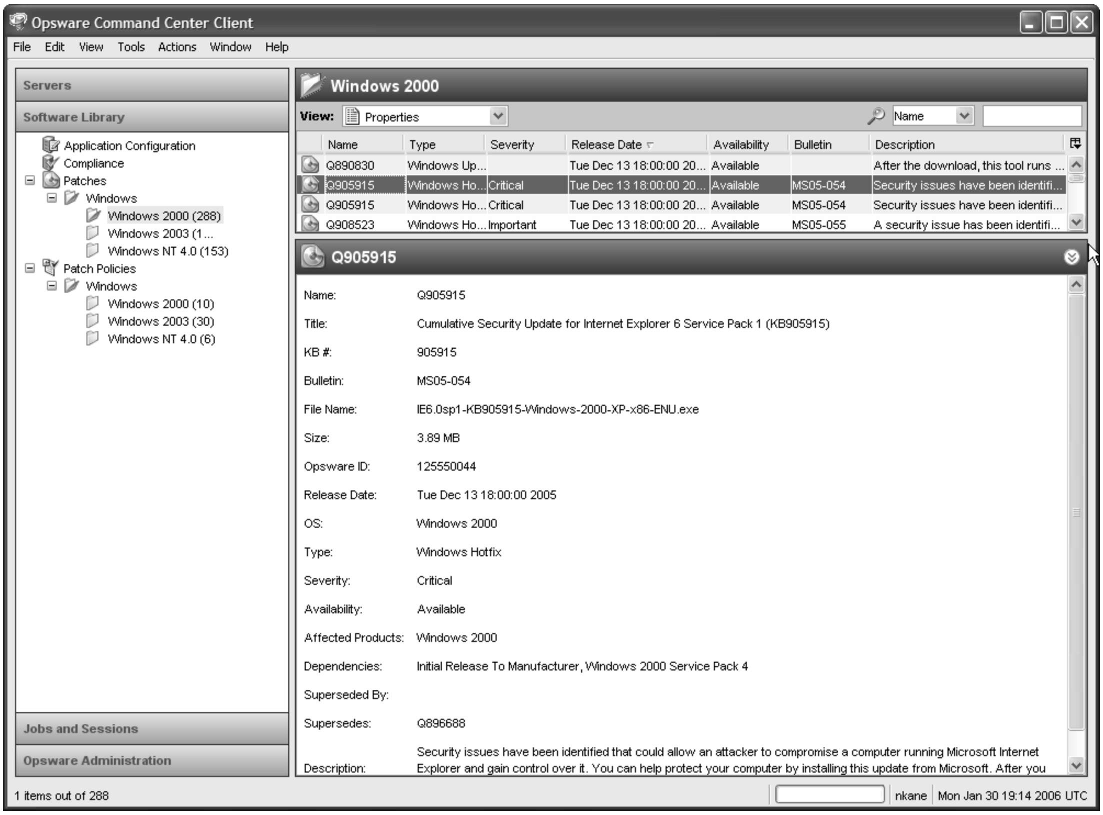

Name	Type	Severity	Release Date	Availability	Bulletin	Description
Q890830	Windows Up...		Tue Dec 13 18:00:00 20...	Available		After the download, this tool runs ...
Q905915	Windows Ho...	Critical	Tue Dec 13 18:00:00 20...	Available	MS05-054	Security issues have been identifi...
Q905915	Windows Ho...	Critical	Tue Dec 13 18:00:00 20...	Available	MS05-054	Security issues have been identifi...
Q908523	Windows Ho...	Important	Tue Dec 13 18:00:00 20...	Available	MS05-055	A security issue has been identifi...

Name:	Q905915
Title:	Cumulative Security Update for Internet Explorer 6 Service Pack 1 (KB905915)
KB #:	905915
Bulletin:	MS05-054
File Name:	IE6.0sp1-KB905915-Windows-2000-XP-x86-ENU.exe
Size:	3.89 MB
Opsware ID:	125550044
Release Date:	Tue Dec 13 18:00:00 2005
OS:	Windows 2000
Type:	Windows Hotfix
Severity:	Critical
Availability:	Available
Affected Products:	Windows 2000
Dependencies:	Initial Release To Manufacturer, Windows 2000 Service Pack 4
Superseded By:	
Supersedes:	Q896688
Description:	Security issues have been identified that could allow an attacker to compromise a computer running Microsoft Internet Explorer and gain control over it. You can help protect your computer by installing this update from Microsoft. After you

Patch metadata includes the following information:

- **Title**: The title of the Microsoft Knowledge Base article for this patch.
- **KB #**: The Microsoft Knowledge Base article ID number for this patch.
- **Bulletin** (Optional): The Microsoft Security Bulletin ID number for this patch.
- **File Name**: The name of the .exe file for this patch.
- **Opsware ID**: The Opsware SAS unique ID for the patch.
- **Release Date**: The date that Microsoft released this patch.
- **OS**: The Windows operating systems that are known to be affected by this patch.

- **Type**: The type of patch, such as Windows Hotfix or Windows Update Rollup.
- **Severity** (Optional): One of following Microsoft severity ratings for this patch:
  - **Critical**: A patch whose exploitation could allow the propagation of an internet worm, without user action.
  - **Important**: A patch whose exploitation could result in a compromise of the confidentiality, integrity, or availability of user data, or of the integrity or availability of processing resources.
  - **Moderate**: Exploitability is mitigated to a significant degree by certain factors, such as default configuration, auditing, or difficulty of exploitation.
  - **Low**: A patch whose exploitation is extremely difficult, or whose impact is minimal.
- **Availability**: The status of a patch within Opsware SAS, which can be one of the following:
  - **Not Imported**: The patch is listed in the Microsoft Patch Database, but has not been imported (uploaded) into Opsware SAS.
  - **Limited**: The patch has been imported into Opsware SAS but cannot be installed. This is the default patch availability.
  - **Available**: The patch has been imported into Opsware SAS, tested, and has been marked available to be installed on managed servers.
  - **Deprecated**: The patch cannot be added to patch policies or set as a patch policy exception but can still be installed.
- **Affected Products**: Information from MBSA that identifies other Microsoft software that is known to be affected by this patch.
- **Dependencies**: Microsoft products that this patch requires. The patch cannot be installed if these products do not already exist on the server.
- **Superseded By**: A list of patches that this patch is obsoleted (replaced) by. This relationship does not apply to MBSA 1.2.1 patches.
- **Supersedes**: A list of patches that this patch obsoletes (replaces). This relationship does not apply to MBSA 1.2.1 patches.

Patch Dependencies and Supersedence

Patch metadata identifies all known dependency and supersedence relationships between patches and Windows products, and between patches and other patches. Dependency relationships identify Windows products that must already exist on a server before you can install a certain patch. Supersedence relationships identify patches that obsolete (supersede) or are obsoleted (superseded) by other patches. In Patch Management, *supersedes* means that one patch replaces another and *superseded by* means that the patch you are installing is obsoleted by another patch.

For all MBSA 2.0 patches, Patch Management analyzes this information to determine the viability of a patch installation. For example, if you are reconciling patches and a superseding patch is already installed, the patch will not be installed. If you try to install a superseded patch and the superseding patch is available and included in a patch policy, the superseded patch will not be installed. In these situations, Patch Management will display a warning message indicating that there is a patch relationship conflict. Patch Management does not analyze this information for MBSA 1.2.1 patches.

Patch Management does not detect whether two patches are mutually exclusive—where either one can be installed but not both. Subsequently, Patch Management does not prevent you from installing both patches on a server. This means that you may be able to install both a superseded patch and a superseding patch on a server.

Viewing Windows Patches

The OCC Client displays information about Microsoft Windows patches that have been imported into Opsware SAS.

To view information about a patch, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.

2. Expand Patches and select a specific Windows operating system.

   The Content pane will display all of the patches listed in the Microsoft Patch Database for the Windows operating system that you selected.

3. (Optional) Use the column selector to sort the patches according to Name, Type, Severity, Availability, Release Date, and Bulletin Number.

4 In the Content pane, open a patch to view its properties in the Patch Window.

Editing Patch Properties

You can edit a patch's Description, Availability, Install Parameters, and Uninstall parameters. Due to the nature of the type of patch, some properties are not editable. For example, you cannot turn the reboot on install option of a Windows Service Pack off.

The Availability property indicates the status of the patch in Opsware SAS. If the Availability is Not Imported, you cannot change this property.

You can set the install and uninstall parameters on either the patch properties page or in the Patch Actions only when you are installing or uninstalling one patch at a time. The parameters on the properties page are saved in the Model Repository, but the parameters in Patch Actions are used only for that action. The parameters in Patch Actions override those on the patch properties page.

To edit the patch properties, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.

2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.

3 In the Content pane, open a patch to view its properties in the Patch Window.

4 Edit any of the following fields: Description, Availability, and the Install and Uninstall parameters.

5 From the **File** menu, select **Save** to save your changes.

Importing Custom Documentation for a Patch

The Custom Documentation view of a patch displays text files that have been imported from the local file system. Non-plain text file types, such as HTML or .doc, are not supported.

To import your own documentation for a patch, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.

2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.

3. In the Content pane, open a patch to view its properties in the Patch Window.
4. From the View pane, select Custom Documentation.
5. From the **Actions** menu, select **Import Custom Documentation**.
6. From Import Custom Documentation window, locate a text file and click **Import**.

Finding Vendor-Recommended Patches

To find out which patches Microsoft recommends for a particular server (based on MBSA 2.0), perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and then select All Managed Servers.
2. From the View drop-down list, select Patches.
3. From the Content pane, select a server that is running Opsware Agent 5.5 and a Windows 2000 with Service Pack 3 (or higher) operating system or a Windows 2003 operating system.
4. From the Preview pane, select Patches Recommended By Vendor in the drop-down list to display these types of patches for the selected server.

Finding Servers That Have a Patch Installed

To find out which servers have a particular patch installed, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.
2. Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
3. From the Content pane, select a patch.
4. From the View drop-down list in the Content pane, select Servers.
5. From the Show drop-down list for the selected patch, select Servers with Patch Installed.

   You can browse a server in this list to view a list of all installed patches. Please note that this list may display a more complete list of installed patches than the list you will find in the Windows Add or Remove Programs utility.

Finding Servers That Do Not Have a Patch Installed

To find out which servers do not have a particular patch installed, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.
2. Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
3. From the Content pane, select a patch.
4. From the View drop-down list, select Servers.
5. From the Show drop-down list, select Servers without Patch Installed.

Importing a Patch

Windows patches are downloaded from the Microsoft web site and then imported (uploaded) into Opsware SAS. To see if a patch has been imported, view the patch's Availability property. The Availability of an imported patch is either Limited, Available, or Deprecated. A patch can be imported with the OCC Client or with a script. For information about the script, see "Automatically Importing Windows Patches" in the *Opsware® SAS Administration Guide*.

To import a patch with the OCC Client, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.
2. Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
3. From the Content pane, select a patch.
4. To import a patch directly from the Microsoft web site, from the **Actions** menu, select **Import ➤ Import from Vendor**.

   The Import from Vendor window displays the URL of the patch's location on the Microsoft web site. You can override this URL, as needed.

   Or

   To import a patch that has already been downloaded to your local file system, from the **Actions** menu, select **Import ➤ Import from File**.

In the file browser window, locate the patch.

5 Click **Import**.

Exporting a Patch

To export a patch from Opsware SAS to the local file system, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.

2 Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.

3 From the Content pane, select a patch.

4 From the **Actions** menu, select **Export**.

5 In the Export Patch window, enter the *folder* name that will contain the patch file in the File Name field.

6 Click **Export**.

Exporting Patch Information

You can export the following information about patches that are installed on a server, patches that are recommended by the vendor, and patches with model information on the selected server (such as patch policies or patch policy exceptions) *and* are also recommended by the vendor into a comma-separated value (CSV) file:

- Server Name
- OS
- Service Pack (This is the service pack level of the server being reported, such as Service Pack 0, Service Pack 1, and so on.)
- KB#
- Bulletin (This is the MSYY-XXX ID associated with a hotfix, such as MS05-012, MS06-012, and so on. If the MSYY-XXX ID is unknown, this column will be blank.)
- Description
- Time Queried (This is the last software registration by the Agent.)
- Time Installed (This is the time the patch was installed.)
- Type (This is the patch type.)

- Compliance Level (This is an integer that represents the compliance level.)
- Compliance (This is a text description that displays when you place your cursor over the Compliance column in the Patch Preview pane.)
- Exception Type
- Exception Reason

Patch Management will display all of the text, including commas, from the Description field displayed in the Patch Properties window in the Description column in the .csv file. To preserve commas in the Description column and keep all text together in that column, double quotes will be converted to single quotes. This does not distort the semantics of the patch description.

To ensure that all of the text about a patch displays in the Description field in the .csv file, Patch Management surrounds the entire description (that you see in the Patch Properties window) with double quotes.

To export the patch information to a CSV file, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and then select All Managed Servers.
2. From the Content pane, select one or more managed servers.
3. From the **Actions** menu, select **Export Patch Info to CSV**.
4. In the Export to CSV window, navigate to a folder and enter the file name.
5. Verify that the file type is Comma Separated Value Files (.csv). If you did not include the .csv extension in the file name field, Patch Management will append it only if you have the .csv file type selected.
6. Click **Export** to save the patch information in a .csv file or click **Cancel** if you do not want to export the patch information.

Deleting a Patch

This action removes a patch from Opsware SAS, but does not uninstall the patch from managed servers. A patch cannot be deleted if it is attached to a policy or if an exception has been set for it.

Do not delete all of the patches from Opsware SAS. If you do so accidentally, contact your Opsware, Inc. support representative for assistance in importing the patches back into Opsware SAS.

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.
2. Expand Patches and select a specific Windows operating system. The Content pane will display all patches associated with that operating system.
3. From the Content pane, select a patch.
4. From the **Actions** menu, select **Delete Patch**.

Policy Management

In Patch Management, patch policies and patch policy exceptions enable you to customize patch distribution in your environment. Policies and exceptions define which patches should be installed or not installed on certain managed servers.

You can choose to have patching in your server environment comply to the model that these policies and exceptions define or you can choose to deviate from this model. If you choose to leverage from the patch policies and exceptions and you perform ad hoc patch installs, you need to reconcile patches to get the applicable patches installed on servers.

Patch Policy

A patch policy is a group of patches that you want to install on Opsware SAS managed servers. All patches in a patch policy must apply to the same Windows operating system.

A patch policy provides broad flexibility in how you distribute patches. For example, you can create a patch policy that contains security patches that you want to distribute only to servers used by your sales force. You can also create a patch policy that contains security patches that are applicable to specific software that is already installed on a server, such as Exchange Server, Internet Information Services (IIS), SQL Server, and so on. Or, you can create a patch policy that includes all patches ranked critical (by Microsoft) to install them on all servers that are used by everyone in your organization.

If you do not want to create a patch policy, you can use the vendor-recommended set of patches (by operating system) as a default patch policy, such as the patches provided by MBSA.

You can attach as many patch policies as you want to servers or server groups. If several policies are attached to one server, the installation logic is cumulative–all patches listed in all attached policies will be installed on the server. The Reconcile Patches task window allows you to select an individual patch policy to reconcile. You do not have to reconcile all policies attached to a server. You cannot nest patch policies.

If a description of the patch policy is defined, it is recorded in the server's patched state (in the Model Repository). This information enables Patch Management to report on patch policies for patch compliance purposes–to explain how patch policies compare with corresponding patch policy exceptions.

Patch Management supports the following types of patch policies:

- **User-defined patch policy**: This allows an Opsware SAS user to specify which patches are included in a policy. User-defined patch policies can be edited or deleted by a user who has permissions.

  A user-defined patch policy allows a policy setter to opt out of patches. The policy setter can create a (user-defined) patch policy that is a subset of all available patches (that are in a vendor-recommended patch policy) to apply only those patches that their environment needs.

- **Vendor-recommended patch policy**: Membership of patches is defined by what MBSA recommends on a server-by-server basis. Vendor-recommended patch policies are system defined and cannot be edited or deleted by a user.

You can only export user-defined patch policies. You cannot export vendor-recommended patch policies.

Patch policies have the following characteristics:

- All patches in a patch policy must apply to the same operating system, such as Windows.
- A patch policy is associated with an operating system version, such as Windows 2003.

- A patch policy has a name and can (optionally) include a description that explains its purpose.
- A patch policy can be either user defined or vendor defined.
- A patch policy does not have sub-policies. There is no inheritance.
- A patch policy is Customer Independent, which means that patches in the policy can be installed on any managed server, no matter what customer it is associated with. For more information, see “Customer Accounts in Opsware SAS” on page 146.
- A patch policy is always public.
- A patch policy can be attached to zero or more servers or public server groups.
- More than one patch policy can be attached to a server or public server group.
- A patch policy can be created, edited, and deleted by users who have permissions. Only user-defined patch policies can be created, edited, and deleted by a user who has permissions.

Patch Policy Exception

A patch policy exception identifies a single patch that you want to explicitly include in or exclude from a specific managed server, along with an optional reason for why the exception exists. The patch in a patch policy exception must apply to the same Windows operating system that the established patch policy is attached to.

A patch policy exception allows you to deviate from an established patch policy (one that is already attached to a server or server group) by deselecting or adding individual patches for a server. Since patch policy exceptions override all patch policies attached to a server, you can use them to intentionally deviate from a patch policy on a server-by-server basis.

If a reason for a patch policy exception is defined, the description is recorded in the server’s patched state (in the Model Repository). This information enables Patch Management to report on patch policy exceptions for patch compliance purposes—to explain how patch policy exceptions compare with corresponding established patch policies. All users who have access to the managed server can view attached patch policy exceptions.

Patch Management supports the following types of patch policy exceptions:

- **Always Installed**: The patch should be installed on the server, even if the patch is not in the policy.

- **Never Installed**: The patch should not be installed on the server, even if the patch is in the policy.

If you ever need to override a patch policy exception, you can manually install a patch.

The following information summarizes detailed characteristics of a patch policy exception:

- A patch policy exception can (optionally) include a description that explains its purpose.
- A patch policy exception can have a rule value of Never Installed or Always Installed.
- A patch policy exception can be set for one patch and one server of the same operating system version. If a patch policy exception is set for a public server group and a server in that group does *not* match the operating system version specified in the patch policy exception, the patch policy exception is *not* applied.
- A patch policy exception can be set, copied, and removed by users who have permissions.

Precedence Rules for Applying Policies

By creating multiple patch policies and patch policy exceptions (that are either directly attached to a server or attached to a server group), you control which patches should be installed or not installed on a server. A precedence hierarchy in Patch Management delineates how a policy or an exception, whether the policy or exception is attached at the server or server group level, is applied to a patch installation.

The following precedence rules apply to policies and exceptions:

- Patch policy exceptions that are directly attached to a server always take precedence over patch policies that are directly attached to a server.
- Patch policies that are directly attached to a server take precedence over patch policies and patch policy exceptions that are attached to a public server group.
- Patch policy exceptions that are attached to a public server group take precedence over patch policies that are attached to a public server group.

- If a server is in multiple public server groups, a Never Installed patch policy exception type always take precedence over an Always Installed patch policy exception type for the same patch.

Reconcile Patches Process

To ensure patch compliance, Patch Management identifies vulnerable managed servers and simultaneously deploys patches to many servers when a reconcile patches process is performed. The reconcile patches process applies an entire patch policy and even multiple policies to an operation that examines the patch policy and the managed servers that it is attached to. (A policy must be attached to a server or a server group before you can reconcile the policy with that server or server group.)

✓ The reconcile patches process requires that the selected managed server is running Opsware Agent 5.5 and a Windows 2000 Service Pack 3 (or higher) operating system or a Windows 2003 operating system. You cannot use the reconcile patches process if the selected managed server is running a Windows NT4.0 operating system, a Windows 2000 RTM (no service pack), Service Pack 1, or Service Pack 2 operating system, or if the server is not running Opsware Agent 5.5. Use the Install Patch task window to install patches on servers that are running these operating systems or Opsware Agents 4.5 or later.

As a best practice, each time a policy setter reviews the latest Microsoft patch releases and subsequently updates a patch policy (by adding new patches to a policy), a policy reconcile should be performed. In these situations, a reconcile patches process provides demand forecasting information that allows you to determine how patch policy changes would impact servers that this policy is attached to.

If the reconcile discovers any (applicable) missing patches, these patches will be installed on the servers. The reconcile process does not uninstall (remove) patches from a server.

After Opsware SAS determines what packages need to be installed to complete the reconcile operation, reconcile uses a set of standard system utilities to complete the operation. For more information, see “Supporting Technologies for Patch Management” on page 459.

To help you optimally manage the conditions under which patch policies are reconciled, Patch Management allows you to specify reconcile options and pre and post reconcile scripts, and set up email notifications to alert you about the status of the reconcile process. The Reconcile Patches task window guides you through setting up these conditions.

Figure 12-5: Reconcile Patches Task Window

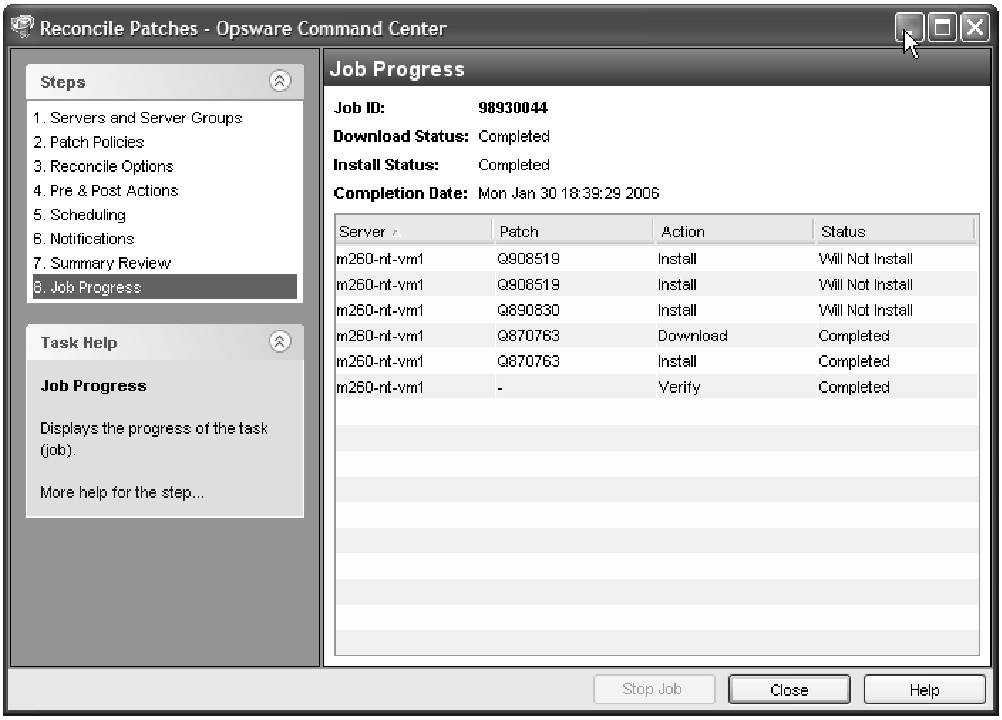

Job Progress

Displays the progress of the task (job).

More help for the step...

Job Progress

Job ID: **98930044**
Download Status: Completed
Install Status: Completed
Completion Date: Mon Jan 30 18:39:29 2006

Server	Patch	Action	Status
m260-nt-vm1	Q908519	Install	Will Not Install
m260-nt-vm1	Q908519	Install	Will Not Install
m260-nt-vm1	Q890830	Install	Will Not Install
m260-nt-vm1	Q870763	Download	Completed
m260-nt-vm1	Q870763	Install	Completed
m260-nt-vm1	-	Verify	Completed

Reconciling Patches

This action installs the patches in a policy that has been attached to managed servers. (This action does not uninstall patches.) A patch policy can be overridden by an exception, which indicates that a patch is either always or never installed on a particular server.

When you invoke the reconcile patches operation for a server group, patches will only be reconciled if any server in the server group is running Opsware Agent 5.5 and a Windows 2000 with Service Pack 3 (or higher) operating system or a Windows 2003 operating system. The Reconcile Patches option is not available in the Actions menu if the selected server is not running Opsware Agent 5.5 and a Windows 2000 with Service Pack 3 (or higher) operating system or a Windows 2003 operating system.

To reconcile a patch policy, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patch Policies

2. Expand Patch Policies and select a specific Windows operating system. The Content pane will display all patch policies associated with that operating system.

3. From the Content pane, open a patch policy.

4. From the View drop-down list, select Servers.

5. From the Show drop-down list in the Content pane, select Servers with Policy Attached.

6. From the Preview pane, select one or more servers.

7. From the **Actions** menu, select **Reconcile Patches**. The first step of the Reconcile Patches task window appears: Servers and Server Groups.
   For instructions on each step, see the following sections:

   – Setting Reconcile Options

   – Setting Reboot Options for a Patch Reconcile

   – Specifying Pre and Post Install Scripts for a Patch Reconcile

   – Scheduling a Patch Installation for a Patch Reconcile

   – Setting Up Email Notifications for a Patch Reconcile

   – Previewing a Patch Reconcile

   After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

8. Click **Start Job** to launch the reconcile job.

   After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

   If the Reconcile Patches task window remains open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select Refresh from the View menu to update information in the Patch Preview pane.

Setting Reconcile Options

You can specify the following types of reconcile policy options:

- Do not interrupt the reconcile process even when an error occurs with one of the policies.
- Perform the reconcile immediately or at a later date and time.

To set these options, perform the following steps:

1. From the Reconcile Patches task window, click **Next** to advance to the Reconcile Options step.
2. Select one of the following Staged Install Options:
   - **Continuous**: Run all phases as an uninterrupted operation.
   - **Staged**: Allow download and install to be scheduled separately.
3. Select the Error Options check box if you want the reconcile process to continue even when an error occurs with any of the patches or scripts. As a default, this check box is not selected.
4. Click **Next** to go to the next step or click **Cancel** to close the Reconcile Patches task window.

Setting Reboot Options for a Patch Reconcile

To minimize the downtime that server reboots can cause, you can control when servers reboot during a patch installation.

You specify the reboot options in the following two places in the OCC Client:

- Install Parameters tab of the patch properties window
- Pre & Post Actions step of the Reconcile Patches task window

When you are selecting reboot options in the Reconcile Patches task window, Opsware, Inc. recommends that you use Microsoft's reboot recommendations, which is the "Reboot servers as specified by patch properties" option in the task window. If it is not possible to use the Microsoft reboot setting, select the single reboot option, which is the "Do not reboot servers until all patches are installed" option in the task window. Failure to do this can result in the MBSA incorrectly reporting which patches are installed on the server until the next reboot occurs (outside of Opsware control).

The following options in the Reconcile Patches task window determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Reconcile Patches task window; they do not change the Reboot Required option, which is on the Install Parameters tab of the patch properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties**: By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install**: Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots**: Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Do not reboot servers until all patches are installed**: If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

1. From the Reconcile Patches task window, click **Next** to advance to the Pre & Post Actions step.
2. Select one of the Rebooting Options.
3. Click **Next** to go to the next step or click **Cancel** to close the Reconcile Patches task window.

Specifying Pre and Post Install Scripts for a Patch Reconcile

For each patch reconcile, you can specify a command or script to run before reconcile or after reconcile. A pre-install script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-install script fails, the patches would not be installed. A pre-install script could also be used to shut down a service or application before it is patched. A post-install script could be used to perform a certain cleanup process on the managed server.

You can specify the following types of scripts to run on the managed server before or after a reconcile process:

- **Pre-Download**: A script that runs before patches are downloaded from Opsware SAS to the managed server. This is available only if you select Staged in the Reconcile Options step.
- **Post-Download**: A script that runs after patches are downloaded from Opsware SAS to the managed server and before the patch is installed. This is available only if you select Staged in the Reconcile Options step.
- **Pre-Install**: A script that runs before patches are installed on the managed server.
- **Post-Instal**l: A script that runs after patches are installed on the managed server.

To specify a pre-install script, perform the following steps:

1. From the Reconcile Patches task window, click **Next** to advance to the Pre & Post Actions step.
2. Select the Pre-Install tab.

   You may specify different scripts and options on each of the tabs.
3. Select the Enable Script check box. This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
4. Select either Saved Script or Ad-Hoc Script from the drop-down list.

   A Saved Script has been previously stored in Opsware SAS with the Opsware Command Center. To specify the script, click **Select**.

   An Ad-Hoc script runs only for this operation and is not saved in Opsware SAS. Select the Type, such as .BAT. In the Script box, enter the contents of the script, including the drive letter of where the script is located, such as echo dir>> C:\temp\preinstall1.log. If you do not enter a drive letter, the default is %SYSTEMDRIVE%, which is where the system folder of Windows is installed.
5. If the script requires command-line flags, enter the flags in the Command text box.
6. In the User section, if the system is not Local, select Name.
7. Enter the system name, your password, and the Domain name.
8. To stop the installation if the script returns an error, select the Error check box.
9. Click **Next** to go to the next step or click **Cancel** to close the Reconcile Patches task window.

Scheduling a Patch Installation for a Patch Reconcile

Since the two phases of patching can be decoupled, you can schedule when you want patches installed (deployed) to occur independently of when patches are downloaded (staged).

To schedule a patch installation, perform the following steps:

1. From the Reconcile Patches task window, select the Scheduling step. To reach this step, you must have completed the Pre & Post Actions step.

   By default, the Scheduling step displays only the scheduling options for the install phase. If you selected Staged in the Reconcile Options step, the scheduling options for the download phase will also be displayed.

2. Select one of the following Install Phase options:

   - **Run Task Immediately**: This enables you to perform the download or install immediately.

   - **Run Task At**: This enables you to specify a date and time that you want the download or install performed.

3. Click **Next** to go to the next step or click **Cancel** to close the Reconcile Patches task window.

Setting Up Email Notifications for a Patch Reconcile

You can set up email notifications to alert users when the download and installation operations complete successfully or with errors.

To set up email notifications, perform the following steps:

1. From the Reconcile Patches task window, click **Next** to advance to the Notifications step.

2. In the Notify field, the default email address that displays is derived from the Profile of your Opsware user. To add email addresses, click **Add Notification** and enter additional email addresses.

3. To set the notification status, from the drop down list, select Completes, Completes Successfully, and Completes with Errors.

4. Click **Next** to go to the next step or click **Cancel** to close the Reconcile Patches task window.

If you previously selected Staged in the Reconcile Options step, the Notifications pane displays notification options for both the download and install phases.

Previewing a Patch Reconcile

The reconcile preview process provides an up-to-date report about the patch state of servers. The reconcile preview is an optional step that lets you see what patches will be installed on managed servers. This preview process verifies whether the servers you selected for the patch installation already have that patch installed (based MBSA 2.0). In some cases, a server could already have the patch installed if a system administrator had manually installed it, which means that Patch Management does not know about it.

The preview process also reports on dependency and supersedence information, such as patches that require certain Windows products, and patches that obsolete other patches or are obsoleted by other patches. If a dependency is not met, Patch Management will display an error message indicating this condition. For example, if a managed server is running Windows 2000 Service Pack 3 (or higher) or Windows 2003, and an Opsware SAS 5.5 Agent, Patch Management will report that a dependency has not been fulfilled. If you try to install a patch for Service Pack 4 and your server is using Service Pack 3, the reconcile preview will display a "Will Not Install" error message to indicate this discrepancy.

In the Preview, the servers, server groups, and patches that are listed in the Summary Step window will be submitted to reconcile when you click Start Job. Patches that are not recommended by the vendor will be excluded from this list. If there are other patches in the policy with the same QNumber, only the vendor-recommended patch is displayed.

This list enables you to see what patches go on what servers (regardless of any patch policy and server group membership changes that may have occurred). If you Preview a reconcile that is scheduled for a future time, this same list of servers, server groups, and patches will be used, even if changes have occurred to the patch policy or server group memberships.

If you modify parameters in the Reconcile Patches task window after you have already clicked Preview, the preview process will produce an invalid summary of simulated patching actions. For example, if you have already clicked Preview and you add patches, patch policies, servers, or server groups, you must click Preview again for results that include your changes.

The reconcile preview does not report on the behavior of the server as though the patches have been applied.

To preview a policy reconcile, perform the following steps:

1. From the Reconcile Patches task window, click **Next** to advance to the Summary Review step.
2. Verify the information displayed for the Servers, Server Groups, and Patches at the top of the window.
3. (Optional) Click **Preview** to see the separate actions that will be performed when the patch is installed. To view the details of a previewed action, select a row in the table.
4. To launch the installation job, click **Start Job**.

   If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected a specific time, the job will run then.
5. The Job Progress displays in the Reconcile Patches task window.

   The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:

   – **Analyze**: Opsware SAS examines the patches needed for the install, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.

   – **Download**: The patch is downloaded from Opsware SAS to the managed server.

   – **Install**: After it is downloaded, the patch is installed.

   – **Final Reboot**: If this action is specified in the Pre & Post Actions step, the server is rebooted.

   – **Run Script**: If this action is specified in the Pre & Post Actions step, a script is run before or after the install.

   – **Install & Reboot**: When a patch will be installed is also when the server will be rebooted.

   – **Verify**: Installed patches will be included in the software registration.
6. To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select Jobs and Sessions to review detailed information about the job. See “Browsing Job Logs” on page 179.

7 Click **Stop Job** to prevent the job from running or click **Close** to close the Reconcile Patches task window.

Verifying Patch Policy Compliance

A patch policy identifies patches that should be installed on a managed server. A patch policy exception identifies a patch that should or should not be installed.

To determine whether a managed server complies with patch policies and exceptions, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Servers and then select All Managed Servers.

2 From the Content pane, select Patches from the View drop-down list and then select a server.

3 The status of patches will display. The Compliance column shows whether a patch is Compliant, Partial, or Non-Compliant.

- **Non-Compliant** The patch is installed on the server, but is not in the policy, or the patch is not installed on the server but is in the policy.
- **Partial**: The policy and exception do not agree, and the exception does not have data in the Reason field.
- **Compliant**: This status indicates one of the following conditions:
  - A patch is installed on the server and is in a policy, or a patch is not installed on the server and is not in a policy.
  - A patch is installed on the server and there are additional patches with the same QNumber in a patch policy or exception. In this case, all patches with the same QNumber are considered installed when Patch Management calculates patch compliance.
  - A patch is not installed on the server and is in a patch policy or has an always install exception, and is not recommended by the vendor. In this case, the patch is considered as if it has a never install exception because it is not recommended by the vendor.

If the icon is Compliant, no further action is required. If the icon is Partial or Non-Compliant, perform the following steps.

- **No Indicator:** The patch is installed on the server and it is not in a policy and there is no exception.

4. Select a server.
5. From the Show drop-down list, select Patches Needed. Non-compliant patches will display.
6. In the Preview pane, select a patch and move the cursor over the icon in the Compliance column to view patch compliance information about a server.
7. To find out which patches are installed on the server but should not be, from the Show drop-down list, select Patches Installed.
8. In the Preview pane, note the patches that have the Non-Compliant icon. Select a patch and move the cursor over the icon to view patch compliance information about a server.

Creating a Patch Policy

A patch policy is a set of patches that should be installed on a managed server. When it is first created, a patch policy contains no patches and is not attached to servers.

To create a patch policy, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patch Policies.
2. Select a specific Windows operating system.
3. From the **Actions** menu, select **Create Patch Policy**.

   The name of the policy you just created is New Patch Policy n, where n is a number based on the number of New Patch Policies already in existence.
4. From the Content pane, open the New Patch Policy.
5. (Optional) In the Name field of the Properties, enter a name that describes the purpose or contents of the policy.

Deleting a Patch Policy

This action removes a patch policy from Opsware SAS but does not remove or uninstall patches from managed servers. You cannot delete a patch policy if it is attached to servers or server groups. You must first detach the policy from the servers or server groups before removing it from Opsware SAS.

To delete a patch policy from Opsware SAS, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patch Policies.
2. Select a specific Windows operating system.
3. From the Content pane of the main window, select a policy.
4. From the **Actions** menu, select **Delete Patch Policy**.

Adding a Patch to a Patch Policy

This action adds a patch to a patch policy, but does not install the patch on a managed server. The patch will be installed when the policy is reconciled.

To add a patch to a patch policy, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.
2. Select a specific Windows operating system and view the list of Windows patches.
3. From the Content pane, select the patch.
4. From the View drop-down list, select Patch Policies.
5. From the Show drop-down list, select Policies without Patch Added.
6. Select a policy. From the **Actions** menu, select **Add to Patch Policy**.
7. In the Add to Patch Policy window, click **Add**.

Removing a Patch from a Patch Policy

This action only removes a patch from a patch policy. This action does not uninstall the patch from a managed server and does not remove the patch from Opsware SAS.

To remove a patch from a patch policy, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.
2. Select a specific Windows operating system and view the list of Windows patches.
3. View the list of Windows patches.
4. From the Content pane, select a patch.
5. From the View drop-down list, select Patch Policies.
6. From the Show drop-down list, select Policies with Patch Added.

7. Select a patch. From the **Actions** menu, select **Remove from Patch Policy**.
8. In the Remove Patch from Policy window, select the policy and click **Remove**.

Attaching a Patch Policy to a Server

This action associates a patch policy with a server (or server group). You must perform this action before you reconcile a policy with a server (or server group).

To attach the policy, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patch Policies.
2. Select a specific Windows operating system and view the list of Windows patch policies.
3. From the Content pane, select a patch policy.
4. From the View drop-down list, select Servers (or Server Groups).
5. From the Show drop-down list, select Servers with Policy Not Attached (or Server Groups with Policy Not Attached).
6. From the Preview pane, select one or more servers.
7. From the **Actions** menu, select **Attach Server**.
8. Click **Attach**.

Detaching a Patch Policy from a Server

This action does not delete the patch policy and does not uninstall patches from a managed server.

To detach the policy, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patch Policies.
2. Select a specific Windows operating system and view the list of Windows patch policies.
3. From the Content pane, select a patch policy.
4. From the View drop-down list, select Servers (or Server Groups).
5. From the Show drop-down list, select Servers with Policy Attached (or Server Groups with Policy Attached).

6. From the Preview pane, select one or more servers.
7. From the **Actions** menu, select **Detach Server**.
8. Click **Detach**.

Setting a Patch Policy Exception

A patch policy exception indicates whether the patch is installed during the Reconcile Patches action. (The Install Patch and Uninstall Patch actions ignore patch policy exceptions.) A patch policy exception overrides the policy. You specify an exception for a particular patch and server (or server group), not for a patch policy.

To set a patch policy exception, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and then select All Managed Servers.
2. Select a server.
3. From the Content pane, select a patch.
4. From the View drop-down list, select Patches.
5. From the Preview pane, select a patch.
6. From the **Actions** menu, select **Set Exception**.
7. In the Set Policy Exception window, select the Exception Type:
   - **Always Install**: The patch should be installed on the server even if the patch is not in the policy.
   - **Never Install**: The patch should not be installed on the server, even if the patch is in the policy.
8. (Optional) In the Reason field, enter an explanation. This explanation is displayed when you move the cursor over the Exception column in the Preview pane display of patches with exceptions.
9. Click **OK**.

Finding an Existing Patch Policy Exception

You can search for managed servers already have patch policy exceptions attached to them, and you can search for patches that have exceptions.

To find an existing patch policy exception, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers and the select All Managed Servers.
2. From the View drop-down list, select Patches.
3. From the Content pane, select a server.
4. From the Show drop-down list, select Patches with Policies or Exceptions or Patches with Exceptions.
5. In the Exception column, move the cursor over the icon to display the reason for this exception. The following icons indicate the type of patch policy exception:

   An always install exception on a patch/server association.

   An always install exception inherited to a server from a server group/patch association.

   A never install exception on a patch/server association.

   A never install exception inherited to a server from a server group/patch association.

Copying a Patch Policy Exception

To copy an exception between servers or server groups, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.
2. Expand the Patches and select a specific Window operating system.
3. From the Content pane, select a patch.
4. From the View drop-down list, select Servers (or Server Groups).
5. From the Show drop-down list, select Servers with Exception (or Server Groups with Exception).
6. From the Preview pane, select a server. This server is the source of the copied exception.
7. From the **Actions** menu, select **Copy Exception**.
8. In the Copy Policy Exception window, select the target servers or server groups.

   These servers are the destinations of the copied exception. If this operation would result in replacing an existing exception, a message displays asking you to confirm whether this is the preferred action.

9. Click **Copy**.

Removing a Patch Policy Exception

To remove a patch policy exception, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.
2. Expand the Patches and select a specific Window operating system.
3. From the Content pane, select a patch.
4. From the View drop-down list, select Servers.
5. From the Show drop-down list, select Servers with Exception.
6. From the Preview pane, select a server.
7. From the **Actions** menu, select **Remove Exception.**

Patch Compliance

Patch Management performs conformance tests (compliance checks) against managed servers and public server groups to determine whether all patches in a policy and a policy exception were installed successfully. To enforce patch compliance, servers are scanned to determine whether they conform to their attached policies and exceptions, based on compliance levels and compliance rules. To optimize patch compliance information for your organization, you can set the patch compliance levels and edit the rules of the customized patch compliance level.

Patch Compliance Scans

A patch compliance scan compares patches that are installed on a server with patch policies and patch policy exceptions that are attached to that server. The results of this scan tell you which servers are in compliance (have all required patches installed) and which servers are out of complies (do not have all required patches installed).

A patch compliance scan occurs only when you request the scan or according to a schedule you have set up. There is no default schedule for running patch compliance scans in Opsware SAS. Opsware, Inc. recommends that you run or schedule patch compliance scans based on the dynamics of your patching environment. For example, if you updated a patch policy or installed a patch outside of (by not using) Opsware SAS, a compliance scan is required because the Opsware model has been changed and the compliance information must now be recalculated. Patch Management indicates these types of conditions as Scan Needed in the GUI. In this case, instead of waiting for the scan schedule to iterate, you can request the scan to run update your patch compliance information.

To indicate whether a server is in compliance, out of compliance, or requires a scan, Patch Management displays the following patch compliance information:

Figure 12-6: Patch Compliance Icons

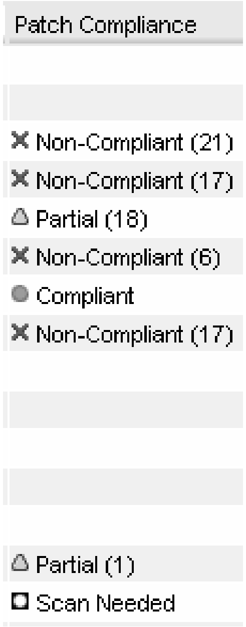

Patch Compliance

✖ Non-Compliant (21)
✖ Non-Compliant (17)
△ Partial (18)
✖ Non-Compliant (6)
● Compliant
✖ Non-Compliant (17)

△ Partial (1)
□ Scan Needed

- **Compliant**: Server is compliant for all patches. Patches in policies attached to the server and are all installed on that server.
- **Partial (#)**: Server is partially compliant for (#) patches. An exception has been set for these patches. The number in parenthesis is the total number of patches affected.
- **Non-Compliant (#)**: Server is not compliant for (#) patches. Patches in policies attached to the server are not installed on that server. The number in parenthesis is the total number of patches affected.
- **Scan Needed**: Compliance information my not be accurate. Please run a patch compliance scan.
- **Scanning**: Compliance information is currently being calculated.

Patch Compliance Levels

Patch compliance levels define your patch compliance rules. Results of a patch compliance scan can include only policies, both policies and exceptions, or your own customized level.

Patch Management supports the following compliance levels:

- **Policy Only**: Verifies whether the patches installed on a server comply with the patch policies.
- **Policy and Exception**: Verifies whether the patches installed on a server comply with the patch policies and any exceptions. The Partial (yellow) icon is displayed if the policy and exception do not agree and the exception does not have data in the Reason field.
- **Customized**: Verifies the rules that you edited for this compliance level.

Patch Compliance Rules

Patch compliance rules are the conditions that determine which compliance icons are displayed in the Managed Server window.

Patch Management supports the following compliance rules:

- **Patch Added to Policy**: The patch has been added to the patch policy.
- **Patch Installed on Server**: The patch has been installed on the managed server.
- **Exception Type**: The Exception Type can have the following values:
  - **Always Installed**: The patch should be installed on the server, even if the patch is not in the policy.
  - **Never Installed**: The patch should not be installed on the server, even if the patch is in the policy.
  - **None**: An exception has not been specified for the patch and server.
- **Exception Reason**: A description entered in the Exception Reason of the Set Policy Exception window. In the Patch Compliance Rules window, the Exception Reason can have the following values.
  - **Yes**: The Exception Reason has data.
  - **No**: The Exception Reason is empty.

  - **N/A**: An exception has not been specified for the patch and server.
- **Compliance Result**: The icon that indicates the result of the patch compliance scan. These icons are displayed in the Managed Server window.

Patch Administration for Windows

You can customize patch administration for Windows to best support your environment in the following manner:

- You can specify whether you want patches immediately available for installation by using a command-line script or the OCC Client.
- You can import the Microsoft patch database (on demand) by using a command-line script or the OCC Client.
- You can track (and import) only patches that apply to certain Microsoft products.
- You can manually launch (on demand) or schedule periodic policy compliance scans to determine the patch state of your managed servers.
- You can customize the icon display of policy compliance scan results.

Setting the Default Patch Availability

You can set the default patch availability with either the OCC Client or a command-line script. The default used by the script overrides the default set by the OCC Client. For information about the script, see “Automatically Importing Windows Patches” in the *Opsware® SAS Administration Guide*.

To set the default value for the Availability of a newly imported patch, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Opsware Administration.
2. Select Patch Configuration.
3. For the Default Availability for Imported Patches, select either Available or Limited. The default is Limited.

If the patch is Available, it can be installed on managed servers. If the patch is Limited, it has been imported into Opsware SAS and can be installed only by a patch administrator who has the required permissions. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

Importing the Microsoft Patch Database

You can import the Microsoft Patch Database by using a command-line script or the OCC Client. For information about the script, see "Automatically Importing Windows Patches" in the *Opsware® SAS Administration Guide*.

To import the database with the OCC Client, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Opsware Administration.

2. Select Patch Configuration.

3. To import the database from the Microsoft web site, click **Import from Vendor**.

   A window appears with the default URL for the location of the database on the Microsoft web site. Click **Import**. To re-import a new version of the Microsoft database that is released monthly, you must use this URL.

4. To import the database from the local file system, click **Import from File**.

   A file browser window appears. Go to the folder containing the `mssecure.cab` (MBSA 1.2) or `wsusscan.cab` (MBSA 2.0) file and click **Import**. These files must have been previously downloaded from the Microsoft web site and copied to the local file system.

To be imported, a patch must be in the Microsoft Patch database that has already been imported into the Software Repository.

Selecting Windows Products to Track for Patching

This operation limits the patches tracked by Opsware SAS to specific Windows products. After performing this operation, the next time the Microsoft Patch Database is imported, any new patches listed by Opsware SAS are limited to the products that you select. Patches that were previously listed by Opsware SAS are still tracked.

To limit the patches tracked to specific Windows operating systems, run the command-line script that automatically imports patches. For information about the script, see "Automatically Importing Windows Patches" in the *Opsware® SAS Administration Guide*.

Many duplicate QNumbers (multiple instances of the same QNumber in the same operating system) can occur if you track patches using both MBSA 2.0 and MBSA 1.2.1. To minimize the volume of duplicate QNumbers, Opsware recommends the following:

- If your servers are running Windows 2000 Service Pack 3 (or higher) operating systems or Windows 2003 operating systems, use MBSA 2.0.
- If your servers are running Windows NT 4.0 operating systems or Windows 2000 RTM (no service pack), Service Pack 1, or Service Pack 2 operating system, use MBSA 1.2.1.

To select the Windows products to track for patching, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Opsware Administration.
2. Select Patch Configuration.
3. Click **Edit**.
4. Select the products whose patches you want to track (from either the Windows MBSA 2.0 tab or the Windows MBSA 1.2.1 tab) and then click **Select**.

   If you select the MBSA 1.2.1 tab and this is a fresh install, the list of Products in Patch Database is empty. Click **Edit** to select the products you want to track patches for.

Scheduling a Patch Policy Compliance Scan

To schedule or modify a patch compliance scan, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Opsware Administration.
2. Select Patch Compliance Rules.
3. In the Patch Policy Compliance Scan Schedule section, click **Edit**.
4. In the Schedule Compliance Scan window, select Enable Compliance Scan.
5. In the Schedule drop-down list, select the frequency of the scans.

   If you select Custom, specify the crontab string with the following values:

   Minute (0-59)

   Hour (0-23)

Day of the month (1-31)

Month of the year (1-12)

Day of the week (0-6 with 0=Sunday)

Any of these fields can contain an asterisk to indicate all possible values. For example, the following crontab string runs the job at midnight every weekday:

0 0 * * 1-5

The crontab string can also handle serial (1,2,3,4) as well as range (1-5) values. For more information, consult the crontab man pages on a Unix computer.

6 In the Start Time field, specify the time you want the job to begin.

7 In the Time Zone drop-down list, select a default time zone for the job execution time or accept the default time zone. The default time shown converts the scheduled time to the time zone set in your user preferences. If you do not set a preferred time zone, the time zone is derived from the Opsware SAS core server, which is typically UTC.

8 In the Day(s) to Run field, select one or more days of the week that you want the scan to run.

9 Click **OK**.

Setting the Patch Policy Compliance Level

The patch policy compliance level defines your patch compliance rules. To view these rules or to set the patch policy compliance level, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Opsware Administration.

2 Select Patch Compliance Rules.

3 Select one of the following compliance levels: Policy and Exception, Policy Only, or Customized.

Editing the Customized Patch Policy Compliance Level

Of the three compliance levels, only the Customized level can be edited. To edit this level, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Opsware Administration.

2 Select Patch Compliance Rules.

3 From the Compliance Level, select Customized.

4. In the Patch Policy Compliance Setting section, click **Edit**.
5. Select the Compliance Level icons that you want to change in the Compliance Result column: Non-Compliant, Compliant, No Indicator, or Partial.
6. Click **Apply** and then click **Close**.

Patch Installation

Patch Management provides the following two phases in the patch installation process:

- **Download Phase**: This is when the patch is downloaded from Opsware SAS to the managed server. This phase is commonly referred to as the staging phase.
- **Installation Phase**: This is when the patch is installed on the managed server. This phase is commonly referred to as the deployment phase.

You can specify whether you want the installation to occur immediately after the patch is downloaded (staged) or you can schedule the installation to occur at a later date and time. Patch Management also supports the need for best-effort installations of multiple patches by allowing you to specify that the patch installation process will continue even when an error occurs with one of the patches.

Patch Management displays the name of the command (.exe file and any predefined command-line arguments) that the Opsware Agent runs on the managed server to install the patch. You can override the default command-line arguments that you want to perform the installation.

To help you optimally manage the conditions under which Windows patches are installed, Patch Management allows you to manage server reboot options, and pre and post installation scripts, simulate (preview) a patch installation, and set up email notifications to alert you about the status of the installation process. The Install Patch task window guides you through setting up these conditions.

Figure 12-7: Install Patch Task Window

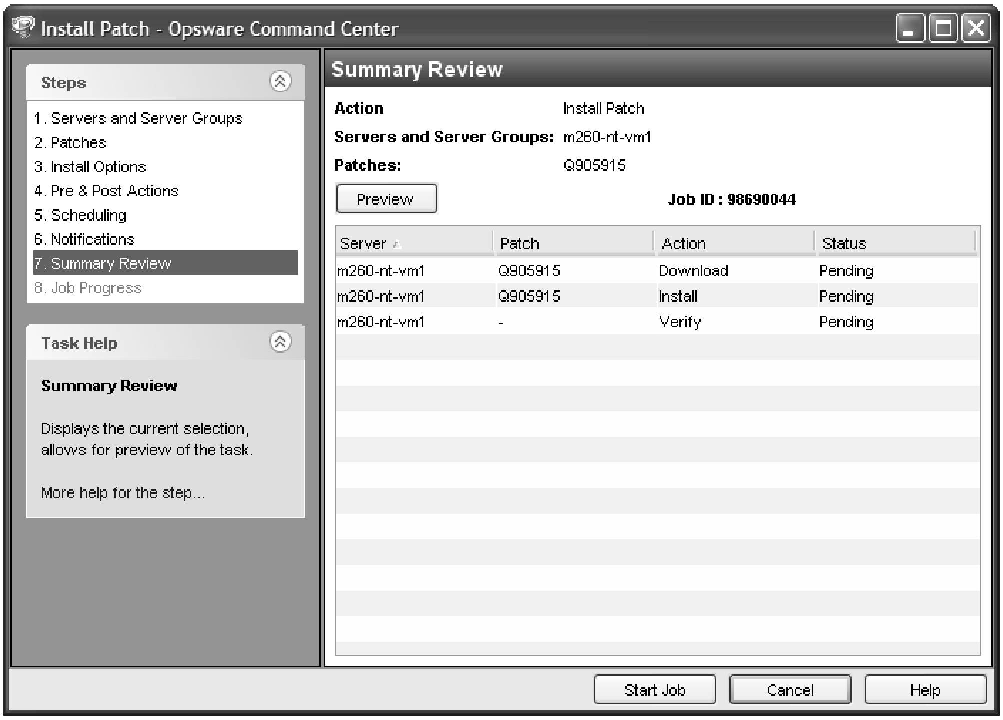

Summary Review

Displays the current selection, allows for preview of the task.

More help for the step...

Action Install Patch
Servers and Server Groups: m260-nt-vm1
Patches: Q905915

Job ID : 98690044

Server	Patch	Action	Status
m260-nt-vm1	Q905915	Download	Pending
m260-nt-vm1	Q905915	Install	Pending
m260-nt-vm1	-	Verify	Pending

Installation Flags

You can specify installation flags that are applied whenever a Windows patch is installed. However, Opsware SAS also uses default installation flags and requires that patches are installed with these flags. You must therefore be certain that you do not specify any installation flags that override or contradict the default flags passed in by Opsware SAS. For information about how to specify commands and flags, see "Setting Install Options" on page 503.

Some Windows hotfixes do not support the -z flag, some do not support the -q flag, and some do not support either. In such cases, you must use a special expression: /-z or /-q or /-z -q respectively, to prevent the Patch Management feature from passing in the -z or

-q or -z -q flag. By default, Opsware SAS adds /z /q to the command line arguments when installing patches. To override this, specify /-z /-q. For example, if you prefer to not suppress the reboot, specify /-z.

The following table lists the default installation flags that Opsware SAS uses.

Table 12-2: Default Installation Flags

WINDOWS PATCH TYPES	FLAGS
Windows Hotfix	`-q -z`
Windows Security Rollup Package (treated identically to a Hotfix by the Patch Management feature)	`-q -z`
Windows OS Service Pack	`-u -n -o -q -z`

Service Packs, Update Rollups, and Hotfixes

When you try to install a Service Pack, Update Rollup, or a Hotfix, there is a known delay and a potential for the installation job to be aborted. For example, when the Opsware Agent is used to install a Service Pack, the OCC Client displays a dialog asking you to confirm. If you do not click OK, the Agent times out and the install job is aborted. For Service Packs and Update Rollups, the Agent will time out if one hour has lapsed and you have not clicked OK in the confirmation dialog. For Hotfixes, the Agent will time out if five minutes have lapsed and you have not clicked OK in the confirmation dialog.

Installing a Patch

Before a patch can be installed on a managed server, it must be imported into Opsware SAS and its status must be Available. Administrators who have the required permissions can install patches that are marked Limited.

You must have a set of permissions to manage patches. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Configuration Guide* for more information.

You can perform the installation by explicitly selecting patches and servers, and you can install a patch even if the patch policy exception is Never Install.

To install a patch on a managed server, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.
2. Expand the Patches and select a specific Window operating system.
3. From the Content pane, select a patch.
4. From the View drop-down list, select Servers (or Server Groups).
5. From the Show drop-down list, select Servers without Patch Installed (or Server Groups without Patch Installed).
6. From the Preview pane, select one or more servers.
7. From the **Actions** menu, select **Install Patch**.

   The first step of the Install Patch task appears: Servers and Server Groups.
   For instructions on each step, see the following sections:

   – Setting Install Options

   – Setting Reboot Options for a Patch Install

   – Specifying Pre and Post Install Scripts for a Patch Install

   – Scheduling a Patch Installation

   – Setting Up Email Notifications for a Patch Install

   – Previewing a Patch Installation

   – Viewing Job Progress of a Patch Install

   After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.
8. When you are ready to launch the installation job, click **Start Job**.

   After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If the Install Patch task window remains open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select Refresh from the View menu to update information in the Patch Preview pane.

For another method of installing a patch, see “Reconciling Patches” on page 478.

Setting Install Options

You can specify the following types of patch installation options:

- Perform the patch installation immediately after the patch is downloaded or at a later date and time.
- Do not interrupt the patch installation process even when an error occurs with one of the patches.
- Use different command-line options to perform the installation.

To set these options, perform the following steps:

1. From the Install Patch task window, click **Next** to advance to the Install Options step.
2. Select one of the following Staged Install Options:
    - **Continuous**: This allows you to run all phases as an uninterrupted operation.
    - **Staged**: This allows you to schedule the download and install to run separately.
3. Select the Error Options check box if you want the patch installation process to continue even when an error occurs with one of the patches. As a default, this check box is not selected.
4. In the Install Command text box, enter command-line arguments for the command (.exe file) that is displayed. By default, Opsware SAS adds /z /q. If you want to override these install flags, enter /-z /-q in the text box.
5. Click **Next** to go to the next step or click **Cancel** to close the Install Patch task window.

Setting Reboot Options for a Patch Install

To minimize the downtime that server reboots can cause, you can control when servers will and will not be rebooted. You can adopt the vendor's reboot assignments, reboot a server each time a patch is installed on it, completely suppress all server reboots, or postpone reboots until all patches have been installed.

When you are selecting reboot options in the Install Patch task window, Opsware, Inc. recommends that you use Microsoft's reboot recommendations, which is the "Reboot servers as specified by patch properties" option in the task window. If it is not possible to use the Microsoft reboot setting, select the single reboot option, which is the "Do not reboot servers until all patches are installed" option in the task window. Failure to do this can result in MBSA incorrectly reporting which patches are installed on the server until the next reboot occurs (outside of Opsware control).

The following options determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Install Patch task window; they do not change the Reboot Required option, which is on the Install Parameters tab of the patch properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties**: By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install**: Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots**: Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Do not reboot servers until all patches are installed**: If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

1. From the Install Patch task window, click **Next** to advance to the Pre & Post Actions step.
2. Select one of the Rebooting Options.
3. Click **Next** to go to the next step or click **Cancel** to close the Install Patch task window.

Specifying Pre and Post Install Scripts for a Patch Install

For each patch, you can specify a command or script to run before installation or after installation. A pre-install script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-install script fails, the patch would not be installed. A pre-install script could also be used to shut down a service or application before it is patched. A post-install script could be used to perform a certain cleanup process on the managed server.

You can also specify the following types of scripts to run on the managed server before or after an installation or download phase:

- **Pre-Download**: A script that runs before patches are downloaded from Opsware SAS to the managed server. This is available only if you select Staged in the Install Options step.
- **Post-Download**: A script that runs after patches are downloaded from Opsware SAS to the managed server and before the patch is installed. This is available only if you select Staged in the Install Options step.
- **Pre-Install**: A script that runs before patches are installed on the managed server.
- **Post-Instal**l: A script that runs after patches are installed on the managed server.

To specify a pre-install script, perform the following steps:

1. From the Install Patch task window, click **Next** to advance to the Pre & Post Actions step.
2. Select the Pre-Install tab. You may specify different scripts and options on each of the tabs.
3. Select Enable Script. This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
4. Select either Saved Script or Ad-Hoc Script.

A Saved Script has been previously stored in Opsware SAS with the Opsware Command Center. To specify the script, click **Select**.

An Ad-Hoc script runs only for this operation and is not saved in Opsware SAS. Select the Type, such as .BAT. In the Script box, enter the contents of the script, including the drive letter of where the script is located, such as echo dir>> C:\temp\preinstall1.log. If you do not enter a drive letter, the default is %SYSTEMDRIVE%, which is where the system folder of Windows is installed.

5. If the script requires command-line flags, enter the flags in the Command text box.
6. Specify the information in the User section. If you choose a system other than Local, enter the User Name, Password, and Domain. The script will be run by this user on the managed server.
7. To stop the installation if the script returns an error, select the Error check box.
8. Click **Next** to go to the next step or click **Cancel** to close the Install Patch task window.

Scheduling a Patch Installation

Since the two phases of patching can be decoupled, you can schedule when you want patches installed (deployed) to occur independently of when patches are downloaded (staged).

To schedule a patch installation, perform the following steps:

1. From the Install Patch task window, click **Next** to advance to the Scheduling step.

   By default, the Scheduling step displays only the scheduling options for the install phase. If you selected Staged in the Install Options step, the scheduling options for the download phase will also be displayed.

2. Select one of the following Install Phase options:

   - **Run Task Immediately**: This enables the system to perform a preview analysis in the Summary Review step. The scheduling option for the download phase is **Run Immediately Following Download**.
   - **Run Task At**: This enables you to specify a later date and time that you want the install or download performed.

3. Click **Next** to go to the next step or click **Cancel** to close the Install Patch task window.

A scheduled patch installation can be cancelled (prior to its execution), even if the patch download has already completed.

Setting Up Email Notifications for a Patch Install

You can set up email notifications to alert users when the download and installation operations complete successfully or with errors.

To set up email notifications, perform the following steps:

1. From the Install Patch task window, click **Next** to advance to the Notifications step.
2. In the Notify text box, the default email address that displays is derived from the Profile of your Opsware user. To add email addresses, click **Add Notification** and enter additional email addresses.
3. Click **Next** to go to the next step or click **Cancel** to close the Install Patch task window.

If you previously selected Staged in the Install Options step, the Notifications pane displays notification options for both the download and install phases.

Previewing a Patch Installation

The installation preview process provides an up-to-date report about the patch state of servers. The installation preview is an optional step that lets you see what patches will be installed on managed servers and what type of server reboots are required. This preview process verifies whether the servers you selected for the patch installation already have that patch installed (based on the MBSA). In some cases, a server could already have the patch installed if a system administrator had manually installed it, which means that Patch Management does not know about it.

The preview process also reports on dependency and supersedence information, such as patches that require certain Windows products, and patches that obsolete other patches or are obsoleted by other patches. If a dependency is not met, Patch Management will display an error message indicating this condition. For example, if a managed server is running Windows 2000 Service Pack 3 (or higher) or Windows 2003, and an Opsware SAS 5.5 Agent, Patch Management will report that a dependency has not been fulfilled. If you try to install a patch for Service Pack 4 and your server is using Service Pack 3, the reconcile preview will display a "Will Not Install" error message to indicate this discrepancy. The Install Patch task window allows superseded patches to be installed.

The installation preview does not report on the behavior of the server as though the patches have been applied.

To preview a patch installation, perform the following steps:

1. From the Install Patch task window, click **Next** to advance to the Summary Review step.
2. Verify the information displayed for the Servers, Server Groups, and Patches at the top of the window.
3. (Optional) Click **Preview** to see the separate actions that will be performed when the patch is installed. To view the details of a previewed action, select a row in the table.
4. Click **Start Job** to launch the installation job or click **Cancel** to close the Install Patch task window without launching the install.

   If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected Run Task At, the job will be launched at the specified time and date.

Viewing Job Progress of a Patch Install

You can review progress information about a patch installation (job), such as whether actions have completed or failed.

To display job progress information, perform the following steps:

1. From the Install Patch task window, click **Next** to advance to the Job Progress step. This will start the install job.

   The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:

   – **Analyze**: Opsware SAS examines the patches needed for the install, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.

   – **Download**: The patch is downloaded from Opsware SAS to the managed server.

   – **Install**: After it is downloaded, the patch is installed.

- **Final Reboot**: If this action is specified in the Pre & Post Actions step, the server is rebooted.
- **Pre/Post Install/Download Script**: If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstall.
- **Install & Reboot**: When a patch will be installed is also when the server will be rebooted.
- **Verify**: Installed patches will be included in the software registration.

2. To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select Jobs and Sessions to review detailed information about the job. See “Browsing Job Logs” on page 179.
3. Click **Stop Job** to prevent the job from running or click **Close** to close the Install Patch task window.

Patch Uninstallation

Patch Management provides granular control over how and under what conditions Windows patches are uninstalled (removed) from managed servers. To minimize problems, you can only uninstall one patch at a time. You cannot use Opsware SAS to uninstall a patch that was not installed by using the Patch Management feature.

To help you optimally manage these conditions, Patch Management allows you to do the following:

- Manage server reboot options, and pre and post installation scripts
- Simulate (preview) a patch uninstallation
- Set up email notifications to alert you about the status of the uninstallation process

The Uninstall Patch task window guides you through setting up these conditions.

Figure 12-8: Uninstall Patch Task Window

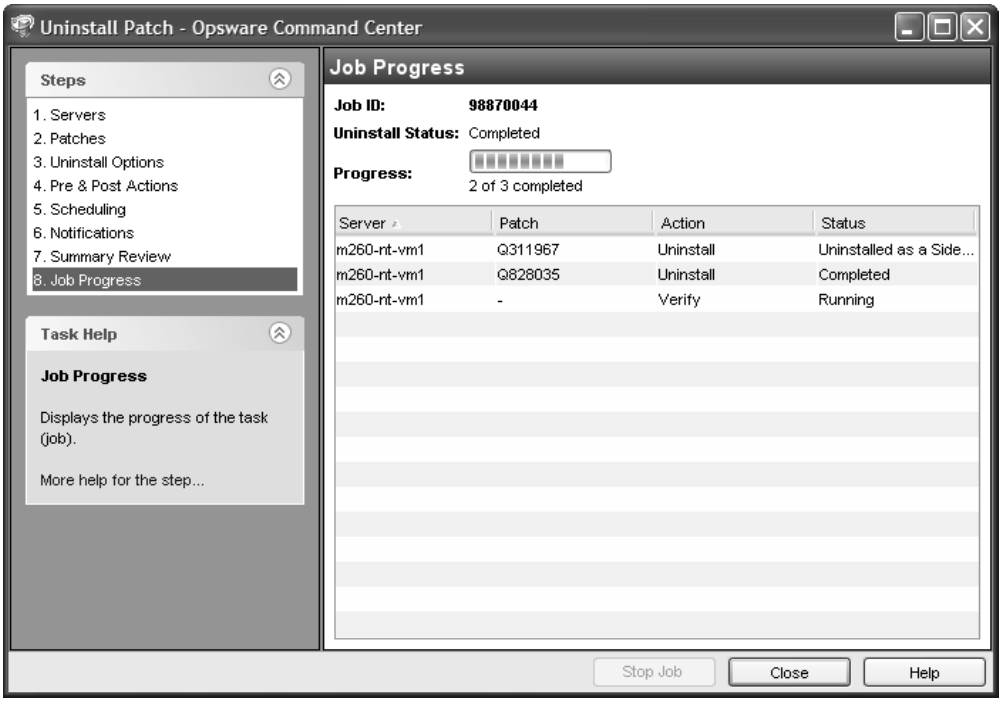

Job Progress

Displays the progress of the task (job).

More help for the step...

Job Progress

Job ID: 98870044

Uninstall Status: Completed

Progress: 2 of 3 completed

Server	Patch	Action	Status
m260-nt-vm1	Q311967	Uninstall	Uninstalled as a Side...
m260-nt-vm1	Q828035	Uninstall	Completed
m260-nt-vm1	-	Verify	Running

Uninstallation Flags

You can specify uninstallation flags that are applied whenever a Windows patch is uninstalled. However, Opsware SAS also uses default uninstallation flags and requires that patches are uninstalled with these flags. You must therefore be certain that you do not specify any uninstallation flags that override or contradict the default flags passed in by Opsware SAS. For information about how to specify commands and flags, see “If the Uninstall Patch task window remains open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select Refresh from the View menu to update information in the Patch Preview pane.” on page 512.

Some Windows hotfixes do not support the -z flag, some do not support the -q flag, and some do not support either. In such cases, you must use a special expression: /-z or /-q or /-z -q respectively, to prevent the Patch Management feature from passing in the -z or -q or -z -q flag. By default, Opsware SAS adds /z /q to the command line arguments when uninstalling patches. To override this, specify /-z /-q. For example, if you prefer to not suppress the reboot, specify /-z.

The following table lists the default uninstallation flags that Opsware SAS uses.

Table 12-3: Default Uninstallation Flags

WINDOWS PATCH TYPES	FLAGS
Windows Hotfix	`-q -z`
Security Rollup Package	`-q -z`
Windows OS Service Pack	Not uninstallable

Uninstalling a Patch

To remove a patch from a managed server, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library and then select Patches.
2. Expand the Patches and select a specific Window operating system.
3. From the Content pane, select a patch.
4. From the View drop-down list, select Servers.
5. From the Show drop-down list, select Servers with Patch Installed.
6. From the Preview pane, select one or more servers.
7. From the **Actions** menu, select **Uninstall Patch**.

   The first step of the Uninstall Patch task appears: Servers.
   For instructions on each step, see the following sections:

   – Setting Reboot Options for a Patch Uninstall

   – Specifying Pre and Post Install Scripts for a Patch Uninstall

   – Scheduling a Patch Uninstallation

   – Setting Up Email Notifications for a Patch Uninstall

   – Viewing Job Progress of a Patch Uninstall

After you have completed a step, click **Next** to advance to the next step. Before you click **Start Job**, you can return to a completed step to make changes by clicking on it in the list of steps.

8 When you are ready to launch the uninstall job, click **Start Job**.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If the Uninstall Patch task window remains open until the job completes, Patch Management updates the Patch Compliance column in the All Managed Servers window with the revised compliance count (in parenthesis) for affected servers. Press F5 or select Refresh from the View menu to update information in the Patch Preview pane.

Setting Uninstall Options

You can specify the following types of patch uninstallation options:

- Do not interrupt the patch uninstallation process even when an error occurs with one of the patches.
- Use different command-line options to perform the uninstall.

To set these options, perform the following steps:

1 From the Uninstall Patch task window, click **Next** to advance to the Uninstall Options step.

2 Select the Error Options check box if you want the patch installation process to continue even when an error occurs with one of the patches. As a default, this check box is not selected.

3 In the Uninstall Command text box, enter command-line arguments for the command (.exe file) that is displayed. By default, Opsware SAS adds /z /q. If you want to override these uninstall flags, enter /-z /-q in the text box.

4 Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch task window.

Setting Reboot Options for a Patch Uninstall

To minimize the downtime that server reboots can cause, you can control when servers will and will not be rebooted. You can adopt the vendor's reboot assignments, reboot a server each time a patch is removed from it, completely suppress all server reboots, or postpone reboots until all patches have been uninstalled.

When you are selecting reboot options in the Uninstall Patch task window, Opsware, Inc. recommends that you use Microsoft's reboot recommendations, which is the "Reboot servers as specified by patch properties" option in the task window. If it is not possible to use the Microsoft reboot setting, select the single reboot option, which is the "Do not reboot servers until all patches are installed" option in the task window. Failure to do this can result in MBSA incorrectly reporting which patches are installed on the server until the next reboot occurs (outside of Opsware control).

The following options determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Uninstall Patch task window; they do not change the Reboot Required option, which is on the Uninstall Parameters tab of the patch properties window. Except for the first option, the following options override the Reboot Required option.

- **Reboot servers as specified by patch properties**: By default, the decision to reboot depends on the Reboot Required option of the patch properties.
- **Reboot servers after each patch install**: Even if the Reboot Required option of the patch properties is not set, reboot the server. If multiple patches are installed, the server reboots multiple times.
- **Suppress all server reboots**: Even if the Reboot Required option of the patch properties is set, do not reboot the server. (Because of vendor settings, some patches ignore the suppress option and force a reboot.)
- **Do not reboot servers until all patches are installed**: If the Reboot Required option is set for some selected patches but not for others, the server is rebooted one time after all patches are installed. If the Reboot Required option is not set for any of the selected patches, the server is not rebooted.

To set reboot options, perform the following steps:

1. From the Uninstall Patch task window, click **Next** to advance to the Pre & Post Actions step.
2. Select one of the Rebooting Options.
3. Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch task window.

Specifying Pre and Post Install Scripts for a Patch Uninstall

For each patch, you can specify a command or script to run before uninstallation or after uninstallation. A pre-uninstall script, for example, could check a certain condition on the managed server. If the condition is not met or if the pre-uninstall script fails, the patch would not be removed from a server. A pre-uninstall script could also be used to shut down a service or application before it is removed from a server. A post-uninstall script could be used to perform a certain cleanup process on the managed server.

You can specify the following types of scripts to run on the managed server before or after a patch uninstall:

- **Pre-Uninstall**: A script that runs before the patch is removed from a managed server.
- **Post-Uninstall**: A script that runs after the patch is removed from a managed server.

To specify a script, perform the following steps:

1. From the Uninstall Patch task window, click **Next** to advance to the Pre & Post Actions step.
2. Select the Pre-Uninstall or Post-Uninstall tab.

   You may specify different scripts and options on each of the tabs.
3. Select Enable Script.

   This option enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
4. Select either Saved Script or Ad-Hoc Script.

   A Saved Script has been previously stored in Opsware SAS with the Opsware Command Center. To specify the script, click **Select**.

   An Ad-Hoc script runs only for this operation and is not saved in Opsware SAS. Select the Type, such as .BAT. In the Script box, enter the contents of the script, including the drive letter of where the script is located, such as echo dir>> C:\temp\preinstall1.log. If you do not enter a drive letter, the default is %SYSTEMDRIVE%, which is where the system folder of Windows is installed.
5. If the script requires command-line flags, enter the flags in Commands.
6. Specify the information in the User section. The script will be run by this user on the managed server.
7. To stop the uninstallation if the script returns an error, select Error.

Scheduling a Patch Uninstallation

You can schedule that a patch will be removed from a server immediately, or at a later date and time.

To schedule a patch uninstall, perform the following steps:

1. From the Uninstall Patch task window, click **Next** to advance to the Scheduling step.
2. Select one of the following Install Phase options:
   - **Run Task Immediately**: This enables you to perform the uninstall in the Summary Review step.
   - **Run Task At**: This enables you to specify a later date and time that you want the uninstall performed.
3. Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch task window.

Setting Up Email Notifications for a Patch Uninstall

You can set up email notifications to alert users when the patch uninstall operation completes successfully or with errors.

To set up email notifications, perform the following steps:

1. From the Uninstall Patch task window, click **Next** to advance to the Notifications step.
2. In the **Notify** text box, the default email address that displays is derived from the Profile of your Opsware user. To add email addresses, click **Add Notification** and enter additional email addresses.
3. Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch task window.

Previewing a Patch Uninstallation

The uninstallation preview process provides an up-to-date report about the patch state of servers. The uninstall preview is an optional step that lets you see what patches will be removed from managed servers. This preview process verifies whether the servers you selected for the patch uninstall have that patch installed (based on the MBSA).

The uninstall preview process does not report or simulate the behavior of a system with patches removed from the server.

To preview a patch uninstall, perform the following steps:

1. From the Uninstall Patch task window, click **Next** to advance to the Summary Review step.
2. Verify the information displayed for the Servers, Server Groups, and Patches at the top of the window.
3. (Optional) Click **Preview** to see the separate actions that will be performed when the patch is uninstalled. To view the details of a previewed action, select a row in the table.
4. Click **Start Job** to launch the job or click **Cancel** to close the Uninstall Patch task window without launching the uninstall.

   If you selected Run Task Immediately in the Scheduling step, the job begins now. If you selected Run Task At, the job will be launched at the specified time and date.

Viewing Job Progress of a Patch Uninstall

You can review progress information about a patch uninstallation (job), such as whether actions have completed or failed.

To display job progress information, perform the following steps:

1. From the Uninstall Patch task window, click **Next** to advance to the Job Progress step. The Progress bar and text indicate how many of the actions listed in the table have been completed. For each server, the following actions can be performed:
   - **Analyze**: Opsware SAS examines the patches needed for the uninstall, checks the managed servers for the most recent patches installed, and determines other actions it must perform.
   - **Uninstall**: The patch is uninstalled.
   - **Final Reboot**: If this action is specified in the Pre & Post Actions step, the server is rebooted.
   - **Pre/Post Uninstall Script**: If this action is specified in the Pre & Post Actions step, a script is run before or after the uninstall.
   - **Uninstall & Reboot**: When a patch will be installed is also when the server will be

rebooted.

- **Verify**: Installed patches will be included in the software registration.

2. To view additional details about a specific action, select the row in the table to display the start and completion times of the job. From the Navigation pane, select Jobs and Sessions to review detailed information about the job. See “Browsing Job Logs” on page 179.

3. Click **Stop Job** to prevent the job from running or click **Close** to close the Uninstall Patch task window.

Chapter 13: Patch Management for Unix

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Patch Management for Unix
- Patch Management Roles for Unix
- Patch Management for Specific Unix Operating Systems
- Patch Uploads for Unix
- Patch Administration for Unix
- Patch Installation and Uninstallation for Unix

Overview of Patch Management for Unix

Opsware SAS automates the key aspects of patch management, while offering a fine degree of control over how and under what conditions patches are installed.

Because patches are often released to address grave security threats, an organization needs to be able to roll out patches quickly, before systems become compromised. At the same time, however, patches can cause serious problems, from performance degradation to server failures.

The Patch Management feature allows you to react quickly to newly discovered threats, but it also provides support for strict testing and standardization of patch installation. And, if patches later cause problems even after being tested and approved, the Patch Management feature also allows you to uninstall the patches in a safe and standardized way.

Patch management is a fully integrated component of Opsware SAS. It leverages the Opsware SAS server automation features. Opsware SAS, for example, maintains a central database (called the Model Repository) that has detailed information about every server under management, the patches and software installed on the servers, and the patches and software available for installation. You can use this data to determine the severity of your exposure to a newly discovered threat, and to help you assess the benefits of rolling out a patch versus the costs in downtime and testing requirements.

By automating the patching procedure, the Patch Management feature can reduce the amount of downtime required for patching. Opsware SAS also allows you to schedule patch activity, so that patching occurs during off-peak hours.

After the patch is integrated into your environment, you can make it part of your standard builds with Opsware templates.

Features of Patch Management for Unix

Opsware SAS automates patch management by providing the following features:

- A central repository where patches are stored and organized in their formats
- A database that includes information on every patch that has been applied
- Customized scripts that can be run before and after a patch is installed
- Advanced search abilities identify servers that require patching
- Auditing abilities so that security personnel can track the deployment of important patches

Opsware SAS Integration

When a server is brought under management by Opsware SAS, the Opsware Agent installed on the server registers the server's hardware and software configuration with Opsware SAS. (The Opsware Agent repeats this registration every 24 hours.) This information, which includes data about the exact OS version, hardware type, installed software and patches, is immediately recorded in the Model Repository. Also, when you first provision a server with Opsware SAS, the same data is immediately recorded.

When a new patch is issued, you can use the Opsware Command Center to immediately identify which servers require patching. Opsware SAS provides a Software Repository where you upload patches and other software. Users access this software from the Opsware Command Center to install patches on the appropriate servers.

After a server is brought under management, you should install all patches by using the Patch Management feature. If you install a patch manually, Opsware SAS does not have data about that patch until the next software registration. If you install a patch manually, it can take as long as 24 hours until the data about that server in the Model Repository is up-to-date.

Whenever you install or uninstall software or patches with Opsware SAS, however, the Opsware Agent immediately updates the information about the server in the Model Repository.

Support for Unix Patch Testing and Installation Standardization

Opsware SAS offers features to minimize the risk of rolling out patches. First, when a patch is uploaded into Opsware SAS, its status is marked as untested and only administrators with special privileges can install it.

The patch administrator then defines patch installation and uninstallation options and tests the patch. Only after the patch is tested and the patch administrator marks it as available for use can other administrators install the patch.

The Patch Management feature allows you to standardize the way that patches are installed and uninstalled, thereby preventing ad-hoc installation procedures. Patch administrators standardize patch installation by providing pre-install and post-install scripts, install and uninstall flags, and instructions on when to reboot and how to handle error codes from the pre-install and post-install scripts.

Patch Management Roles for Unix

Opsware SAS provides support for rigorous change management by assigning the functions of patch management to two different types of administrators:

- The patch administrator (often referred to as the security administrator), who has the authority to upload and test, and edit patch options
- The system administrator, who applies the patches (that have been approved for use) uniformly and automatically according to the options that the patch administrator specifies

Only the patch administrator should have the Patches permission, which gives access to advanced features not available through the Patch Management Wizards. Both administrators must have permissions for the Patch Management Wizards. To obtain these permissions, contact your Opsware administrator. For more information, see the Permissions Reference appendix in the *Opsware® SAS Configuration Guide*.

Patch Administrator

In most organizations, patch administrators are responsible for reviewing the latest security threats and the patches that vendors have released to address these problems. The patch administrators are generally experts in the operating systems and applications that they manage, and are able to assess the necessity of applying patches issued by vendors. They are able to diagnose common problems that arise after patches are installed, allowing them to thoroughly test the patch application process.

In Opsware SAS, patch administrators are granted specific permissions that allow them to upload patches into Opsware SAS, test the patches, and then mark them as available for use. Basic users can upload patches, but they cannot install them or mark them as available. Patch administrators are also able to edit patch options (such as installation scripts) through patch management. Other types of users are not allowed to upload or edit patches.

Typically, the patch administrator uploads patches and then tests them on non-production reference hardware. After testing the patches and determining that the patches are safe to apply to production systems, they mark the patches available in the Opsware Command Center, and then advise the system administrators that they must apply the approved patches.

System Administrator

System administrators are responsible for the day-to-day maintenance of the servers in a deployment. These users are not required to have the same level of expertise in low-level system details as the patch administrator.

Because the patch administrator has set up the patch installation, the system administrators can apply the patches to a large number of servers with a few mouse clicks. They are responsible for searching for the servers that require the approved patch, running the Patch Installation Wizard, and verifying that the patches are installed successfully.

The system administrator cannot use the Patches link on the Opsware Command Center navigation panel. (This link requires the Patches permission.) The system administrator can upload patches through the Upload Patch Wizard, but cannot install a patch until the patch administrator has marked it as available. The system administrator can also uninstall patches.

Patch Management for Specific Unix Operating Systems

The types of patches and their underlying technologies can vary according to the vendor of the operating system. This section discusses the vendor-specific details for Unix patch management in Opsware SAS.

Supported Unix Versions and Patch Types

The Patch Management feature supports all of the operating system versions that Opsware SAS supports, except for Linux.

Linux does not support patches in the ordinary sense. The packages are not patchable. Instead, new versions of the RPM are delivered. Linux systems that Opsware SAS manages are therefore not viewable through the Patch Management feature interfaces. New Linux packages and updates should be managed and applied though the software page.

The following table shows the Unix versions and the patch types that the Patch Management feature supports.

Table 13-1: Supported Unix Versions and Patch Types

UNIX VERSIONS	PATCH TYPES
AIX 4.3	AIX Update Fileset APARs
AIX 5.1	AIX Update Fileset APARs
AIX 5.3	AIX Update Fileset APARs
HP-UX 11.00	HP-UX Patch Filseset HP-UX Patch Products
Solaris 6	Solaris Patch Solaris Patch Cluster
Solaris 7	Solaris Patch Solaris Patch Cluster

Table 13-1: Supported Unix Versions and Patch Types

UNIX VERSIONS	PATCH TYPES
Solaris 8	Solaris Patch Solaris Patch Cluster
Solaris 9	Solaris Patch Solaris Patch Cluster
Solaris 10	Solaris Patch Solaris Patch Cluster

Underlying Technologies for Patch Management on Unix

Behind the scenes, the Patch Management feature uses utilities and technologies that are specific for a particular operating system. Although the utilities vary, Opsware SAS enables you to perform patch management through a single interface, without having to worry about invoking a number of different patching utilities.

Opsware SAS models the way it treats patches on the way the underlying utility treats a patch. For example, if the Solaris patchadd utility is not able to install one patch contained in a patch cluster, the Solaris utility continues to install the remaining patches in the patch cluster. Opsware SAS respects this behavior and allows that patch installation operation to continue. Any patches that are not installed are reported at the end of the installation operation.

The following table shows the patch management and installation tools that are used for each of the supported Unix systems.

Table 13-2: Supporting Technologies for Patch Management on Unix

SOLARIS	AIX	HU-UX
Patchadd installs Solaris patches	Installp installs and uninstalls filesets	Swlist lists patch products, files, products, and filesets
Patchrm uninstalls Solaris patches	Lslpp lists installed LPPs	Swinstall installs a depot

Table 13-2: Supporting Technologies for Patch Management on Unix

SOLARIS	AIX	HU-UX
Showrev lists installed Solaris patches	Instfix lists installed APARs	Swremove removes a depot
Pkgadd installs Solaris packages		
Pkginfo lists installed Solaris packages		

AIX Patches

AIX periodically releases Authorized Program Analysis Reports (APARs), which specify what update filesets (contained in LPPs) are necessary to fix an identified problem. An APAR only specifies the minimum version of an update fileset required to fix a problem; an APAR can therefore be satisfied with later versions of the same filesets. To maintain compatibility, however, Opsware SAS always adopts the fileset with the lowest version number that meets the minimum version that APAR specifies. If a later version of the update fileset is uploaded, Opsware SAS still associates the earlier version of the fileset with the APAR.

When uploading an LPP, Opsware SAS recognizes which APARs the filesets contained in the LPP belong to. An entry is created for the APAR in the Patch Management feature when the first fileset associated with an APAR is uploaded. (In some cases, a fileset is associated with more than one APAR. An entry is created for each APAR the fileset is associated with, if the entry does not already exist.)

If you want to be able to install all LPPs that APAR specifies, you must make certain to upload all of the specified LPPs into the Patch Management feature, either through the Upload Patch Wizard or through the OCLI.

If you do not upload all of the LPPs that APAR specifies, it is still possible for the system administrator to browse for an APAR and install the partial set of LPPs that are uploaded. In such cases, the administrator receives a warning that the filesets for the APAR are not all installed.

The Patch Administrator must first upload and test an LPP before it is generally available in Opsware SAS. The new fileset is integrated into the APAR only after the LPP is tested and approved. Even though the APAR is updated automatically, you still maintain control over the exact filesets that are allowed to be installed on your managed servers.

APAR update filesets cannot be installed on a server if the server does not already have the base filesets for which the update filesets are intended.

If, however, a server has a partial set of the base filesets, the APAR can be applied and only the applicable filesets for the base filesets are installed. For example, if an APAR specifies four update filesets to update four base filesets, and you attempt to apply the APAR to a server that has only three of the base filesets, three of the four update filesets from the APAR are installed.

When installing an AIX Update fileset, the Patch Management feature normally applies the fileset, which allows it to be rejected (uninstalled.) If you want to commit the fileset instead (so that it cannot be removed), use the -c option here.

Solaris Patches

A Solaris patch cluster contains a set of selected patches for a specific Solaris release level. Ordinarily, after a patch cluster is installed, it is not possible to search for a particular patch cluster. The patches do not contain any metadata that relate them to the patch cluster in which they were originally bundled. You can only search for the individual patches.

If you install a Solaris patch cluster by using the Patch Management feature, however, Opsware SAS keeps track of the patch cluster in the Model Repository. You can therefore search for a patch cluster to determine if a full patch cluster is installed. You can also uninstall the patch cluster if you installed it with the Patch Management feature.

HP-UX Patches

HP-UX patches are delivered exclusively as depots, which are patch products that contain patch filesets. The depot is uploaded directly into Opsware SAS by using the Patch Management feature.

If a depot is already uploaded and attached to a node, it cannot be uploaded by using the Patch Management feature. If you want to upload the depot by using the Patch Management feature, you must detach a depot from any nodes that it is attached to, and then delete it from the Software Repository.

For HP-UX 10.20, you can only apply patches through the Install Software Wizard because Opsware SAS recognizes them as software and not patches.

Patch Uploads for Unix

Before a patch can be installed on a managed server with Opsware SAS, the patch must be uploaded into the Software Repository. Uploading patches is the responsibility of the patch administrator.

Patch Uploads for Specific Unix Versions

When a patch is uploaded, you associate the patch with a specific version of an operating system. When you upload a Solaris patch, for example, you must select the version of the Solaris operating system that this patch applies to, such as Solaris 5.6 or 5.9. You can only install this patch on servers that are running that version of the operating system.

If, for any reason, you need to install a given patch across servers running different versions of the same operating system, you need to upload the patch multiple times and associate the patch with each of the operating system versions that the patch applies to.

For example, if the same Solaris patch needs to be installed on servers running Solaris 2.7 and 2.8, you must upload the patch two times. The first time that you upload the patch, you associate it with the Solaris 2.7. You then repeat the procedure and associate the patch with Solaris 2.8. (This procedure also allows you to specify different installation options. The different versions of the same operating system can sometimes require different installation scripts, installation flags, and so forth.)

In the case of application patches, it is even more common that you need to upload a patch multiple times. A Solaris patch for Oracle, for example, often needs to be applied to instances of Oracle running on slightly different versions of the Solaris operating system.

Methods for Uploading a Unix Patch

If you upload a patch through the Opsware Command Center, the Upload Patch Wizard guides you through the process. The Upload Patch Wizard allows you to specify a number of options for the patch, including install and post-install scripts, install and uninstall flags, and other options.

Because the Opsware Command Center is a browser-based interface, you can only upload one patch or patch container (such as a Solaris patch cluster or an HP-UX depot) at a time. If you want to upload multiple patches at the same time, such as a large set of AIX LPPs, you can do so more quickly through the OCLI.

If you upload patches through the OCLI, however, you are not able to specify installation options during the upload process. Instead, you specify these options by editing the patches through the Opsware Command Center.

Preparation for Uploading a Patch for Unix

Before you upload a patch, you must copy it to a location that is accessible to the browser that you are using or the OCLI. If you are using the Opsware Command Center, you specify the path of the patch in the upload wizard, either by entering it directly or by browsing for the patch.

In some cases, you need to install patches in a particular order. After you upload the patches, you can create an installation order dependency by using the Opsware Command Center.

Uploading a Patch for Unix

To upload a patch with the Opsware Command Center, perform the following steps:

1. Launch the Upload Patch Wizard from the Patch Management pane on the Opsware home page. The Select Patch page appears, as Figure 13-1 shows.

Figure 13-1: Select Patch Page

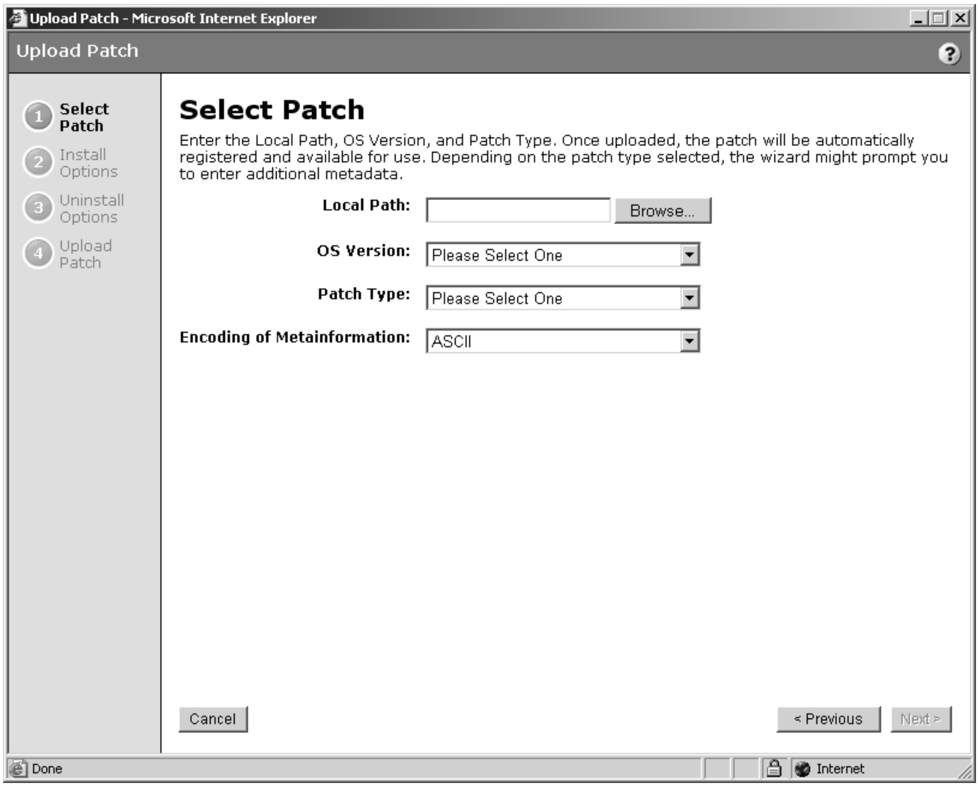

Select Patch

Enter the Local Path, OS Version, and Patch Type. Once uploaded, the patch will be automatically registered and available for use. Depending on the patch type selected, the wizard might prompt you to enter additional metadata.

Local Path:

OS Version: Please Select One

Patch Type: Please Select One

Encoding of Metainformation: ASCII

2. Either enter the fully qualified path of the patch that you want to upload, or click **Browse** and navigate to the patch that you want to upload.

3. Select the OS version of the patch that want to upload. You must be certain to select the correct operating version, or the patch will not be available for the correct operating system.

4. Select the type of patch that you are uploading. You must be careful to select the correct patch type or the patch will be misapplied or uninstallable. The Opsware Command Center only allows you to select patches that are appropriate for the

operating system that you select, but it is still possible to select the wrong kind of patch. (For example, selecting a Solaris patch when you intended to select a Solaris patch cluster.)

5. Click **Next** to continue to the Install Options page. In this page, you can specify a number of installation options:

- **Install Flags passed directly to the patch installer**: The Patch Management feature also passes a number of default flags.

  When installing an AIX Update fileset, the Patch Management feature normally applies the fileset, which allows it to be rejected (uninstalled). If you want to commit the fileset instead (so that it cannot be removed), use the -c option here.

- **Pre-install Script**: Enter the pre-install script into this box. If you want to terminate the installation of the patch if the script returns a non-zero return code, select the check box.

- **Post-Install Script**: Enter the post-install script into this box. If you want to terminate the installation of the patch if the script returns a non-zero return code, select the check box.

- **Reboot on Install**: Select this option if the patch you are removing requires a reboot. Keep in mind that other patches can be directly applied after this patch, so be sure to check this option if it is necessary.

6. Click **Next** to continue to the Uninstall Options page.

   In this page, you can specify the following uninstallation options:

- **Uninstall Flags passed directly to the installer**: The Opsware Patch Management System passes a number of default uninstall flags to the installer.

- **Pre-uninstall Script**: Enter the Pre-uninstall script into this box. If you want to terminate the installation of the patch if the script returns a non-zero return code, select the check box.

- **Post-uninstall Script**: Enter the post-uninstall script into this box. If you want to terminate the installation of the patch if the script returns a non-zero return code, select the check box.

- **Reboot on Uninstall**: Select this option if the patch that you are removing requires an immediate reboot. Keep in mind that other patches can be directly applied after this patch, so be sure to check this option if it is necessary.

7. Click **Next** to upload the patch. A progress bar appears.
8. After the patch is uploaded, you have the option to install the patch. Click **Yes** to install the patch, and then click **Next**. Otherwise, click **No** or click **Close**.

Remember that if you need to upload the same patch for multiple versions of the same operating system, you must repeat this process with the same patch.

Patch Testing for Unix

After you upload the patch, you can install the patch on a testing server with the Patch Install Wizard. As the patch administrator, you can install the patch, even though the patch is automatically set in the Untested state after you upload it the first time. When you finish testing the patch, use the Opsware Command Center to change the patch state to Available for Use so that system administrators can install the patch.

Patch Administration for Unix

The Opsware Command Center allows you to search through all patches that have been uploaded. You can use the Opsware Command Center to edit patch options and create install order dependencies, and change the state of patches to Available for Use so that system administrators can install them. You can also view detailed information about individual patches, such as the number of times the patch has been installed.

Setting the Patch Status for Unix

The patch administrator sets the status of a patch in Opsware SAS. The status determines who can apply the patch, or if the patch can be applied at all.

Table 13-3 describes the statuses that patches in Opsware SAS can have.

Table 13-3: Patch Statuses in Opsware SAS

STATUS	DESCRIPTION
Untested	Initial state of a patch after being uploaded. Only administrators with special privileges can install untested patches.
Available for Use	Has been uploaded and approved by the patch administrator and can be installed on servers.

Table 13-3: Patch Statuses in Opsware SAS

STATUS	DESCRIPTION
Deprecated	The patch is possibly still installed on some systems, but cannot be installed anymore, not even by a user who is a member of the Advanced User role.

To set the patch status, perform the following steps:

1. From the navigation panel of the Opsware Command Center, click Software ➤ Patches. The Patches page appears.
2. Select the filter options from the drop-down menus to display the type of patch whose status you want to change. Select the patch type, the operating system version, and the patch state. See Figure 13-2.

Figure 13-2: Search Example

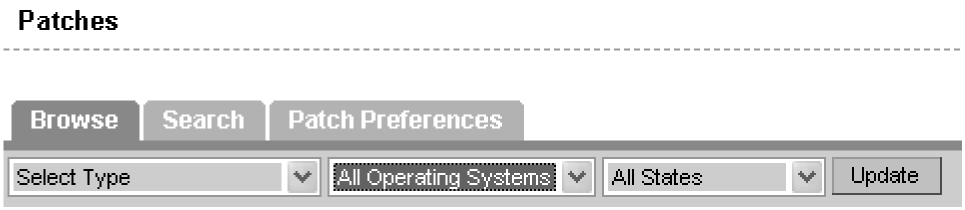

3. Click **Update** to display the list of patches that meet your selection criteria.
4. Locate the patch and click the name of the patch. The View Patch page appears.
5. In the patch summary section of the View Patch page, click **Edit**. The Edit Patch page appears.
6. Select the desired status from the Patch Status drop-down menu and then click **Save**.

Editing Patch Options for Unix

You can edit any of the options that you specified for a patch that you uploaded using the Patch Upload Wizard. Additionally, if you uploaded a patch with the OCLI, you can specify the same options for the patch by editing the patch options.

Some patch option are not editable, due to the nature of the patch type. For example, you cannot set patch status on an HP-UX patch fileset, because you can only set the patch status on the parent HP-UX patch product. (After you change the status of the parent HP-UX patch product, the change is applied to the children filesets.) Other options cannot be set, because they do not apply to the patch type that you are editing.

To edit patch options, perform the following steps:

1. From the navigation panel of the Opsware Command Center, click Software ➤ Patches. The Patches page appears.
2. Select the options from the drop-down menus to display the type of patch that you want to edit. You select the operating system version, the patch type, and the patch state.
3. Click **Update** to display a list of patches that match your selected criteria.
4. Locate the patch that you want to edit and click the link for the patch name. The View Patch page appears.
5. Click **Edit** to edit the patch options. (Click **Edit** in the Install Options or in the Uninstall Options, as appropriate.)
6. Add or modify the patch install or uninstall options and click **Save**.

If you are modifying the options of a patch that you already marked as Available for Use, consider resetting the status of the patch back to Untested. Test the patch again with the new options, and set the status back to Available when you determine that it is safe to install the patch again.

Setting Patch Installation Order Dependencies for Unix

For some patch types, install order dependencies can be set. Before setting patch dependencies, you must first upload the patch using the Patch Upload Wizard.

To set patch installation order dependencies, perform the following steps:

1. From the navigation panel of the Opsware Command Center, click Software ➤ Patches. The Patches page appears.
2. Select the options from the drop-down menus to display the type of patch that you want to edit. Select the operating system version, the patch type, and the patch state.

3. Click **Update** to display a list of patches that match your selected criteria.

4. Locate the patch that you want to edit and click the link for the patch name. The Patch Summary page appears.

5. Click **Edit** in the Install Order section. See Figure 13-3.

Figure 13-3: Install Order Section

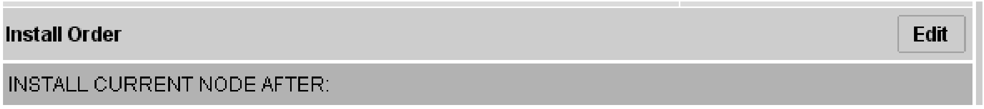

Install Order

INSTALL CURRENT NODE AFTER:

6. Click **Add** to select the type of software that must be installed before the selected path. See Figure 13-4.

Figure 13-4: Software Type

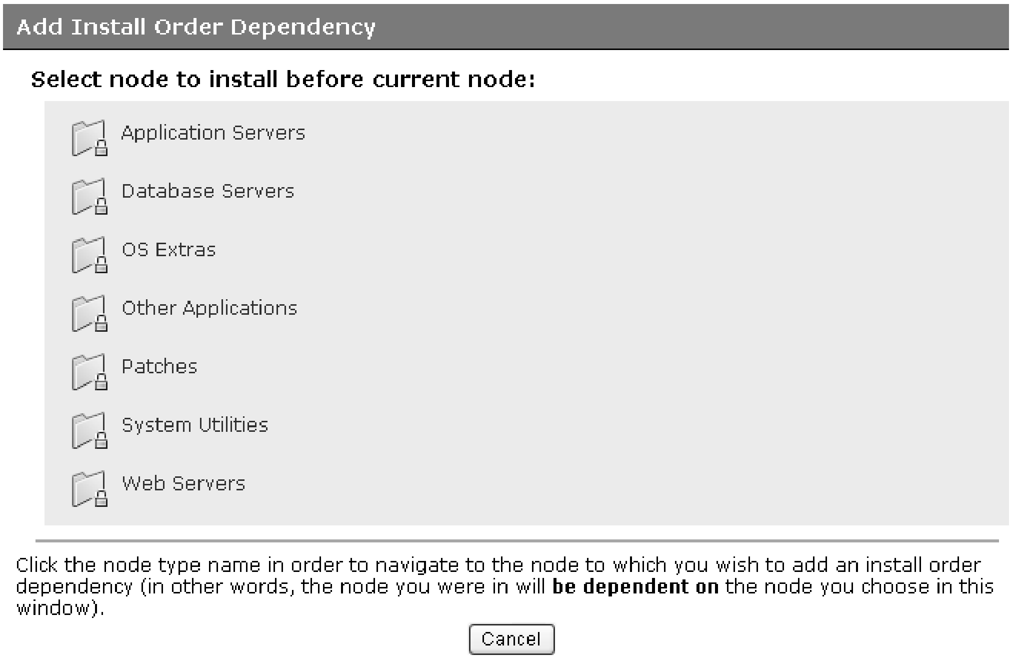

Add Install Order Dependency

Select node to install before current node:

- Application Servers
- Database Servers
- OS Extras
- Other Applications
- Patches
- System Utilities
- Web Servers

Click the node type name in order to navigate to the node to which you wish to add an install order dependency (in other words, the node you were in will **be dependent on** the node you choose in this window).

7. Browse for the software package or patch that must be installed before your selected patch.

8 Select the check box next to the desired software package and then click **Add**. See Figure 13-5.

Figure 13-5: Adding Install Order Dependency

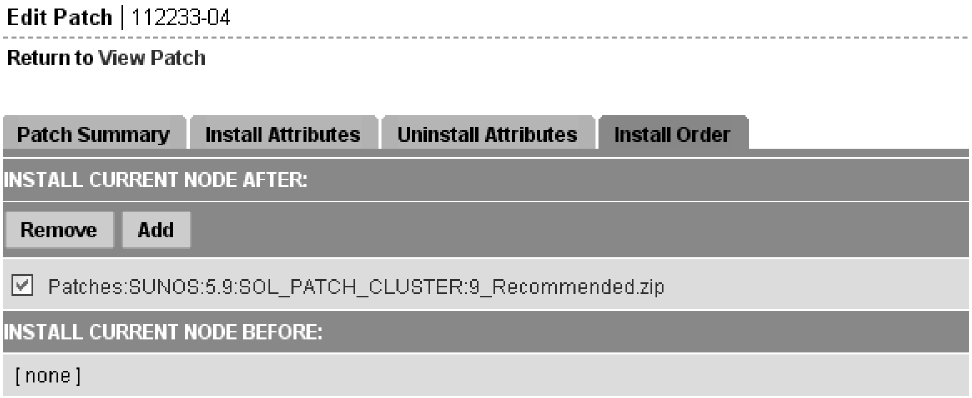

Edit Patch | 112233-04

Return to View Patch

INSTALL CURRENT NODE AFTER:

☑ Patches:SUNOS:5.9:SOL_PATCH_CLUSTER:9_Recommended.zip

INSTALL CURRENT NODE BEFORE:

[none]

9 Confirm the dependency by clicking **Add** in the View Patch page. If you click **Add** again, the confirmation page does not appear. You will instead see the Add Install Order Dependency page, which allows you to add more packages.

10 Repeat the process if other dependencies must be expressed.

Patch Installation and Uninstallation for Unix

Installing and uninstalling patches are responsibilities of the systems administrator.

Typically, after the patch administrator has tested and approved a patch, the patch administrator notifies the system administrator that the patch is ready to be installed. The system administrator installs the patches with the Install Patch Wizard.

Patching and uninstalling patches frequently causes disruptions in service. Often, installing or uninstalling a patch will cause a server to reboot. In order to minimize the effects of service disruption, you can schedule the installation of patches.

The Patch Management feature allows you to install both operating system patches and application patches. The procedure for installing application patches is slightly different than the process for installing operating system patches, because you must perform an additional search to find the servers where the application is installed.

The Install and Uninstall Patch Wizards allow you to select patches and servers by either browsing or searching.

Application Patches for Unix

The Patch Management feature does not allow you to apply a patch to an operating system for which the patch is not intended. When you are installing an application patch, however, the Patch Management feature does not automatically filter out servers that do not have the application installed that the patch is intended for. Although the Patch Management feature does not prevent you from doing so, you must not attempt to apply application patches to servers that do not have the necessary applications installed.

If an application patch is intended for an application that is running on more than one version of the same operating system, you cannot apply the patch to all of the servers at the same time. An application patch is associated with only one operating system version. You must first select the patch for one operating system, select the servers where the application is installed, and apply the patch. You must repeat this process for each version of the operating system where the application is installed.

Similarly, when uninstalling application patches that are installed on multiple versions of the same operating system, you cannot uninstall all of the patches at the same time. You must repeat the uninstallation process for each version of the operating system where the patch is installed.

Installation Scripts for Patches on Unix

When you upload a patch, you can specify the following types of scripts:

- Pre-installation scripts that are executed before a patch is installed
- Post-installation scripts that are executed after a patch is installed
- Pre-uninstallation scripts that are executed before a patch is uninstalled
- Post-uninstallation scripts that are executed after a patch is uninstalled

A typical use of a pre-installation script is to shut down a process before you apply a patch to the application that the process belongs to. The post-installation script then restarts the process after the patch is applied.

The type of script must be supported by the operating system of the server where the patch is to be installed. The required shells and utilities must be installed on the servers where you plan to run the scripts. (Also, environment variables might need to be set.) For example, a pre-installation shell script can run a Python script, but Python must be installed on the servers where you want to run the scripts. If you want to run an unwrapped Python script, be sure to specify the `python` command before the script name.

Installation and Uninstallation Flags for Patches on Unix

You can specify installation and uninstallation flags that are applied whenever a patch is installed or uninstalled.

Opsware SAS, however, also uses default installation and uninstallation flags. Opsware SAS requires that patches are installed and uninstalled with these flags. You must therefore be certain that you do not specify any installation or uninstallation flags that override or contradict the default flags passed in by Opsware SAS.

The following table lists the default installation flags that Opsware SAS uses.

Table 13-4: Default Installation Flags

OPERATING SYSTEM/PATCH TYPES	FLAGS
AIX	`-a -Q -g -X -w`
HP-UX	None

The following table lists the default uninstallation flags that Opsware SAS uses.

Table 13-5: Default Uninstallation Flags

OPERATING SYSTEM/PATCH TYPES	FLAGS
AIX	`-u -g -X`
AIX Reject Options	`-r -g -X`
HP-UX	None

Installing a Unix Patch

You install Unix patches with the Install Patch Wizard of the Opsware Command Center.

The wizards are flexible; for example, when you select multiple servers and patches, Opsware SAS will install the correct patches on the servers even if you selected different OS versions for the servers and patches in the respective steps. Server Groups are not associated with a specific OS or customer; therefore, you can add a patch for any OS to the group. (Patches are always associated with the customer, Customer Independent.) If Opsware SAS cannot find a match, an error message "No valid devices found" appears at the end of the wizard; therefore, use caution when modifying these default values.

To install a Unix operating system patch, perform the following steps:

1. From the Opsware Command Center home page, click the Install Patch link in the Tasks panel.

   Or

   From the navigation panel, click Software ➤ Patches. The Patches page appears. Select the patch and click **Install** in the View Patch page.

2 The Select Patches page appears, as Figure 13-6 shows.

Figure 13-6: Select Patches Page

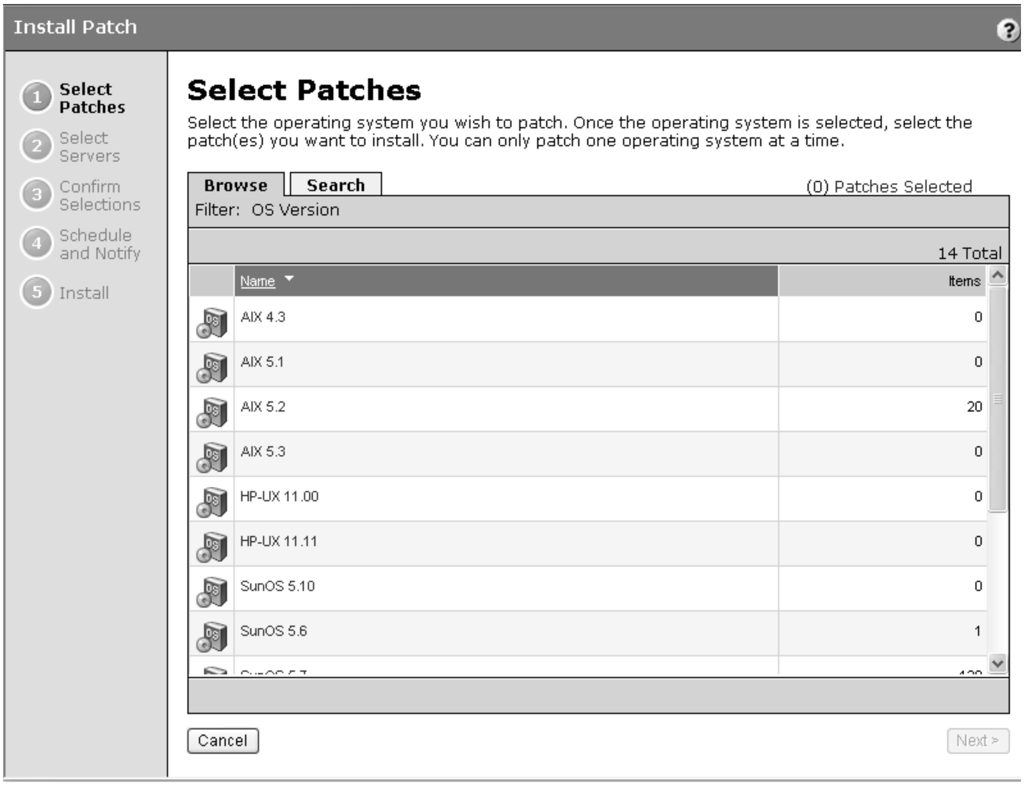

Install Patch

1 Select Patches
2 Select Servers
3 Confirm Selections
4 Schedule and Notify
5 Install

Select Patches

Select the operating system you wish to patch. Once the operating system is selected, select the patch(es) you want to install. You can only patch one operating system at a time.

(0) Patches Selected

Filter: OS Version

14 Total

Name	Items
AIX 4.3	0
AIX 5.1	0
AIX 5.2	20
AIX 5.3	0
HP-UX 11.00	0
HP-UX 11.11	0
SunOS 5.10	0
SunOS 5.6	1

3 Select the operating system version for the patch that you want to apply.

4 Select one or more patches that you want to install and then click **Next**.

5. Select the servers or groups that you want to patch and then click **Next**. The Confirm Selection page appears, as Figure 13-7 shows.

Figure 13-7: Confirm Selection Page

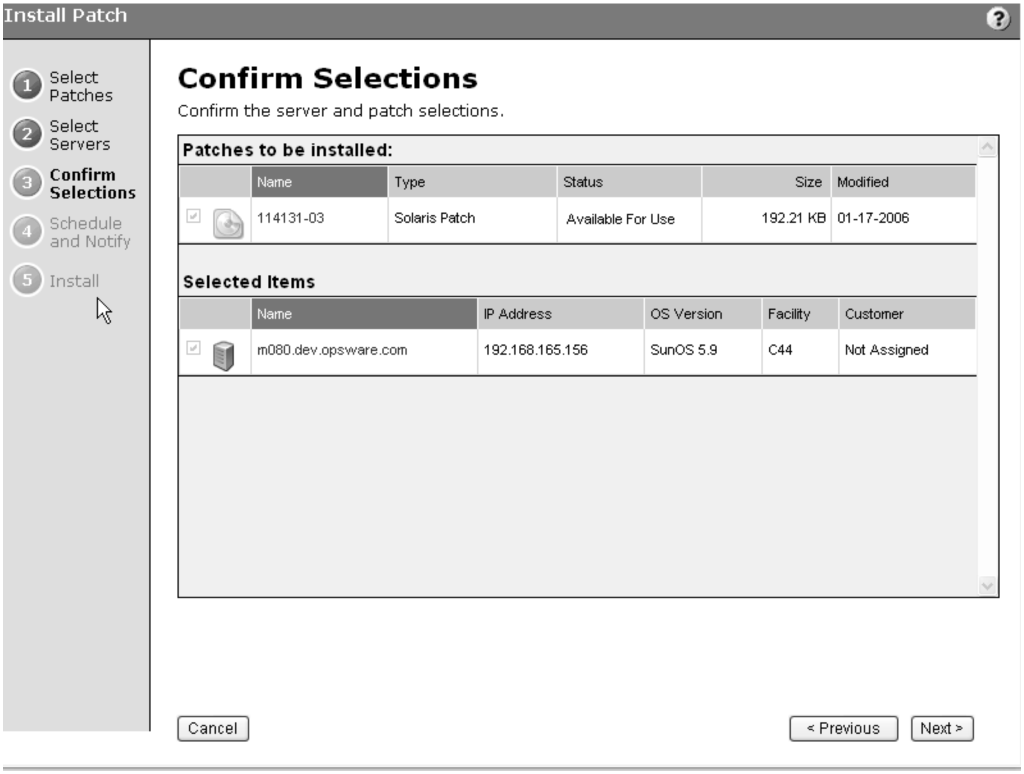

Confirm Selections

Confirm the server and patch selections.

Patches to be installed:

	Name	Type	Status	Size	Modified
☑	114131-03	Solaris Patch	Available For Use	192.21 KB	01-17-2006

Selected Items

	Name	IP Address	OS Version	Facility	Customer
☑	m080.dev.opsware.com	192.168.165.156	SunOS 5.9	C44	Not Assigned

6. Review your selections.

7. On the Schedule and Notify page, you have the following options:

   - **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.

   - **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

     When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have "Allow Run Refresh Jobs" permissions, you will see additional options. Select either of the following options:

- **Option 1**: Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
- **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

If you decided to install the patches immediately by clicking **Install**, a progress bar appears, as Figure 13-8 shows.

Figure 13-8: Install Progress

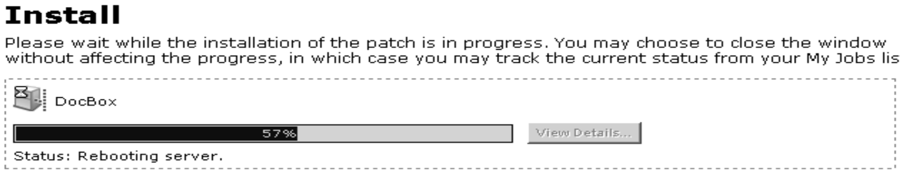

Install

Please wait while the installation of the patch is in progress. You may choose to close the window without affecting the progress, in which case you may track the current status from your My Jobs list.

DocBox

57%

Status: Rebooting server.

8 After the installation completes, click **View Details** for more information about the results of the installation operation.

9 Click **Close** to end the wizard.

Closing the wizard does not stop the installation if it is still running. After you close the wizard, you can view the progress of the installation by viewing My Jobs.

See “Server Management Scheduling and Notification” on page 276 in Chapter 7 for information about the My Jobs feature.

Installing an Application Patch on Unix

You install application patches on Unix with the Install Patch Wizard of the Opsware Command Center.

The wizards are flexible; for example, when you select multiple servers and patches, Opsware SAS will install the correct patches on the servers even if you selected different OS versions for the servers and patches in the respective steps. Server Groups are not associated with a specific OS or customer; therefore, you can add a patch for any OS to the group. (Patches are always associated with the customer, "Customer Independent.") If Opsware SAS cannot find a match, an error message "No valid devices found" appears at the end of the wizard; therefore, use caution when modifying these default values.

To install an application patch on a Unix server, perform the following steps:

1. From the Opsware Command Center home page, click the Install Patch link in the Tasks panel.

   Or

   From the navigation panel, click Software ➤ Patches. The Patches page appears. Select the patch and click **Install** in the View Patch page.

2. Select the operating system version of the servers where the application is installed. A list of all uploaded patches for that operating system appears.

3 Select one or more application patches (for a given application) that you want to install and then click **Next**. The Select Servers page appears, as Figure 13-9 shows.

Figure 13-9: Select Servers

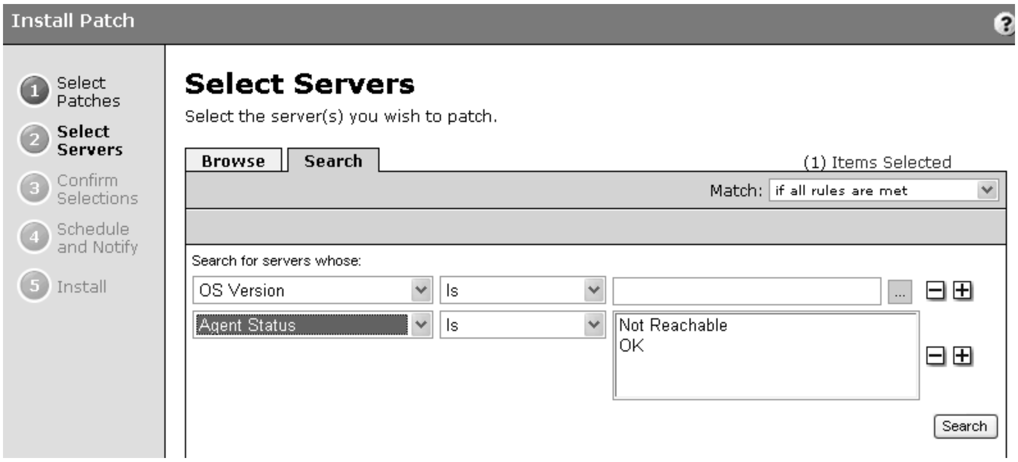

Install Patch

1 Select Patches
2 Select Servers
3 Confirm Selections
4 Schedule and Notify
5 Install

Select Servers

Select the server(s) you wish to patch.

Browse | Search

(1) Items Selected

Match: if all rules are met

Search for servers whose:

OS Version | Is |

Agent Status | Is | Not Reachable
OK

Search

4 Select the Search tab and search for the servers where the application that you want to patch is installed and then click **Search**.

5 From the search results, select the servers or groups running the application that you want to patch and then click **Next**. The Confirmation page appears. Review your selections.

6 On the Schedule and Notify page, you have the following options:

- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.

- **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

  When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

- **Option 1**: Membership is determined based on the "Time of Confirm Selection." Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
- **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

If you decided to install the patches immediately by clicking **Install**, a progress bar appears.

7 After the installation completes, click **View Details** for more information about the results of the installation operation.

8 Click **Close** to end the wizard.

Closing the wizard does not stop the installation if it is still running. After you close the wizard, you can view the progress of the installation by viewing My Jobs.

See "Server Management Scheduling and Notification" on page 276 in Chapter 7 for information about the My Jobs feature.

Uninstalling a Unix Patch

To uninstall a Unix operating system patch, perform the following steps:

1. From the Opsware Command Center home page, click the Uninstall Patch link in the Tasks panel.

   Or

   From the navigation panel, click Software ➤ Patches. The Patches page appears. Select the patch and click **Uninstall** in the View Patch page.

2. The Select Patches page appears, as Figure 13-10 shows.

Figure 13-10: Select Patches Page

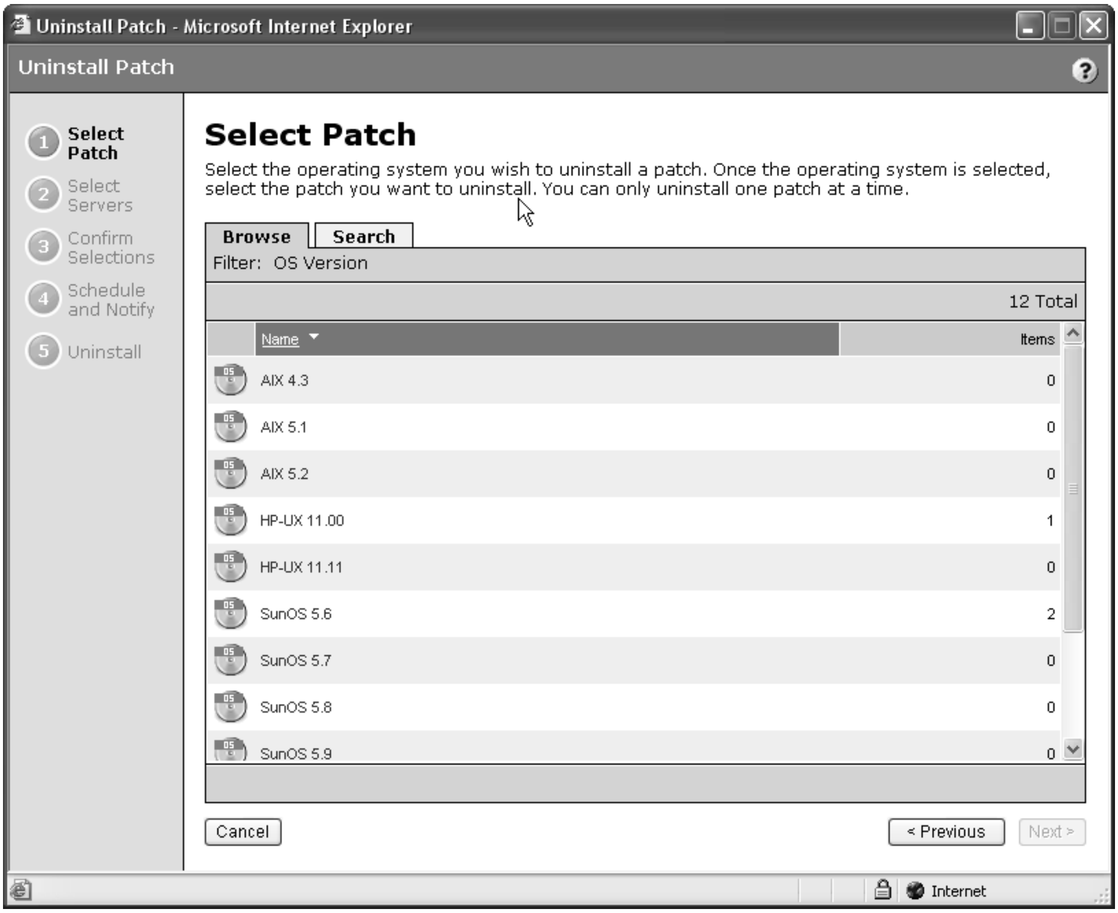

Uninstall Patch - Microsoft Internet Explorer

Uninstall Patch

1 Select Patch
2 Select Servers
3 Confirm Selections
4 Schedule and Notify
5 Uninstall

Select Patch

Select the operating system you wish to uninstall a patch. Once the operating system is selected, select the patch you want to uninstall. You can only uninstall one patch at a time.

Browse | Search

Filter: OS Version

12 Total

Name	Items
AIX 4.3	0
AIX 5.1	0
AIX 5.2	0
HP-UX 11.00	1
HP-UX 11.11	0
SunOS 5.6	2
SunOS 5.7	0
SunOS 5.8	0
SunOS 5.9	0

3. Select the operating system version for the patch to uninstall. A list of all installed patches for that operating system appears. You can also search for the patch. (Patches that Opsware SAS did not install will display, but you cannot select them.

4 Select the patch that you want to uninstall and then click **Next**. (You can only select one patch to uninstall.) A list of servers that have the selected patch installed appears.

5 Select the servers that you want to uninstall the patches from and then click **Next**. The Confirm Selection page appears.

6 Review your selections.

7 On the Schedule and Notify page, you have the following options:

- **Schedule**: Click either **Run Now** to execute the operation immediately, or click **Specify Time** to schedule the operation for a later time.
- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.

If you decide to uninstall the patches immediately by clicking **Uninstall**, a progress bar appears.

8 After the uninstallation completes, click **View Details** for more information about the results of the uninstallation operation.

9 Click **Close** to end the wizard.

Closing the wizard does not stop the uninstallation if it is still running. After you close the wizard, you can view the progress of the uninstallation by viewing My Jobs.

See "Server Management Scheduling and Notification" on page 276 in Chapter 7 for information about the My Jobs feature.

Uninstalling an Application Patch on Unix

To uninstall an application patch on a Unix server, perform the following steps:

1 From the Opsware Command Center home page, click the Uninstall Patch link in the Tasks panel.

Or

From the navigation panel, click Software ➤ Patches. The Patches page appears. Select the patch and click **Uninstall** in the View Patch page.

2 Select the operating system version of the servers where the application is installed. A list of all uploaded patches for that operating system appears.

3. Select the application patches that you want to uninstall and then click **Next**. The Select Server page appears.

4. Select the Search tab and search for the servers that have the patch that you want to remove and then click **Search**.

5. From the search results, select the servers or groups whose application patch you want to remove and then click **Next**. The Confirmation page appears.

6. Review your selections.

7. On the Schedule and Notify page, you have the following options:

   - **Schedule**: Click either **Run Now** to execute the operation immediately, or click **Specify Time** to schedule the operation for a later time.

   - **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.

   If you decide to uninstall the patches immediately by clicking **Uninstall**, a progress bar appears.

8. After the uninstallation completes, click **View Details** for more information about the results of the uninstallation operation.

9. Click **Close** to end the wizard.

   Closing the wizard does not stop the uninstallation if it is still running. After you close the wizard, you can view the progress of the uninstallation by viewing My Jobs.

See “Server Management Scheduling and Notification” on page 276 in Chapter 7 for information about the My Jobs feature.

Chapter 14: Software Provisioning

IN THIS CHAPTER

This section discusses the following topics:

- Installing and Uninstalling Software
- Software Installation and Uninstallation Options
- Template Installation

Installing and Uninstalling Software

This section provides information on how to install and uninstall software by using Opsware SAS and contains the following topics:

- Overview of Installing and Uninstalling Software
- Software Package Management in the Model Repository
- Ways to Install Software
- Opsware SAS Categories for Software
- Software Provisioning Functionality
- Platform-Specific Reconcile

Overview of Installing and Uninstalling Software

Opsware SAS automates the time-consuming process of installing and uninstalling software on servers. Using Opsware SAS, you can quickly deploy applications across a large number of servers with a minimum amount of downtime. Opsware SAS allows you to select the packages from the Software Repository, select the servers that you want to install the software on, preview the results of the installation, and install the software all in a single operation.

Opsware SAS provides detailed feedback about which packages are installed, what events occurred (such as a server reboot), the output of scripts, and any errors that occurred. Opsware SAS also provides the same degree of control for uninstalling software.

You can uninstall any software that you installed through Opsware SAS. You can also preview the results of an installation and then schedule the installation for some later time. If, for example, Opsware SAS reports that the software installation requires multiple reboots, you can schedule the installation for a time when the reboots cause the least disruption to your services.

Software Package Management in the Model Repository

In Opsware SAS, packages reside in a central Software Repository. Opsware users upload the software and also specify options that ensure that the software is correctly installed. The users add pre- and post-install and uninstall scripts to packages that help control the way that the software is installed.

See the *Opsware® SAS Configuration Guide* for information about how to upload packages to and manage packages in the Software Repository.

Opsware SAS maintains detailed information about the state of every server under management in a central database called the Model Repository.

Ways to Install Software

Opsware SAS provides several ways to install software. You can:

- Select a single package and install it.
- Select a template that includes a number of different (and usually related) software that you can install in a single operation.
- Use Visual Packager to prepare installable software packages. See Chapter 9, "Visual Packager" on page 363 of this guide for more information.

Opsware users create and test the templates, and they provide the same level of control and consistency that is available when you install individual software. Templates, for example, allow you to install a set of applications, such as a Web server, an Oracle database, and related software that allow you to quickly deploy a server in a fully operational state.

Opsware SAS Categories for Software

The Model Repository usually stores information about hundreds of different types of packages. To make the selection of software easier, Opsware SAS organizes software into the following categories:

- Application Servers
- Database Servers
- Operating System Extras
- Other Applications
- System Utilities
- Web Servers

Each of these categories can contain software for any of the operating systems that Opsware SAS supports. You can browse through the list of software or use advanced search capabilities in Opsware SAS to locate the packages or templates that you want to apply.

Software Provisioning Functionality

The Software Provisioning feature has the following functionality:

- A central location (the Software Repository) for package storage
- Consistency of installation with user-created scripts that specifies options to ensure that software is installed in a uniform way
- The ability to preview the installation and uninstallation processes to see what packages will be installed and what server operations are required (such as rebooting)
- The ability to install individual packages quickly across a large number of servers
- The ability to quickly apply templates, which includes bundles of applications, across a large number of servers
- Auditing of all changes made to managed servers and quick uninstallation in case problems result

Platform-Specific Reconcile

The Reconcile step is the last step that you perform when installing software. During that step, Opsware SAS provides you with information about software that is currently installed on a server, along with information about the software that is about to be installed, and software that is about to be removed. See Appendix A for information about reconcile.

Table 14-1 gives details about reconcile operations for various operating systems. This table shows the utilities used during the reconcile process for all supported operating systems.

Table 14-1: Reconcile Operation by OS

OPERATING SYSTEM	SEE THIS PAGE
AIX	See “AIX Reconcile” on page 707.
HP-UX	See “HP-UX Reconcile” on page 708.
Linux	See “Linux Reconcile” on page 708.
Solaris	See “Solaris Reconcile” on page 708.

Software Installation and Uninstallation Options

When you upload software to the Software Repository, a user can specify a number of options that affect what happens when you install software by using the Install Software Wizard. The following sections explain the consequences of these options:

- Script Error Conditions
- Inheritance and the Install and Uninstall Software Wizards
- Installing Software with the Install Software Wizard
- Package Installation on Servers with Low Disk Space
- Uninstalling Software with the Uninstall Software Wizard

Script Error Conditions

A package that you select in the Install Software Wizard can include installation scripts. A user can specify what happens if the pre- or post-install script encounters an error condition. The user can specify an option that allows the software installation or uninstallation to proceed in spite of an error, or the user can select an option that prevents the software installation if the install script returns an error (for example, exits with a non-zero return code). You can view the error message when you preview the installation or uninstallation.

Inheritance and the Install and Uninstall Software Wizards

Software in Opsware SAS is organized into a Software Tree. Each branch of the tree supports inheritance. A node lower in the branch inherits all the software and attributes of the nodes above it in the same branch.

See the *Opsware® SAS Configuration Guide* for information about how to set up inheritance for software nodes.

Ordinarily, a branch in the tree is set up for one package. In some cases, however, the branch might include further nodes. For example, the child node of an application node might contain patches for the application.

This same inheritance tree displays in the Install and Uninstall Software Wizards. The branches of the Software Tree display as a familiar folder hierarchy.

In most cases, the last level of the folder hierarchy contains a particular application that you want to install or uninstall. Keep in mind, however, that if the folder hierarchy continues with additional software (such as patches), all software above the node that you select is also installed or uninstalled. For example, if you are using the Install Software Wizard and you select a subfolder that contains a patch for an Oracle database – and the parent folder contains the Oracle database itself – selecting the patch for the database also causes the database itself to be installed.

Installing Software with the Install Software Wizard

The wizards are flexible; for example, when you select multiple servers and applications, Opsware SAS will install the correct applications on the servers even if you selected different OS versions for the servers and applications in the respective steps. Opsware SAS also matches the customer association for applications, operating systems, and templates with the customer association of servers. If Opsware SAS cannot find a match, an error message appears at the end of the wizard; therefore, use caution when modifying these default values.

Perform the following steps to install software with the Install Software Wizard:

1. From the Opsware Command Center home page, click the Install Software link in the Tasks panel.

   Or

   From the navigation panel, click Software ➤ Applications. The Application page appears. Select the application and click **Install**.

2. The Select Software page appears, as Figure 14-1 shows.

3. Select the category of software that you want to install (Application Server, Database Server, and so forth), or select the Search tab and search for the software. Select one or more software to install. Figure 14-1 shows an example of possible software selections.

Figure 14-1: Select Software Page

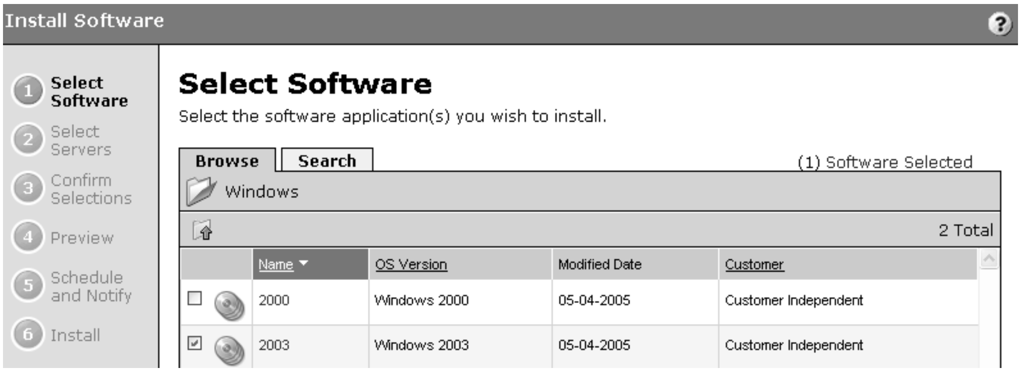

Select Software

Select the software application(s) you wish to install.

(1) Software Selected

Windows

2 Total

	Name	OS Version	Modified Date	Customer
□	2000	Windows 2000	05-04-2005	Customer Independent
☑	2003	Windows 2003	05-04-2005	Customer Independent

4. (Optional) Click the name of any of the software to display additional information about the software, as Figure 14-2 shows.

Figure 14-2: Software Details Page

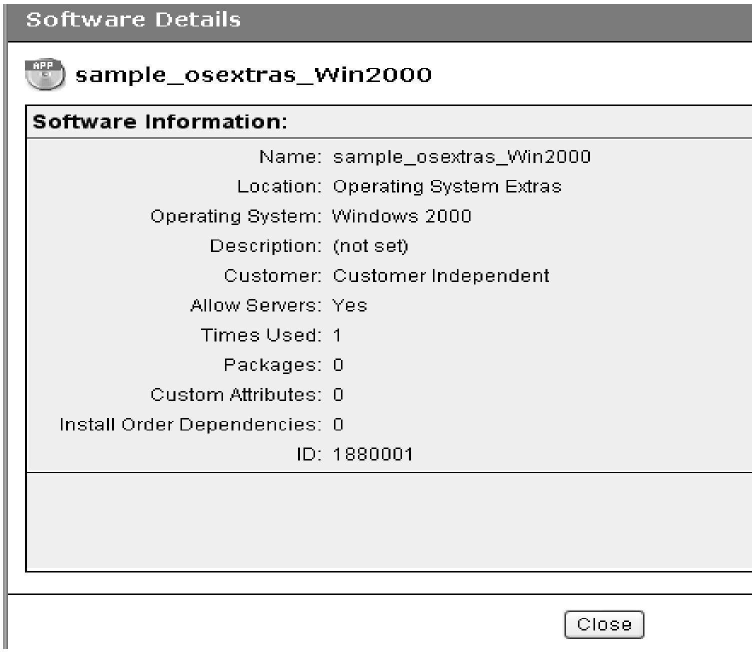

Software Details

sample_osextras_Win2000

Software Information:

Name: sample_osextras_Win2000
Location: Operating System Extras
Operating System: Windows 2000
Description: (not set)
Customer: Customer Independent
Allow Servers: Yes
Times Used: 1
Packages: 0
Custom Attributes: 0
Install Order Dependencies: 0
ID: 1880001

5 Click **Next** to continue. The Select Servers page appears, as Figure 14-3 shows.

Figure 14-3: Select Servers Page

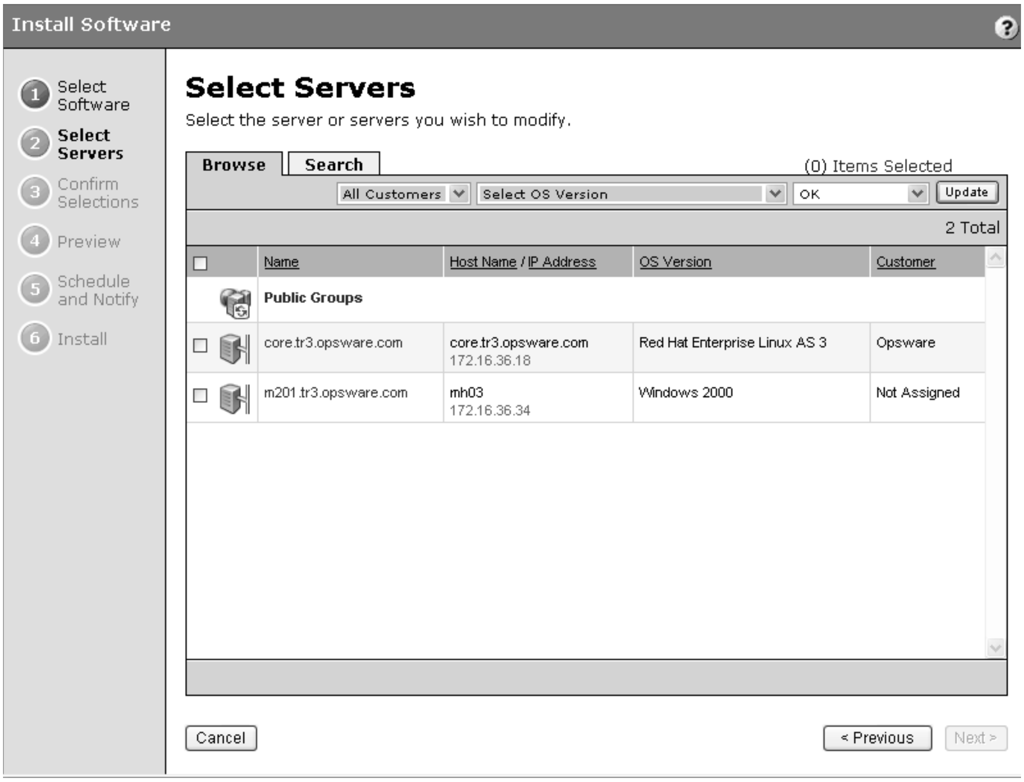

Install Software

1 Select Software
2 **Select Servers**
3 Confirm Selections
4 Preview
5 Schedule and Notify
6 Install

Select Servers

Select the server or servers you wish to modify.

Browse | Search — (0) Items Selected

All Customers | Select OS Version | OK | Update

2 Total

□	Name	Host Name / IP Address	OS Version	Customer
	Public Groups			
□	core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	Red Hat Enterprise Linux AS 3	Opsware
□	m201.tr3.opsware.com	mh03 172.16.36.34	Windows 2000	Not Assigned

The selection of servers is defaults to common operating system and customer of the software that you selected.

6 Select the servers or groups where you want to install the software by browsing or searching. Click **Next** to continue. The Confirm Selection page appears that displays the servers and software that you selected in previous steps.

7 (Optional) Review your selections and click **Preview** to continue. The Preview Page appears. The Preview Page initially displays status bars for each server that indicates the progress of the preview process.

Or

Click **Skip Preview** to skip the preview process.

8 On the Schedule and Notify page, you have the following options:

- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.
- **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

  When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

  - **Option 1**: Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
  - **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

9. (Optional) When the process is completed, click **View Details** to see what packages were installed or removed, and the output, if any, from the installation scripts.

10. Click **Close** to end the wizard.

    Closing the wizard does not stop the installation if it is still running. After you close the wizard, you can view the progress of the installation by viewing My Jobs.

See “Server Management Scheduling and Notification” on page 276 in Chapter 7 for information about the My Jobs feature.

Package Installation on Servers with Low Disk Space

If a managed server does not have enough disk space to download packages from the Software Repository to be installed, you can specify locations such as shared network drives, or a CD-ROM, where Opsware SAS should look for packages to install.

To accomplish this task you need to specify the location where Opsware SAS should look for packages to install by defining a new custom attribute. See See "Custom Attributes for Servers" on page 282 in Chapter 7 for more information.

You can enter any number of paths that Opsware SAS attempts to use to find the package. Opsware SAS tries each path, one after the other, until the package is located. If the package is not found in any of the specified locations, Opsware SAS looks in the Software Repository. In that case, Opsware SAS attempts to download the package if there is enough disk space. If Opsware SAS calculates that there is not enough disk space, an error message appears, and the packages are not downloaded.

Check that permissions on files are set so that Opsware SAS has read access.

Table 14-2 shows the operating systems and package types that you can use with the Low Disk Space feature.

Table 14-2: Supported Operating Systems for the Low Disk Space Feature

OPERATING SYSTEM	FILE TYPE
Sun Solaris	RPM only
Linux	RPM
IBM AIX	RPM, APAR, Base Fileset, Update Fileset, Maintenance Level,
HP-UX	Depots, disk format only

Special Operating System Requirements

Some operating systems have additional requirements when you use this low disk space feature:

- HP-UX patches must be in disk format.
- When you add new install packages to an AIX installation directory, always run the `inutoc` command to ensure that the installation subsystem recognizes the new packages. This command creates a new .toc file.

  Run the `inutoc` command in the AIX installation directory or with the installation directory as the only parameter.

Write access as root to the installation directory is required for NFS-mounted installation directories. For example, if the installp packages are in /tmp/sys/inst.images, run the command as follows:

```
cd /tmp/sys/inst.images ; inutoc
```

Or

```
inutoc /tmp/sys/inst.images
```

Specifying Paths for Package Installation

Perform the following steps to specify paths for package installation:

1. Log into the Opsware Command Center.
2. From the navigation panel, select Servers ➤ Manage Servers.
3. Navigate to the servers where you are defining a new custom attribute, and select the check box next to the name of each server that you are defining a new custom attribute for.
4. Select the Custom Attributes tab. The Manage Servers: Custom Attributes | [server name] page appears.
5. Click **New** to open a new Custom Attribute form.
6. Enter `OPSWpackage_paths` in the name field. Be sure to use this exact spelling and case for each server for which you specify paths for package installation.
7. In the Value field, enter each path where Opsware SAS should look for the package to install. For example:

   ```
   /mnt/cdrom
   ```

   Or

   ```
   /networkshare/packages/SunOS/5.6/
   ```

```
/mnt/cd0
```

You can specify the Software Repository as a path by using opsware_repository as one of the values. This is useful in cases where you enter a number of pathnames and want to disable the feature temporarily without having to re-enter the pathnames when you are ready to enable the feature again. Enter this value at the top of the list of values.

Uninstalling Software with the Uninstall Software Wizard

The process for uninstalling software is nearly identical to the process of installing software. When you use the Uninstall Software Wizard, you can select any number of servers running the same version of the same operating system to remove software from. (The servers must, however, be assigned to the same customer as the customer associated with the software.)

Perform the following steps to uninstall software with the Uninstall Software Wizard:

1. From the Opsware Command Center home page, click the Uninstall Software link in the Tasks panel.

   Or

   From the navigation panel, click Software ➤ Applications. The Application page appears. Select the application and click **Uninstall**.

2 The Select Servers Page appears. Select the OS version of the servers that you want to uninstall software from by browsing or searching. The type of server is defined by its operating system type and operating system version, as Figure 14-4 shows.

Figure 14-4: OS Versions Page

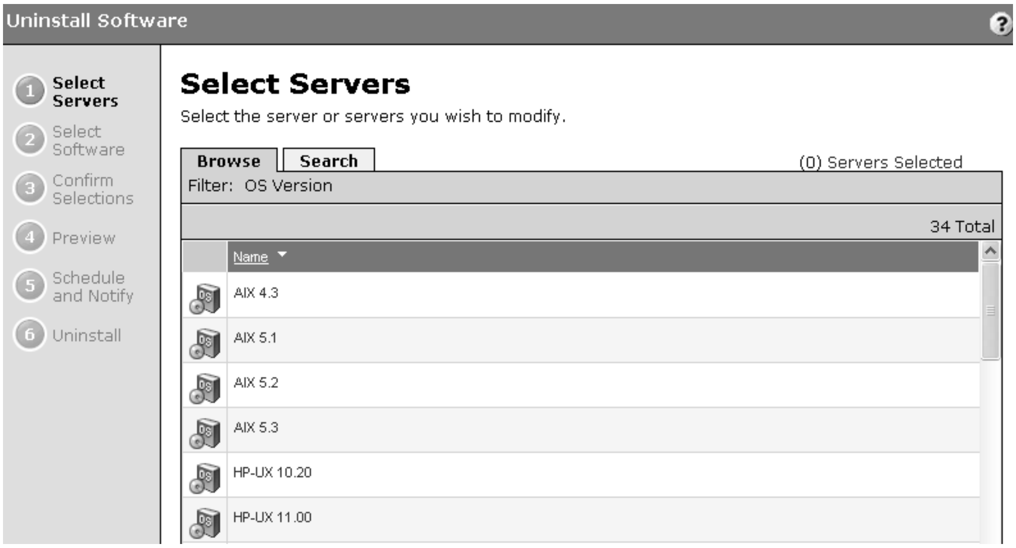

Select Servers

Select the server or servers you wish to modify.

(0) Servers Selected

Filter: OS Version

34 Total

	Name
	AIX 4.3
	AIX 5.1
	AIX 5.2
	AIX 5.3
	HP-UX 10.20
	HP-UX 11.00

After you select the OS version, the Select Servers page appears. The Select Servers page presents a list of all servers of the type that you selected.

3 Select servers or groups and click **Next**. The servers that you select, however, must all have at least one common software node assigned to them in order for you to select software to uninstall.

Alternatively, you can use the search function to find servers that all have a particular package installed on them, as Figure 14-5 shows.

Figure 14-5: Search Servers Page

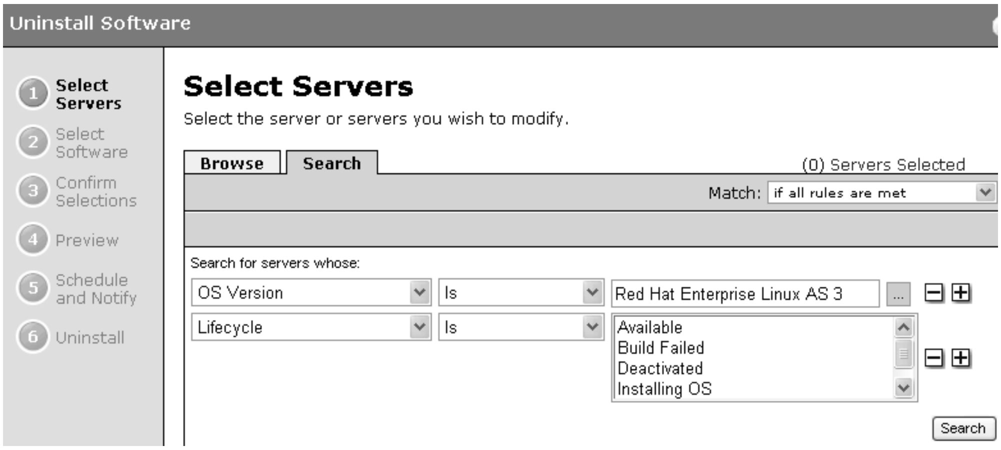

Uninstall Software

Select Servers

Select the server or servers you wish to modify.

(0) Servers Selected

Match: if all rules are met

Search for servers whose:

OS Version | Is | Red Hat Enterprise Linux AS 3

Lifecycle | Is | Available, Build Failed, Deactivated, Installing OS

When searching, the first search condition, OS Version, is preselected and cannot be changed. This is because you already selected the operating system version in the previous step.

After you click **Next**, the Select Software page appears. The Select Software page shows a list of software packages common to all the servers that you selected.

If you selected a group of servers that has no software nodes in common, this list is empty.

4 Select the software that you want to uninstall and click **Next** to continue. The Confirm Selection page appears. The Confirm Selection page displays the servers and software that you selected to uninstall.

5 Review your selection and click **Preview** to continue. The Preview page appears. The Preview page displays a progress bar for each server, which shows the status of the preview process.

Or

Click **Skip Preview** to skip the preview process.

6 (Optional) When the preview process completes, click **View Details** to see what occurs when you uninstall software.

From the summary, you can see what packages will be uninstalled, including packages that you have not selected but that will be uninstalled as a consequence of uninstalling the package that you selected. You can also see if any reboots are required.

7 On the Schedule and Notify page, you have the following options:

- **Schedule**: Click either **Run Now** to execute the operation immediately, or click **Specify Time** to schedule the operation for a later time.
- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.

If you chose to uninstall immediately, progress bars appear that indicate the status of the uninstallation process. You can watch the progress bars to see when the software is uninstalled.

8 (Optional) When the process is complete, click **View Details** to see what packages were uninstalled or removed, and the output, if any, for the uninstallation scripts.

9 Click **Close** to end the wizard.

Closing the wizard does not stop the uninstallation if it is still running. After you close the wizard, you can view the progress of the uninstallation by viewing My Jobs.

See “Server Management Scheduling and Notification” on page 276 in Chapter 7 for information about the My Jobs feature.

Template Installation

This section provides information on how to install templates within Opsware SAS and contains the following topics:

- Overview of Installing Templates
- Installing Templates with the Install Templates Wizard

Overview of Installing Templates

Templates are predefined suites of software that you can install on a large number of servers running the same version of the same operating system in a single operation. Templates can include both applications and patches. Using templates, you can quickly deploy a suite of applications that together provide a particular service, such as a Web server and the applications that run on top of it that provide a Web service.

See the *Opsware® SAS Configuration Guide* for information about how to create templates in Opsware SAS.

You cannot uninstall the packages in a template as a group through an Opsware wizard. You can, however, use the Opsware Uninstall Software Wizard to uninstall the individual packages that a template contains.

Installing Templates with the Install Templates Wizard

Perform the following steps to install a template with the Install Templates Wizard:

1. From the Opsware Command Center home page, click the Install Template link in the Tasks panel.

   Or

   From the navigation panel, click Servers ➤ Manage Servers. The Manage Server page appears. Select the servers and click **Tasks ➤ Install ➤ By Template**.

2. The Select Template page appears, as Figure 14-6 shows.

Figure 14-6: Install Templates Wizard Page

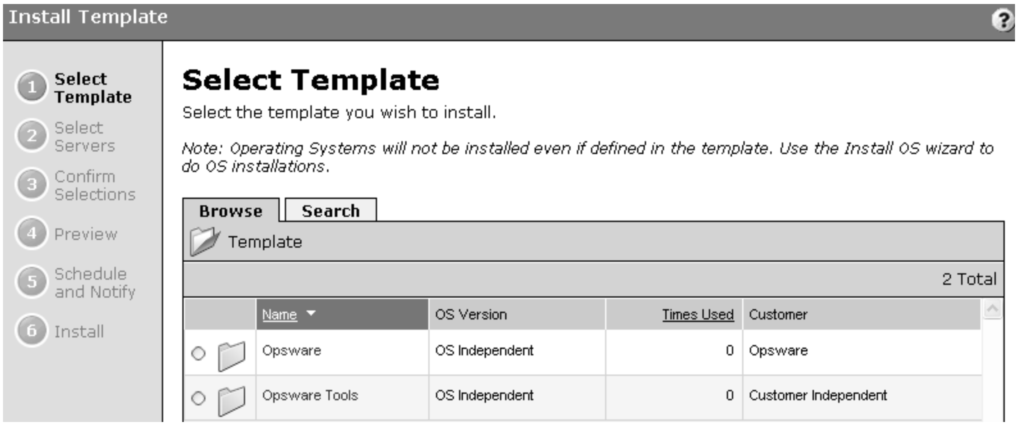

Select Template

Select the template you wish to install.

Note: Operating Systems will not be installed even if defined in the template. Use the Install OS wizard to do OS installations.

	Name ▼	OS Version	Times Used	Customer
○	Opsware	OS Independent	0	Opsware
○	Opsware Tools	OS Independent	0	Customer Independent

3. Select the template that you want to install by browsing or searching. Templates are organized by operating system type and operating system version and by customer. You can select only one template at a time to install. Select the template and click **Next** to continue.

4. The Select Servers page appears. The Select Servers page displays a list of servers limited to the operating system type and version and the customer type of the template that you selected. Select the servers or groups that you want to apply the template to by browsing or searching, and click **Next** to continue.

5. The Confirm Selection page appears. The Confirm Selection page displays the servers and template that you selected to install.

6. Review your selection and click **Preview** to continue. The Preview page appears. The Preview page displays one progress bar per server that shows the status of the preview process.

   Or

   Click **Skip Preview** to skip the preview process.

7. (Optional) When the preview process completes, click **View Details** to see the templates that will be installed. You can also see if any reboots are required.

8. On the Schedule and Notify page, you have the following options:

   - **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.

   - **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

     When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

     - **Option 1**: Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.

- **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

If you choose to install immediately, progress bars appear that indicate the status of the installation process on each server that you selected.

9 (Optional) When the process is complete, click **View Details** to see what templates were installed and the output, if any, from the installation scripts.

10 Click **Close** to end the wizard.

Closing the wizard does not stop the installation if it is still running. After you close the wizard, you can view the progress of the installation by viewing My Jobs.

See "Server Management Scheduling and Notification" on page 276 in Chapter 7 for information about the My Jobs feature.

Chapter 15: Operating System Provisioning

IN THIS CHAPTER

This section discusses the following topics:

- Supported Operating Systems for OS Provisioning
- OS Provisioning
- Hardware Preparation
- New Server Booting
- OS Installation with the Opsware Command Center

Before users can install operating systems on servers with the OS Provisioning feature, the operating systems must be defined in Opsware SAS. The OS media must be made available in Opsware SAS. Additionally, OS definitions must be created in Opsware SAS.

Supported Operating Systems for OS Provisioning

The OS Provisioning feature supports installation of the following versions of Red Hat Linux, Sun Solaris, and Microsoft Windows operating systems:

- Red Hat Linux 6.2
- Red Hat Linux 7.1
- Red Hat Linux 7.2
- Red Hat Linux 7.3
- Red Hat Linux 8.0
- Red Hat Linux Enterprise Linux 2.1 AS
- Red Hat Linux Enterprise Linux 2.1 ES
- Red Hat Linux Enterprise Linux 2.1 WS

- Red Hat Linux Enterprise Linux 3 AS
- Red Hat Linux Enterprise Linux 3 WS
- Red Hat Linux Enterprise Linux 3 ES
- Red Hat Linux Enterprise Linux 4 AS
- Red Hat Linux Enterprise Linux 4 WS
- Red Hat Linux Enterprise Linux 4 ES
- Sun Solaris 7
- Sun Solaris 8
- Sun Solaris 9
- Sun Solaris 10 (See the Note following this list.)
- Fujitsu Solaris 8
- Fujitsu Solaris 9
- Fujitsu Solaris 10
- SUSE Linux Standard Server 8
- SUSE Linux Enterprise Server 8
- SUSE Linux Enterprise Server 9
- Windows NT 4.0
- Windows 2000
- Windows 2003

For Solaris on SPARC, the OS Provisioning feature only supports hardware that is supported by Solaris 10 Update 1, regardless of the version (7, 8, 9, or 10) of Solaris being provisioned.

The OS Provisioning feature works with floppy, CD, or network booting.

The OS Provisioning feature does not provision HP-UX or AIX operating systems. However, you can integrate Opsware SAS with Network Installation Management (NIM) to provision AIX and Ignite-UX to provision HP-UX. See the *Opsware® SAS Configuration Guide* for more information about how to integrate Opsware SAS with HP-UX and AIX OS provisioning systems.

The OS Provisioning feature supports a large variety of hardware models from different manufacturers out of the box. You can also configure the OS Provisioning feature to support additional hardware models. See the *Opsware® SAS Configuration Guide* for more information about how to extend the OS Provisioning feature to support new hardware.

OS Provisioning

This section provides information on OS Provisioning within Opsware SAS and contains the following topics:

- Overview of OS Provisioning
- Server Life Cycle for OS Provisioning

Overview of OS Provisioning

In Opsware SAS, OS Provisioning is installation-based instead of image-based. The OS Provisioning feature uses Red Hat Linux Kickstart, Sun Solaris JumpStart, and Microsoft Windows unattended installation to install operating systems on servers.

The OS Provisioning feature is fully integrated with Opsware SAS; users can install an OS on the following types of servers:

- A bare metal server that does not have an OS installed
- A server that Opsware SAS already manages
- A server that is running in the environment but Opsware SAS does not manage it

The OS Provisioning feature facilitates installing operating systems on servers in the following ways:

- Each OS definition in the OS Provisioning feature contains all the information necessary to build and maintain a server with that OS.

- When installing an OS on a server, the OS Provisioning feature displays information about server hardware and which operating systems are compatible with that hardware architecture.

You need a specific set of feature permissions for OS Provisioning. You'll also need permissions to access the servers associated with customers, facilities, or server groups. To obtain these permissions, contact your Opsware administrator. For more information, see the Permissions Reference appendix in the *Opsware® SAS Configuration Guide*.

Server Life Cycle for OS Provisioning

Opsware SAS is designed to enable multiple teams to work together to provision servers. The OS Provisioning feature allows IT teams to separate the tasks of readying servers for provisioning (such as racking servers, connecting them to power and a network) from provisioning the servers with operating systems.

Someone mounts a new server in a rack and connects it to the Opsware build network. Then they boot the server for the first time by using an Opsware Boot Floppy or CD or by using the network. At a later time, a different system administrator can select the available server from the Server Pool list and provision it with an OS. In the available state, servers do not have an OS installed and might not have access to disk resources.

See Table 15-1 for the life cycle values for servers. During OS provisioning, servers progress through the following Opsware life cycle state changes:

Unprovisioned ➤ Available ➤ Installing OS ➤ Managed

Table 15-1: Opsware SAS *Life Cycle Values for Servers*

OPSWARE LIFE CYCLE VALUE	DESCRIPTION
Server Pool Values	

Table 15-1: Opsware SAS *Life Cycle Values for Servers*

OPSWARE LIFE CYCLE VALUE	DESCRIPTION
Available	Indicates a server on which the OS Build Agent was installed and is running, but an OS has not been installed on the server. The OS Build Agent is a small agent that can run in the memory of the bare metal server. See “Installation of OS Build Agents” on page 581 in this chapter for more information.
Installing OS	Indicates that a user is installing an OS on the server. The server stays in the Server Pool list until the installation process finishes successfully; then, the server moves to the Manage Server list.
Build Failed	Indicates a server on which the OS Build Agent was installed and is running, but the installation of an OS failed. The server will remain in the Server Pool list with this status for 7 days before Opsware SAS deletes the entry. See “Recovering when an OS Installation Fails” on page 592 in this chapter for more information.
Managed Server Values	
Managed	Indicates a server that Opsware SAS is managing. Opsware SAS performs reachability checks for managed servers. After a server reaches this life cycle state, the entry for the server moves from the Server Pool list to the Manage Servers list.
Deactivated	Indicates a server was previously managed by Opsware SAS but is no longer managed by Opsware SAS. However, the server’s history still exists in Opsware SAS. Deactivated servers are not reachable.

OS Provisioning

This section provides information about the OS provisioning process within Opsware SAS and contains the following topics:

- Overview of OS Provisioning

- Solaris OS Provisioning
- Linux OS Provisioning
- Windows OS Provisioning

Overview of OS Provisioning

The process for provisioning new servers of all supported operating systems includes the following steps in the OS Provisioning feature:

1. A system administrator unpacks a server, mounts it in a rack, and attaches the server to power and a network that can reach Opsware SAS.

2. The system administrator prepares the hardware for OS provisioning.

   See "Hardware Preparation" on page 575 in this chapter for more information.

3. If necessary, the system administrator inserts a bootable floppy or CD provided with Opsware SAS. Using a bootable floppy or CD is not necessary for Intel-based servers that support PXE or Unix servers that support DHCP because these types of servers are capable of booting over a network.

   See "New Server Booting" on page 576 in this chapter for more information.

4. The system administrator turns the server on.

   For servers capable of booting over the network, powering the server on causes the server to initiate its network boot process. For example, the server sends a boot request to a PXE server.

   The Opsware OS Build Manager responds to this network boot request by delivering the Opsware OS Build Agent, a small agent that can run in the memory of the bare metal server. (For servers not capable of booting over the network, the Opsware OS Build Agent is on the bootable floppy or CD.)

   The Opsware OS Build Agent constructs an inventory of the server (including server manufacturer, server model, MAC address, available memory, and available storage) and delivers that information to the Opsware OS Build Manager.

5. In the Opsware Command Center, the system administrator sees this server and its hardware inventory in a list of available servers ready to be provisioned.

   See "Verifying Installation of an OS Build Agent" on page 581 in this chapter for more information.

6 The system administrator selects the OS or a complete server baseline (which can include a base OS, a set of OS patches, system utilities, and middleware software) to provision.

The system administrator selects to install the OS or complete server baseline on the server at that time or schedule the installation for some time in the future.

See "Installing an OS by Using a Template" on page 585 in this chapter for more information.

See "Installing an OS by Using a Custom Installation" on page 590 in this chapter for more information.

The OS Provisioning feature installs the selected software onto the server.

7 The system administrator uses Opsware SAS to configure networking for the newly provisioned server.

See "Configuring Networking for an Opsware Managed Server" on page 294 in Chapter 7 for more information.

Additionally, the system administrator might choose to reprovision servers running Red Hat Linux or Sun Solaris operating systems by using the OS Provisioning feature.

See "Reprovisioning a Solaris or Linux Server" on page 596 in this chapter for more information.

Solaris OS Provisioning

The OS Provisioning feature includes a DHCP-based JumpStart configuration that hides the complexity of JumpStart from the end user. Unlike typical JumpStart systems, the OS Provisioning feature does not require configuration updates to the JumpStart server for each installation that you provision.

Instead, an OS definition exists in the OS Provisioning feature for each version of the Solaris OS that you want to install on servers in your environment.

The process for Solaris OS provisioning follows the general OS provisioning process that the OS Provisioning feature established.

See the *Opsware® SAS Configuration Guide* for more information about the detailed steps that occur during the Solaris build process.

Linux OS Provisioning

The OS Provisioning feature includes a Kickstart or YaST2 system that hides the complexity of Kickstart or YAST2 from the end user.

Mapping a specific installation client to a particular Kickstart or YaST2 configuration is a simple procedure in the OS Provisioning feature. The OS Provisioning feature allows users to easily choose a particular Kickstart or YaST2 configuration through the Opsware Command Center at installation time.

The process for Linux OS provisioning follows the general OS provisioning process that the OS Provisioning feature established.

See the *Opsware® SAS Configuration Guide* for more information about the detailed steps that occur during the Linux build process.

Windows OS Provisioning

In the OS Provisioning feature, system administrators can perform unattended, scripted installations of Windows NT, Windows 2000, and Windows 2003 on bare metal servers.

The installation-based approach allows system administrators to adapt to variations in hardware. The OS Provisioning feature can be set up to install Windows operating systems on the known hardware makes and models in the managed environment. At build time, the OS Provisioning feature provisions the server with the correct hardware-specific software and drivers based on the hardware signature of the server about to be provisioned.

The process for Windows OS provisioning follows the general OS provisioning process that the OS Provisioning feature established.

See the *Opsware® SAS Configuration Guide* for more information about the steps that occur during the Windows build process.

Hardware Preparation

Before you use the OS Provisioning feature to install an OS on a server, the server must meet certain requirements, or the hardware must be prepared in certain ways, as Table 15-2 shows.

Table 15-2: Required Hardware Preparation for Servers Managed by Opsware SAS

OPERATING SYSTEM	HARDWARE REQUIREMENTS
Microsoft Windows	Before you install Windows on a server, you need to prepare the hardware by performing the following tasks: • If the hardware has a SCSI RAID controller, you must extend the Windows OS media distribution based on vendor specific requirements. The Windows OS media from Microsoft Corporation does not include the necessary drivers for these SCSI-RAID controllers. • Depending on the version of Windows, create a FAT16 partition or FAT32 on which to install the Windows OS You can create this required partition when using the Windows Boot images or PXE to boot a server the first time. The boot image contains the functionality to create the required partition. See “Booting a Windows or Linux Server with PXE” on page 577 in this chapter for more information. See “Booting a Windows or Linux Server” on page 579 in this chapter for more information.
Sun Solaris	To install Solaris on a server, the hardware must meet the following requirements: • Have a DHCP-capable PROM (older servers can be upgraded to DHCP-capable PROM) • Be part of the sun4u system architecture (platform group) You do not need to perform any Opsware SAS-specific preparation of the hardware before you install Solaris on a server.
Linux	Before you install Linux on a server, you need to prepare the hardware by configuring valid, logical drives for RAID.

New Server Booting

This section provides information on booting new servers with Opsware SAS and contains the following topics:

- Booting New Servers with Different Operating Systems
- OS Build Agent
- Booting a Windows or Linux Server with PXE
- Booting a Windows or Linux Server
- Booting a Solaris Server Over the Network
- Installation of OS Build Agents
- Verifying Installation of an OS Build Agent
- Recovering when an OS Build Agent Fails to Install

Booting New Servers with Different Operating Systems

On Intel-based servers, you can boot a new server over a network in a hands-off fashion by using PXE. For environments with servers that do not support network boot technology, Opsware SAS supports floppy- or CD-based booting.

For Windows and Linux servers, the Opsware Boot Floppy and CD respectively contains a small operating system, network drivers, software required to mount a network drive, and the Opsware OS Build Agent. The Opsware Boot Floppy or CD has the software that is otherwise delivered over the network as part of the network boot process.

For Solaris servers, you can provision an OS over the network by using DHCP.

To boot servers over the network, the installation client must be able to reach the Opsware DHCP server on the Opsware core network. If the installation client is running on a different network than the Opsware core network, your environment must have a DHCP proxy (IP helper). Alternatively, for Linux and Windows installation clients, you can boot the servers by using an Opsware Boot CD or Floppy instead of booting the servers over the network.

OS Build Agent

The OS Provisioning feature de-couples the task of readying a server for provisioning from provisioning the server with an OS. This de-coupling of tasks is possible because of the OS Build Agent.

Booting a new server for the first time installs an OS Build Agent on the server; however, the server does not have the target OS installed and might not have access to disk resources. Opsware SAS can still communicate with the server and perform commands on it remotely because the OS Build Agent is running an OS installed in memory.

The OS Build Agent performs the following functions:

- Registers the server with Opsware SAS when the OS Build Agent starts
- Listens for command requests from Opsware SAS and performs them

The OS Build Agent can perform commands even though the target OS is not installed.

Booting a Windows or Linux Server with PXE

Perform the following steps to boot a Windows or Linux Server with PXE:

1. After you mount the new server in a rack and connect it to the Opsware build network, set up the server to boot by using PXE.

   See the hardware vendor's documentation on how to prepare a server to boot by using PXE.

2. Power on the server and select the option to boot the server with PXE.

   The Opsware SAS menu appears and prompts you to select the type of Opsware Build Agent to install on the server.

   ```
   windows   - Windows Build Agent (DOS 6.22)
   undi      - Windows Build Agent (DOS 6.22 + UNDI)
   win98     - Windows Build Agent (DOS 7.01)
   undi98    - Windows Build Agent (DOS 7.01 + UNDI))
   linux     - Linux Build Agent
   localdisk - Normal boot from localdisk (default after 10 sec)
   ```

   Which version of the Windows Build Agent you should select depends on the type of x86 hardware being provisioned. The images for the Windows Build Agents vary in terms of the memory management software, disk partitioning capabilities, and network drivers – DOS or universal network device interface (UNDI) – that they contain.

For example, if you are provisioning a server that has more than 2GB of RAM, you should select the `Win98` or `Undi98` Boot Image. If an incompatible Boot Image is selected for the hardware, an error message appears at the console. The error message can appear at any point during the provisioning process; for example, it might appear when the Windows Build Agent is booting and DOS is loading or it might appear later in the process when the Windows Installer is loading. See Table 15-3 for the differences between images for the Windows Build Agents.

Table 15-3: Differences Between Images for the Windows Build Agents

BOOT IMAGE	NETWORK DRIVERS	MEMORY MANAGEMENT SOFTWARE	DISK PARTITIONING CAPABILITIES
`windows`	DOS	DOS 6.22	FAT16
`undi`	UNDI	DOS 6.22	FAT16
`win98`	DOS	Windows 98	FAT32
`undi98`	UNDI	Windows 98	FAT32

If you do not select an option after 10 seconds, the server defaults to booting from local disk.

If you select Windows as the option for booting the server, an additional set of Opsware SAS menus appear on the console so that you can partition the hardware disk.

3. For Windows servers only, select the menu choices to partition the disk based on your specifications.

   After the booting process finishes successfully, a message appears on the console that indicates that the server is ready for OS provisioning. An OS Build Agent was installed on the server and the server appears in the Server Pool list in the Opsware Command Center.

4. (Optional) Record the MAC address of the server so that you can locate the server in the Server Pool list in the Opsware Command Center.

You should verify that the newly racked server shows up in the Opsware Command Center and is ready to hand off for OS installation.

See "Verifying Installation of an OS Build Agent" on page 581 in this chapter for more information.

When booting a Linux or Windows server by using PXE, the DHCP relay must be running on the router of the build network for PXE to function properly.

Booting a Windows or Linux Server

You can boot different types of x86 hardware by using an Opsware Boot Floppy (Windows, Windows 98) or by using an Opsware Boot CD (Red Hat Linux or SUSE Linux) because a Boot Floppy or CD can contain multiple NIC drivers.

When you boot a Windows server with a boot floppy, select the Windows or Windows 98 boot floppy based on the server's memory and disk partitioning requirements.

See Table 15-3 for the differences between the Windows and Windows 98 OS build images.

Perform the following steps to boot a Windows or Linux server:

1. After you mount the new server in a rack and connect it to the Opsware build network, insert the Windows Boot Floppy or Linux Boot CD (depending on which OS you want to install on the server).

2. Power on the server. A hardware-vendor specific message appears on the console.

   If you selected Windows as the option for booting the server, Opsware menus appear on the console so that you can partition the hardware disk.

3. For Windows servers only, select the menu choices to partition the disk based on your specifications.

   After the booting process finishes successfully, a message appears on the console that indicates that the server is ready for OS provisioning. An OS Build Agent was installed on the server and the server appears in the Server Pool list in the Opsware Command Center.

4. (Optional) Record the MAC address of the server so that you can locate the server in the Server Pool list in the Opsware Command Center.

You should verify that the newly racked server shows up in the Opsware Command Center and is ready to hand off for OS installation.

See "Verifying Installation of an OS Build Agent" on page 581 in this chapter for more information.

Booting a Solaris Server Over the Network

When Opsware SAS was installed in your facility, the OS Provisioning feature was set up so that the Opsware Boot Server listens for broadcast requests from new servers and it responds by using DHCP.

Perform the following steps to boot a Solaris server over the network:

1. Mount the new Solaris server in a rack and connect it to the network.

   The installation client on this network must be able to reach the Opsware DHCP server on the Opsware core network. If the installation client is running on a different network than the Opsware core network, your environment must have a DHCP proxy (`IP helper`).

2. Enter one of the following commands at the prompt:

   ```
   ok boot net:dhcp - install
   ```

   Or

   ```
   ok boot net:dhcp - install <interface_setting>
   <buildmgr=hostname|IP_address>
   ```

   Where `<interface_`*`setting`*`>` is one of the following options:

   ```
   autoneg, 100fdx, 100hdx, 10fdx, 10hdx
   ```

   You can include an interface setting with the boot command to set the network interface to a specific speed and duplex during OS provisioning. When Opsware SAS was installed in the local facility, a default value was provided for this interface setting. Specifying this boot argument allows you to override the default interface setting.

   To continue setting the network interface with a specific speed and duplex, you can use a variety of methods, including using a Solaris build customization script or specifying the values in a Solaris Package or RPM in the OS media.

   See the *Opsware® SAS Configuration Guide* for more information.

Ways that the OS Build Agent Locates the Opsware Build Manager

For Solaris OS provisioning, the JumpStart build script runs the OS Build Agent, which contacts the Opsware Build Manager (via the Agent Gateway in the core). The Solaris `begin` script attempts to locate the Opsware Build Manager in the following ways:

- By using information that the Opsware DHCP server provided
- By looking for the host name `buildmgr` in DNS as configured by the DHCP server

You can override the way that the OS Build Agent contacts the Opsware Build Manager by specifying a boot argument at the prompt when you boot a new Solaris server:

```
ok boot net:dhcp - install [buildmgr=hostname|IP_address]:port
```

Installation of OS Build Agents

After you install an OS Build Agent on a server by booting the server with PXE or an Opsware Boot Image (Windows and Linux) or by using the network (Solaris), the server appears in the Server Pool list.

You should verify that the newly racked server shows up in the Opsware Command Center and is ready to hand off for OS installation.

The Server Pool list displays the servers that have registered their presence with Opsware SAS but do not have the target OS installed on them. From here, you can install an OS by selecting the server and clicking **Install OS**.

Verifying Installation of an OS Build Agent

Perform the following steps to verify the installation of an OS build agent:

1. Log into the Opsware Command Center.
2. From the navigation panel, click Servers ➤ Server Pool. The Server Pool page appears, as Figure 15-1 shows.

Figure 15-1: Server Pool List in the Opsware Command Center

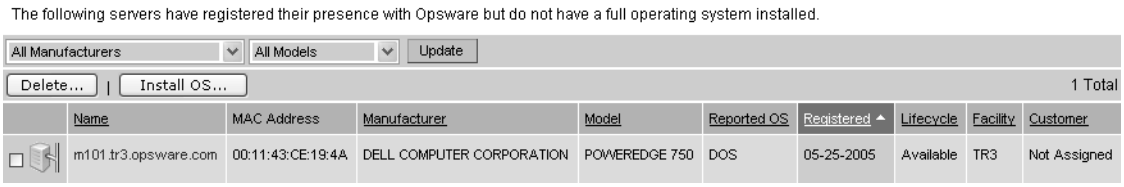

The following servers have registered their presence with Opsware but do not have a full operating system installed.

1 Total

	Name	MAC Address	Manufacturer	Model	Reported OS	Registered ▲	Lifecycle	Facility	Customer
□	m101.tr3.opsware.com	00:11:43:CE:19:4A	DELL COMPUTER CORPORATION	POWEREDGE 750	DOS	05-25-2005	Available	TR3	Not Assigned

3. (Optional) From the drop-down lists, select the manufacturer, model, or facility of the server that you want to verify and click **Update**.
4. For Intel x86 processor-based servers, locate the MAC address of the server that you just booted.

   Or

   For Sun SPARC processor servers, locate the chassis ID of the server that you just booted.

The chassis ID for Sun SPARC processor servers appears in the MAC Address column in the Server Pool list.

The Life cycle column indicates the progress or success of the OS Build Agent installation. If the OS Build Agent was successfully installed, the Life cycle column indicates that the server is available for OS provisioning.

See "Server Life Cycle for OS Provisioning" on page 570 in this chapter for more information.

Recovering when an OS Build Agent Fails to Install

When an OS Build Agent fails to install on a server, the server does not appear in the Server Pool list.

You can check the server console for error messages and try to boot the server again with PXE or by using the Opsware Boot Floppy or CD.

If all errors were successfully resolved, the initial boot occurs, the OS Build Agent is installed on the server, the server appears in the Server Pool list, and the Life cycle column indicates that the server is available.

If you are unable to resolve the error condition and install the OS Build Agent on the server so that it appears in the Server Pool list, contact your Opsware administrator for troubleshooting assistance.

OS Installation with the Opsware Command Center

This section provides information on OS installation with the Opsware Command Center and contains the following topics:

- Servers and the Server Pool
- Ways to Install Operating Systems on Servers
- Installing an OS by Using a Template
- Installing an OS by Using a Custom Installation
- Recovering when an OS Installation Fails
- Network Configuration for Servers after OS Provisioning
- Requirements for Reprovisioning Solaris and Linux Servers
- Reprovisioning a Solaris or Linux Server

Servers and the Server Pool

You begin the OS provisioning process by reviewing the servers in the Server Pool list. Servers in the Server Pool have registered their presence with Opsware SAS but do not have the target OS installed. From there, you can install an OS by clicking **Install OS**. See Figure 15-2.

Figure 15-2: Server Pool List in the Opsware Command Center

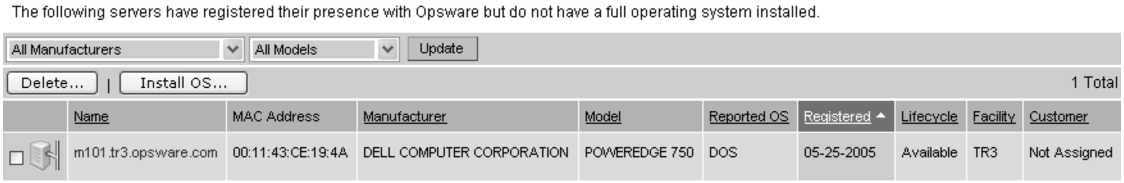

The following servers have registered their presence with Opsware but do not have a full operating system installed.

	Name	MAC Address	Manufacturer	Model	Reported OS	Registered ▲	Lifecycle	Facility	Customer
□	m101.tr3.opsware.com	00:11:43:CE:19:4A	DELL COMPUTER CORPORATION	POWEREDGE 750	DOS	05-25-2005	Available	TR3	Not Assigned

The Server Pool provides the following information about each server waiting to be provisioned with the target OS:

- The host name set by booting the server the first time over the network or by using an Opsware Boot Floppy or CD
- The MAC address or chassis ID
- The manufacturer and model of the server
- The OS of the OS Build Agent (Windows, Red Hat Linux, or Solaris)

  You use this information to select the target OS for servers.
- The life cycle associated with the server
- The customer and facility association

- Additional hardware information. (Clicking the server name opens a window that displays specific hardware information, as Figure 15-3 shows.)

Figure 15-3: Information Displayed in the Edit Server Page for a Server in the Server Pool

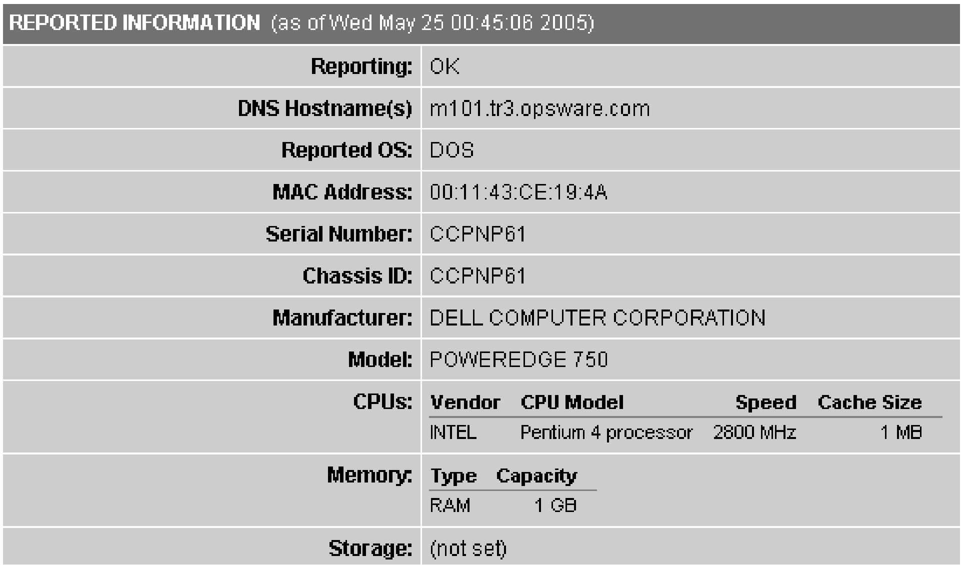

REPORTED INFORMATION (as of Wed May 25 00:45:06 2005)	
Reporting:	OK
DNS Hostname(s)	m101.tr3.opsware.com
Reported OS:	DOS
MAC Address:	00:11:43:CE:19:4A
Serial Number:	CCPNP61
Chassis ID:	CCPNP61
Manufacturer:	DELL COMPUTER CORPORATION
Model:	POWEREDGE 750
CPUs:	**Vendor** **CPU Model** **Speed** **Cache Size** INTEL Pentium 4 processor 2800 MHz 1 MB
Memory:	**Type** **Capacity** RAM 1 GB
Storage:	(not set)

Ways to Install Operating Systems on Servers

You can install an OS on a server by using one of the following methods:

- Selecting a pre-defined template, which is a pre-packaged collection of installable software

  The template includes the base operating system, and can include software to provision the entire software stack, such as the latest set of operating system patches, system utilities (SSH or the latest JVMs), middleware including databases, Web servers, and application servers, and so on, up to the custom business applications the server ultimately runs.

Figure 15-4: Example of an OS Provisioning Baseline

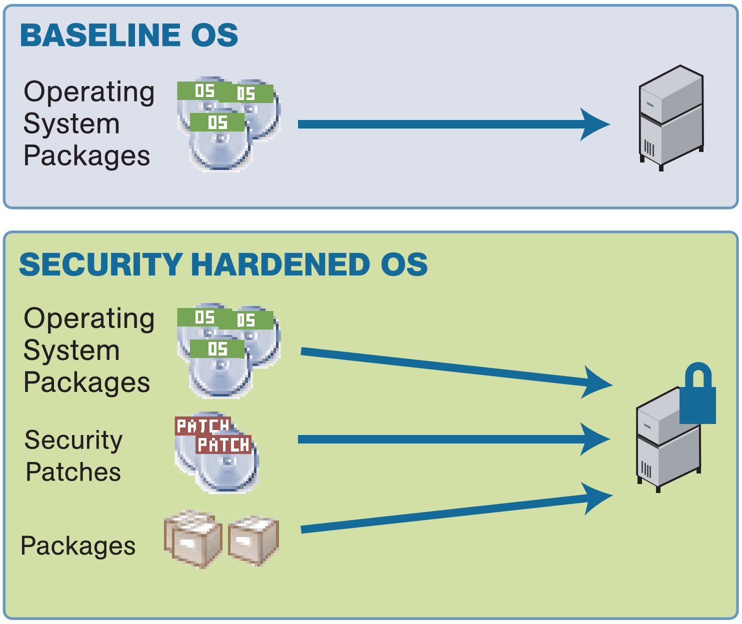

- Performing a custom installation, which includes defining the installation on-the-fly by selecting an OS definition, patches, and other applications to install

  After performing a custom installation, you can then save the selections in a new template for later use on other servers.

After the installation has begun on a set of servers, you can view progress or results either in the Opsware wizard itself or through the My Jobs interface.

See “Templates and Server Management” on page 75 in Chapter 1 for information about how Opsware SAS works with templates.

See the *Opsware® SAS Configuration Guide* for more information about how to create and manage templates.

Installing an OS by Using a Template

During OS installation, you cannot select a server from the Server Pool list that has the status Installing OS.

Perform the following steps to install an OS by using a template:

1. From the Opsware Command Center home page, click the Install OS link in the Tasks panel.

   Or

   From the navigation panel, click Servers ➤ Server Pool. The Server Pool page appears. Select the servers that you want to provision and click **Install OS**.

   Or

   From the navigation panel, click Software ➤ Operating Systems. the Operating System page appears. Click the Install OS link.

   When you select multiple servers to provision, you must select servers with similar hardware architecture – x86-processor-based hardware or SPARC. If you select servers that have different hardware architecture, an error message appears.

   The Install Operating System Wizard appears. See Figure 15-5.

Figure 15-5: Overview Page in the Install Operating System Wizard

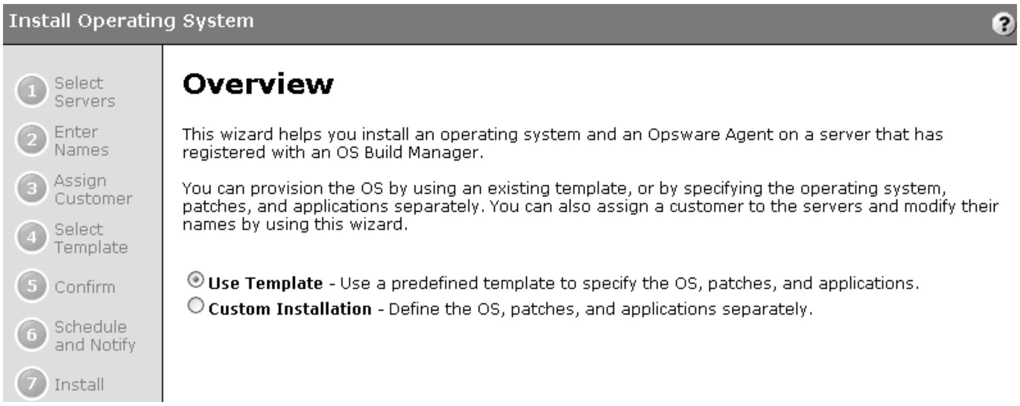

Overview

This wizard helps you install an operating system and an Opsware Agent on a server that has registered with an OS Build Manager.

You can provision the OS by using an existing template, or by specifying the operating system, patches, and applications separately. You can also assign a customer to the servers and modify their names by using this wizard.

● **Use Template** - Use a predefined template to specify the OS, patches, and applications.

○ **Custom Installation** - Define the OS, patches, and applications separately.

2. If necessary, select the Use Template option and click **Start**. If you did not select servers in Step 1, the Select Servers page appears.

3. If prompted, select the servers or groups that you want to provision and click **Next**. You can find the servers that you want to provision by browsing the list or by searching.

   Servers in the Server Pool list waiting to be provisioned are identified by MAC address or chassis ID because they are still using DHCP addresses.

You must select servers with similar hardware architecture (based on the value in the Reported OS column in the Server Pool list) or an error message appears.

4 Enter a name for each server and click **Next**. The Assign Customer page appears.

By default, the OS Provisioning feature entered the server host name in this field. You can enter new names that adequately describe each server.

The name that you enter appears as the display name for the servers in the Opsware Command Center UI.

5 Select a customer for the servers and click **Next**. The Select Template page appears.

In the Assign Customers page, you only see customers listed that you have the permission to access with your user account. Additionally, the customer that you select controls which templates you can use to install an OS and software on the servers. Depending on the customer that you select, you see only the templates associated with that customer, that are customer independent, or are not assigned to a customer.

See “Customer Accounts in Opsware SAS” on page 146 in Chapter 4 for information about the distinction between the customers Customer Independent and Not Assigned.

6 Select the template to use to provision the servers and click **Next**. Templates appear in the list only if they meet these requirements:

- The templates were created for the same OS as the servers that are being provisioned.
- The templates contain OS software to install on servers.

  See “Templates and Server Management” on page 75 in Chapter 1 for information about how Opsware SAS works with templates.

  See the *Opsware® SAS Configuration Guide* for more information about how to create and manage templates.

To view the details about the templates before you select one, click the template name. A window appears that displays general information about the template, as Figure 15-6 shows.

Figure 15-6: Template Details Page in the Install Operating System Wizard

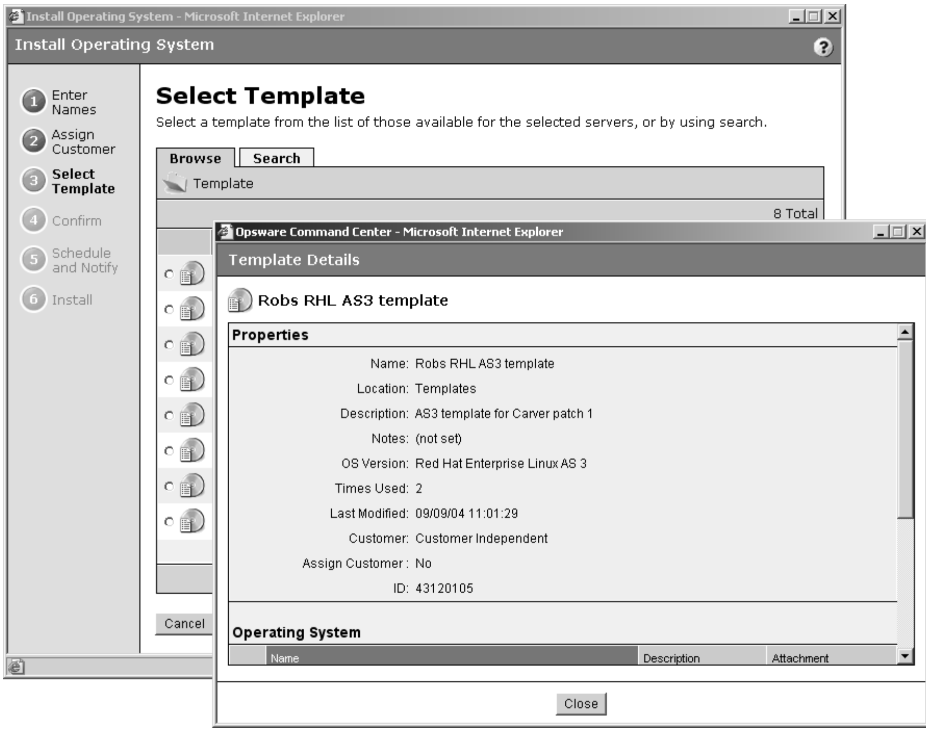

Name:	Robs RHL AS3 template
Location:	Templates
Description:	AS3 template for Carver patch 1
Notes:	(not set)
OS Version:	Red Hat Enterprise Linux AS 3
Times Used:	2
Last Modified:	09/09/04 11:01:29
Customer:	Customer Independent
Assign Customer :	No
ID:	43120105

The Select Template page only displays the templates available for the OS that you chose to install.

To view detailed information about the OS that the template installs, such as the settings in the configuration file, click Software ➤ Operating Systems from the navigation panel, then select the Installation tab. The installation resources for the OS appear. You can view information about the OS installation, such as the contents of the configuration file.

After you click **Next**, a confirmation page appears that shows details about the template that will be installed on the servers.

7 On the Schedule and Notify page, you have the following options:

- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.
- **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

  When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have "Allow Run Refresh Jobs" permissions, you will see additional options. Select either of the following options:

  - **Option 1**: Membership is determined based on the "Time of Confirm Selection." Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
  - **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

> The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

If you select to run the job at that time, a progress bar appears that shows the progress of the OS installation.

8. (Optional) When the installation finishes, click **View Details** to see progress or the results of the installation.
9. Click **Close** to end the wizard.

   Closing the wizard does not stop the installation if it is still running. After you close the wizard, you can view the progress of the installation by viewing My Jobs.

See "Server Management Scheduling and Notification" on page 276 in Chapter 7 for information about the My Jobs feature.

Installing an OS by Using a Custom Installation

During OS installation, you cannot select a server from the Server Pool list that has the status Installing OS.

Perform the following steps to install an OS by using a custom installation:

1. From the Opsware Command Center home page, click the Install OS link in the Tasks panel.

   Or

   From the navigation panel, click Servers ➤ Server Pool. The Server Pool page appears. Select the servers that you want to provision and click **Install OS**.

   Or

   From the navigation panel, click Software ➤ Operating Systems. the Operating System page appears. Click the Install OS link.

   You must select servers with similar hardware architecture (x86-processor-based hardware or SPARC) or an error message appears.

   The Install Operating System Wizard appears, as Figure 15-7 shows.

Figure 15-7: Overview Page in the Install Operating System Wizard

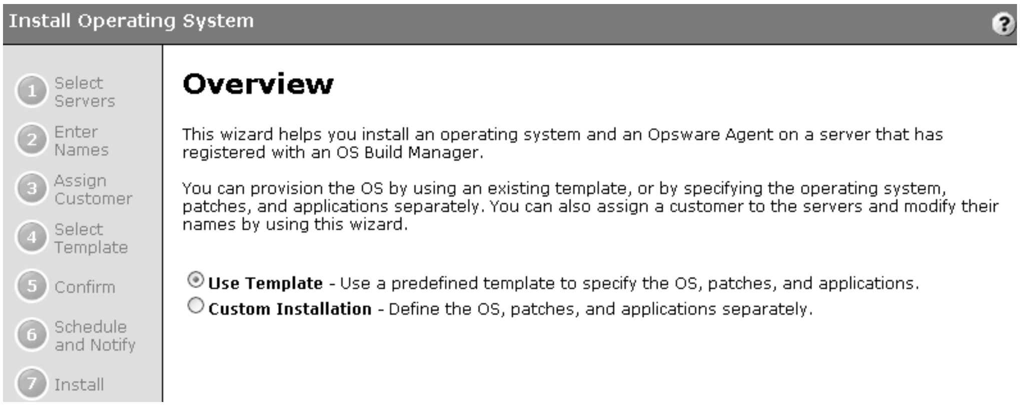

Overview

This wizard helps you install an operating system and an Opsware Agent on a server that has registered with an OS Build Manager.

You can provision the OS by using an existing template, or by specifying the operating system, patches, and applications separately. You can also assign a customer to the servers and modify their names by using this wizard.

● **Use Template** - Use a predefined template to specify the OS, patches, and applications.

○ **Custom Installation** - Define the OS, patches, and applications separately.

2. Select the Custom Installation option and click **Start**. If you did not select servers in Step 1, the Select Servers page appears.

3. If necessary, select the servers or groups that you want to provision and click **Next**. The Select Operating System page appears.

   You must select servers with similar hardware architecture (based on the value in the Reported OS column in the Server Pool list) or an error message appears.

4 Enter a name for each server and click **Next**. The Assign Customer page appears.

By default, the OS Provisioning feature entered the server host name in this field. You can enter new names that adequately describe each server.

The name that you enter appears as the display name for the servers in the Opsware Command Center UI.

5 Select a customer for the servers and click **Next**.

The customer that you select controls which operating systems and applications you can select to install on the servers. Depending on the customer that you select, you see only the operating systems and applications associated with that customer, that are customer independent, or are not assigned to a customer. Patches are always customer independent.

6 Select the OS for the servers and click **Next**.

7 Select the OS patches that you want to apply to the servers and click **Next**. The Select Applications page appears.

See “Patch Management for Windows” on page 453. or “Patch Management for Unix” on page 519 for information on how Opsware SAS performs patch management.

8 Select any applications that you want to install on the servers and click **Next**.

9 On the Schedule and Notify page, you have the following options:

- **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.
- **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

  When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have “Allow Run Refresh Jobs” permissions, you will see additional options. Select either of the following options:

- **Option 1**: Membership is determined based on the "Time of Confirm Selection." Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
- **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

If you select to run the job at that time, a progress bar appears that shows the progress of the OS installation.

10 (Optional) When the installation finishes, click **View Details** to see progress or the results of the installation.

11 Click **Close** to end the Wizard.

Closing the Wizard does not stop the installation if it is still running. After you close the Wizard, you can view the progress of the installation by viewing My Jobs.

See "Server Management Scheduling and Notification" on page 276 in Chapter 7 for information about the My Jobs feature.

Recovering when an OS Installation Fails

Servers waiting to be provisioned appear in the Server Pool list with the status available. When an OS installation fails for an unprovisioned server, the server status changes to Build Failed.

Perform the following steps to recover when an OS installation fails:

1 From the Opsware Command Center home page, click the link for the failed installation in the My Jobs panel.

Or

From the navigation panel, click My Jobs. The My Jobs page appears and displays the operations that you performed with Opsware SAS, as Figure 15-8 shows.

Figure 15-8: My Jobs Panel in the Opsware Command Center Home Page

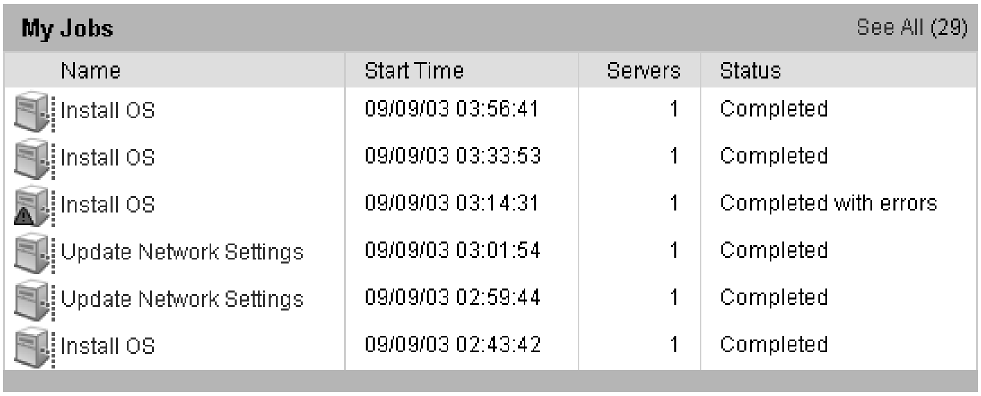

My Jobs See All (29)

Name	Start Time	Servers	Status
Install OS	09/09/03 03:56:41	1	Completed
Install OS	09/09/03 03:33:53	1	Completed
⚠ Install OS	09/09/03 03:14:31	1	Completed with errors
Update Network Settings	09/09/03 03:01:54	1	Completed
Update Network Settings	09/09/03 02:59:44	1	Completed
Install OS	09/09/03 02:43:42	1	Completed

2 Locate the failed OS installation that you initiated and click the link for the job to open a window for the job. An error message appears in the window.

3 Click **View Details** to see detailed information about the job. The My Jobs information contains a build log for the OS installation. This build log contains any error messages that the OS Provisioning feature generated, as Figure 15-9 shows.

Figure 15-9: Details Page in the My Job Window

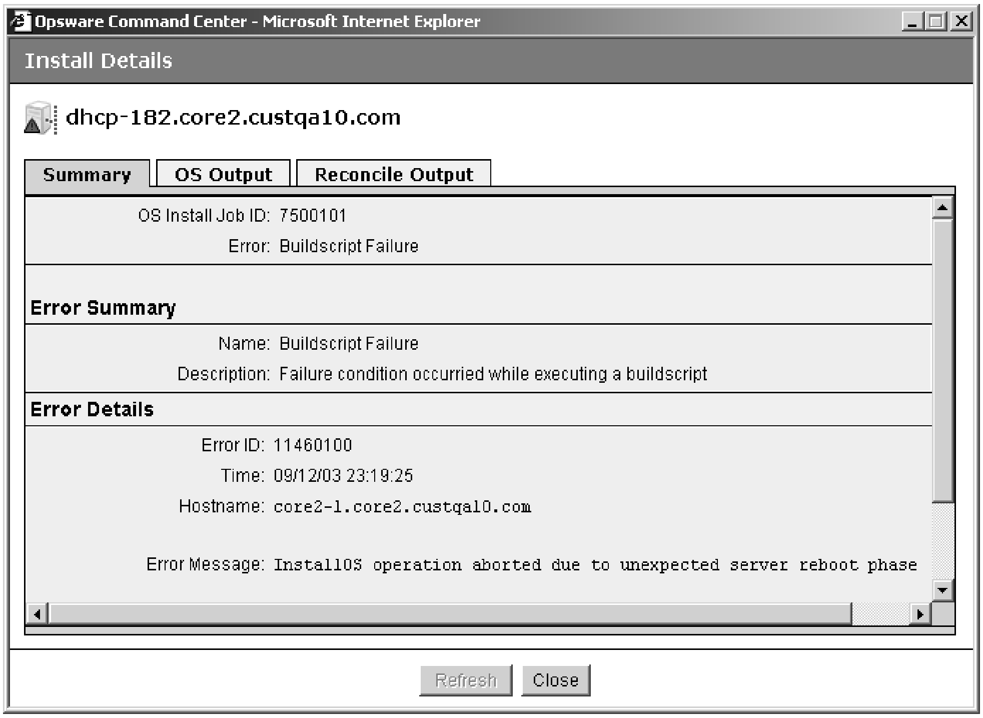

dhcp-182.core2.custqa10.com

OS Install Job ID: 7500101
Error: Buildscript Failure

Error Summary

Name: Buildscript Failure
Description: Failure condition occurried while executing a buildscript

Error Details

Error ID: 11460100
Time: 09/12/03 23:19:25
Hostname: `core2-1.core2.custqa10.com`

Error Message: `InstallOS operation aborted due to unexpected server reboot phase`

4 Review and fix any errors that occurred during the build process.

5 (Optional) Delete the failed OS installation from the Server Pool list by selecting the server and clicking **Delete**.

You can leave the entry for the failed OS installation in the Server Pool list because Opsware SAS automatically replaces the entry when you reboot the server.

Alternatively, Opsware SAS removes the entry from the list after 7 days and all information about the server is removed from Opsware SAS.

6 Reboot the server by using the network (PXE for Windows and Linux servers or DHCP for Solaris servers) or by using an Opsware Boot Floppy (for Windows) or by using an Opsware Boot CD (for Linux servers). If successful, the server appears in the Server Pool list with the status available.

7 Install the OS on the server by using a template or custom installation.

See “Installing an OS by Using a Template” on page 585 in this chapter for more information.

See “Installing an OS by Using a Custom Installation” on page 590 in this chapter for more information.

If all errors were successfully resolved, the OS is installed on the server and the server moves from the Server Pool list to the Manage Server list.

See “Server Management Scheduling and Notification” on page 276 in Chapter 7 for information about using My Jobs to obtain a history of your operations.

Network Configuration for Servers after OS Provisioning

The OS Provisioning feature provisions servers with an OS by using DHCP addresses. By using the Opsware Command Center, users can configure network settings, including a static IP address, host name, default gateway, DNS server addresses, subnet masks, and so on, after the base operating system is installed on servers.

Because DHCP servers often assign temporary IP addresses to servers that boot over a network, system administrators typically need to assign static IP addresses (and other network properties) before the servers can be put into service. Opsware SAS enables system administrators to do this through the Opsware Command Center rather than logging onto the server manually after OS provisioning is complete.

See “Configuring Networking for an Opsware Managed Server” on page 294 in Chapter 7 for information about how Opsware SAS configures networking for servers.

Requirements for Reprovisioning Solaris and Linux Servers

When you choose to preserver network configuration for a server you are about to reprovision, the following requirements apply:

- The server must be in a DHCP-enabled network when the reprovisioning begins.
- The server must be in its original network after the OS installation is complete.

Because of these requirements, re-provisioning Solaris and Linux servers functions best when the build and production networks are overlaid (namely, the production network is running a DHCP server).

Reprovisioning a Solaris or Linux Server

You can reprovision servers built or managed by Opsware SAS with this feature. You can reprovision Solaris and Linux servers so that they are running another version of the same OS if the hardware supports that new version of the OS.

You cannot reprovision a Linux server so that it runs a Windows OS.

Perform the following steps to reprovision a Solaris or Linux server:

1. From the navigation panel, click Servers ➤ Manage Servers. The Manage Servers page appears.

2. Select the servers that you want to reprovision from the list.

   You must select servers with similar hardware architecture (x86-processor-based hardware or SPARC) or an error message appears.

   The check boxes for servers running the Windows OS are unavailable.

3. From the **Tasks** menu, choose **Install ➤ Operating System**.

   The Install Operating System Wizard appears. The wizard contains a warning that you are about to reprovision the servers. See Figure 15-10.

Figure 15-10: Overview Page for the Install Operating System Wizard When Re-Provisioning

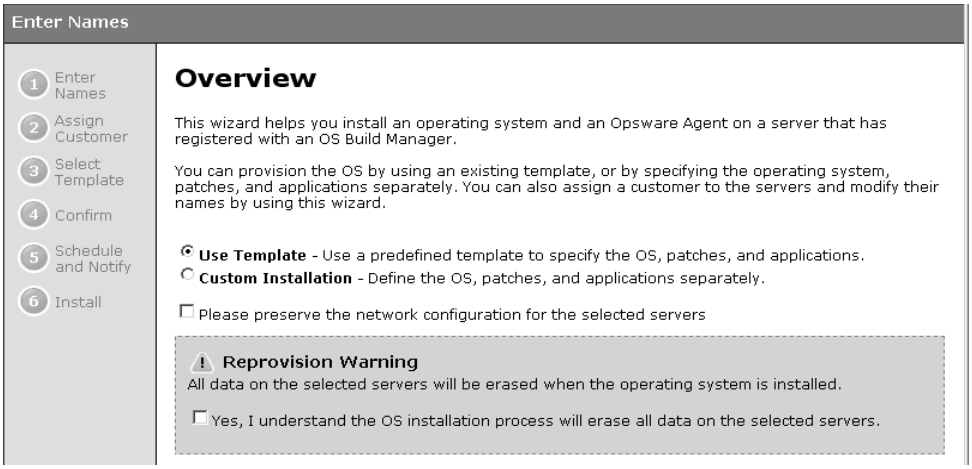

Overview

This wizard helps you install an operating system and an Opsware Agent on a server that has registered with an OS Build Manager.

You can provision the OS by using an existing template, or by specifying the operating system, patches, and applications separately. You can also assign a customer to the servers and modify their names by using this wizard.

● **Use Template** - Use a predefined template to specify the OS, patches, and applications.
○ **Custom Installation** - Define the OS, patches, and applications separately.

□ Please preserve the network configuration for the selected servers

Reprovision Warning
All data on the selected servers will be erased when the operating system is installed.

□ Yes, I understand the OS installation process will erase all data on the selected servers.

4. Select the check box in the warning (indicating that you understand that Opsware SAS erases all data from the servers).

5. (Optional) Select the check box to preserve the network settings for the server.

   See "Configuring Networking for an Opsware Managed Server" on page 294 in Chapter 7 for information about how Opsware SAS configures networking for servers.

6. Select the Use Template option or the Custom Installation option and click **Start**.

   The installation proceeds normally through the Install Operating System Wizard.

7. Complete the installation by using the Install Operating System Wizard.

   See "Installing an OS by Using a Template" on page 585 in this chapter for more information.

   See "Installing an OS by Using a Custom Installation" on page 590 in this chapter for more information.

Boot Arguments for Reprovisioning Solaris Servers

When you reprovision a Solaris server, the OS Provisioning feature reboots the server automatically. Therefore, you cannot provide boot arguments at the prompt (for example, to specify the interface setting or the location of the Opsware Build Manager).

Alternatively, you can provide a boot argument by including a custom attribute for the server named `reboot_command` that has the value for the interface setting that you want to use.

See "Booting a Solaris Server Over the Network" on page 580 in this chapter for information about how to use interface settings when you boot a Solaris server.

Chapter 16: Application Configuration Management

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Application Configuration Management (ACM)
- Application Configuration Creation and Use
- ACM Components
- Application Configuration Inheritance
- Sequence Merging and Inheritance
- Using ACM

Overview of Application Configuration Management (ACM)

Opsware Application Configuration Management (ACM) enables you to create templates that help manage configuration files associated with applications. Using ACM, you can manage and update application configuration files from a central location. This ensures that applications in your facility are accurately and consistently configured.

For example, you can create an Application Configuration and set its values, and then push those values to all instances of that application in your facility, whether that application resides on a single server or on groups of servers. You can also check live servers against your Application Configuration and view any differences between the desired state of the application’s configuration and the actual state of the application’s configuration. If you would like to make a change, simply edit the values and push the changes.

In addition, ACM supports rollback. Because ACM creates a record of the application configuration before the configuration change is made, you can rollback to the original Application Configuration.

ACM also allows you to configure different instances of the same application in your facility. Because an Application Configuration can be attached to several application instances across multiple servers, you can modify default values by customer and facility. For example, you can create default application configuration values across your entire facility, and then make changes to specific instances of the application configuration contained in different facilities and for specific customers.

Application Configuration Creation and Use

Using an Application Configuration enables you to manage and modify configuration files for applications on your Opsware-managed servers. The process of using ACM follows these general steps:

1. **Determine which applications you want to manage**: Your first step is to choose applications to manage. For example, for the iPlanet Web server, you might want to manage the following configuration files: password.conf, obj.conf, mimetypes, and magnus.conf. To manage these iPlanet configuration files with ACM, you need to make templates out of each configuration file.
2. **Create CML files**: For each application file, create a CML file based upon the actual configuration file you want to manage.
3. **Create configuration templates**: Once you have created all of your CML files from the configuration files, create a Application Configuration Template for each CML file inside the OCC Client.
4. **Create application configuration to hold templates**: Once all the configuration files associated with an application have Application Configuration Templates, add them to an Application Configuration. An Application Configuration is a container that houses multiple Application Configuration Templates.
5. **Set the default values**: Next, set the Application Configuration's default values at various levels in the Application Configuration hierarchy, such as at the customer or facility level, or individually at the application instance level on a server.
6. **Attach application configuration to a server (or group)**: Once you have created and configured your Application Configuration, attach it to each server (or server group) where you are managing application files.

7. **Compare the actual configuration files with the configuration template**: You can easily compare a Application Configuration Template with the actual configuration file on the server and see if any changes have been made. This comparison shows manually changed configuration files or configuration values that have been changed, but not pushed.
8. **Push changes**: No changes are made to the actual configuration files on the server until you push those changes to the server where the Application Configuration files are stored. Application configuration changes can be pushed to individual servers or groups of servers

Figure 16-1: Application Configuration Creation and Usage Process

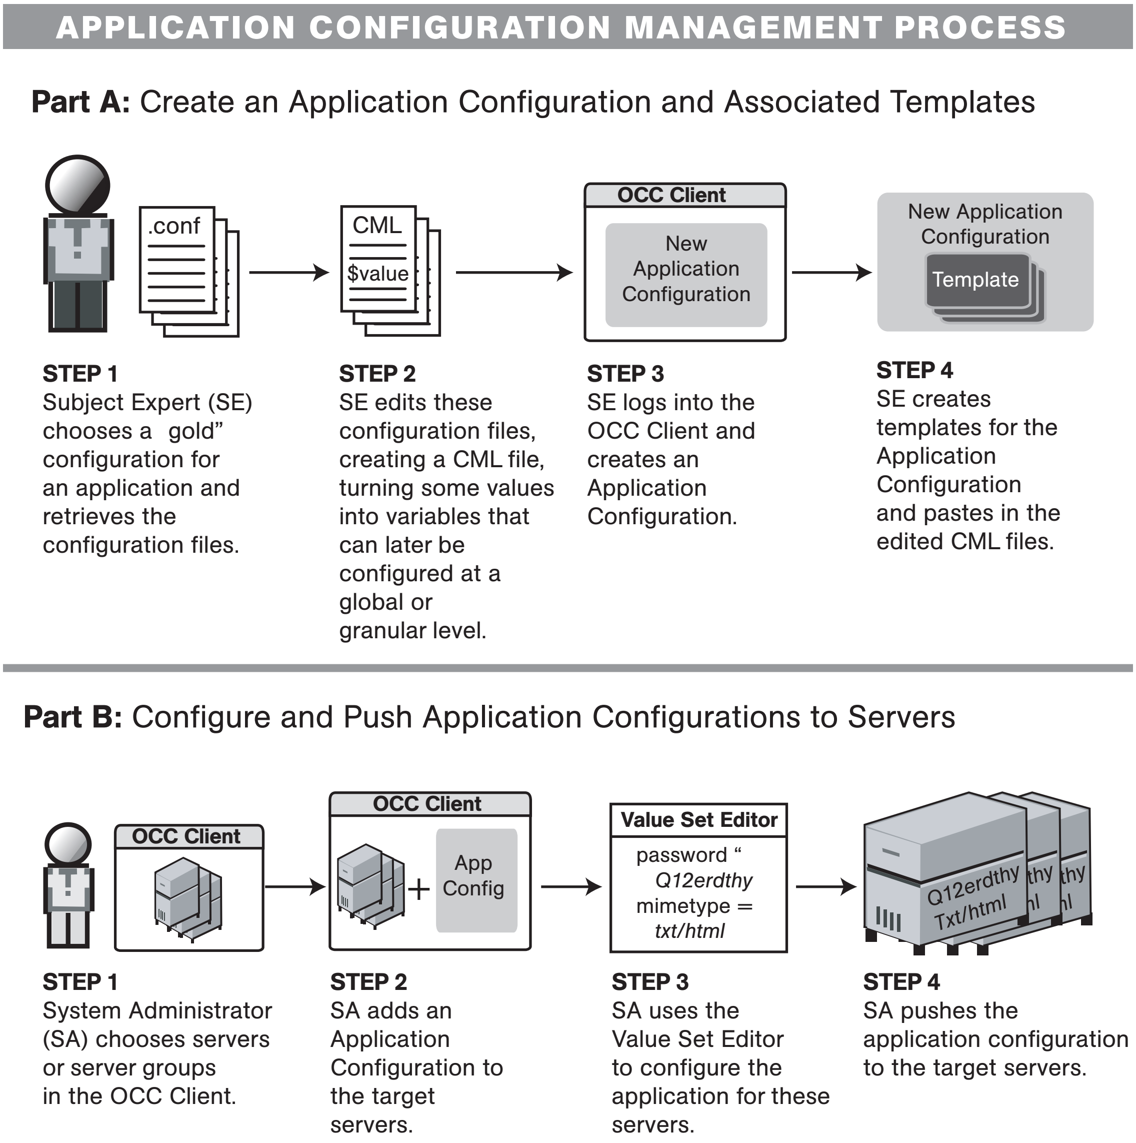

APPLICATION CONFIGURATION MANAGEMENT PROCESS

Part A: Create an Application Configuration and Associated Templates

STEP 1
Subject Expert (SE) chooses a gold" configuration for an application and retrieves the configuration files.

STEP 2
SE edits these configuration files, creating a CML file, turning some values into variables that can later be configured at a global or granular level.

STEP 3
SE logs into the OCC Client and creates an Application Configuration.

STEP 4
SE creates templates for the Application Configuration and pastes in the edited CML files.

Part B: Configure and Push Application Configurations to Servers

STEP 1
System Administrator (SA) chooses servers or server groups in the OCC Client.

STEP 2
SA adds an Application Configuration to the target servers.

STEP 3
SA uses the Value Set Editor to configure the application for these servers.

STEP 4
SA pushes the application configuration to the target servers.

ACM Components

Application Configuration Management consists of the following main components:

- Configuration Template
- Application Configuration

- Value Set Editor
- Configuration Markup Language (CML)

Configuration Template

An Application Configuration Template is set of values that represent the configuration file of an application. Using the ACM tool in the OCC Client, you can edit the values in the configuration template and push those changes to the actual configuration file on the server.

Using Opsware's Configuration Markup Language (CML), an application expert modifies an application's configuration file and turns it into a Application Configuration Template. In this form, Opsware SAS can then make changes to the actual configuration file on the server. When you make changes to the Application Configuration Template and then push the changes to a server, ACM replaces a section of text in the configuration template with the desired value.

Application Configuration

In many cases, an application has multiple configuration files. For each application managed by ACM, create an Application Configuration that holds all Application Configuration Templates associated with the application. The application configuration aggregates all those templates in a single location.

In addition to configuration templates, an Application Configuration can be configured to contain and execute pre/post configuration scripts.

Value Set Editor

The Value Set Editor enables you to specify the values for each configuration file. Each entry inside a configuration file is represented inside the value set editor as an element, which consists of a name-value pair. The entire collection of elements in a configuration file is referred to as the configuration file's value set – that is, all the elements and their names and values in the file.

You can edit value set elements for an application configuration at two of the following levels:

- **Default Values Level**: The value set elements you edit at this level are applied across all instances of the application that the application configuration is attached to. (These can, however, be overridden by customer or facility.) You access the value set editor at the default level by selecting the Application Configuration feature from inside the OCC Client and double-clicking an Application Configuration.

Figure 16-2: Application Configuration Default Values

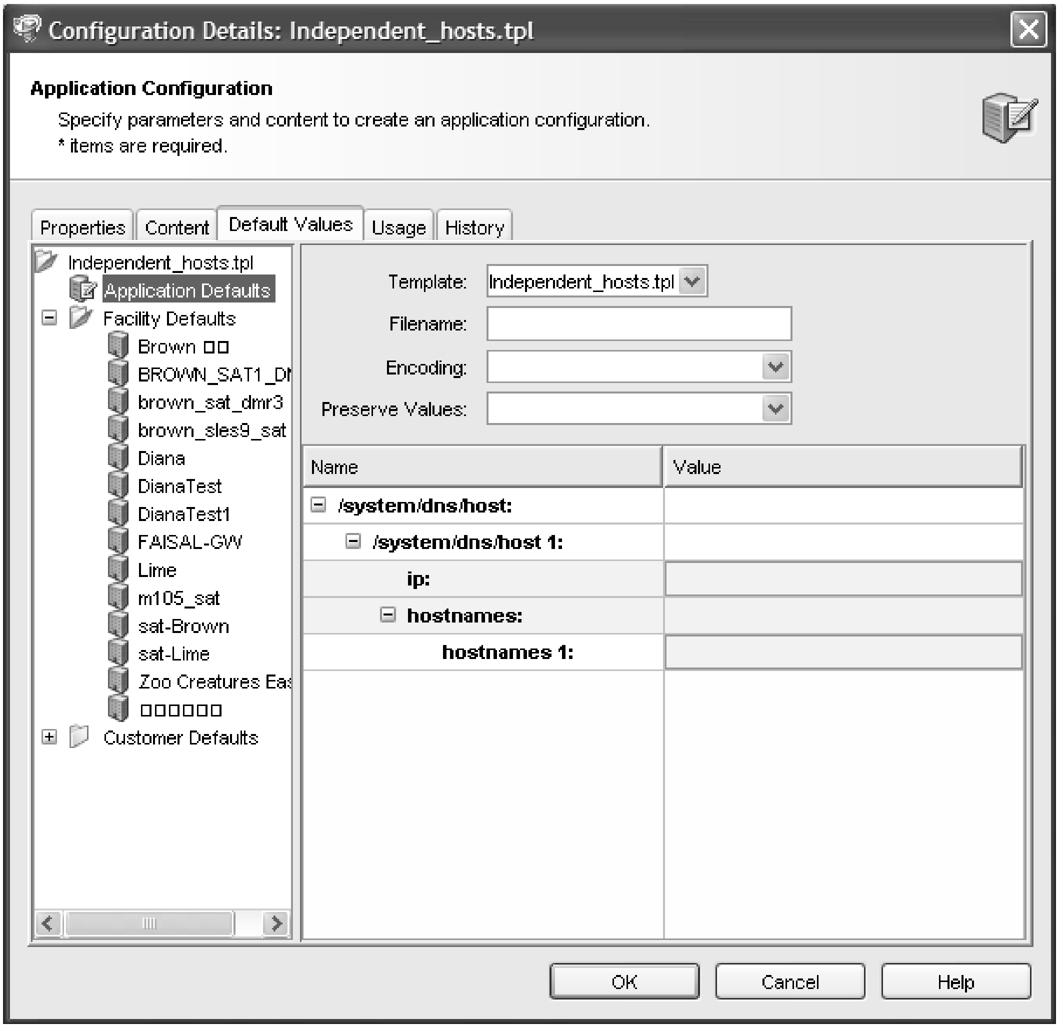

Configuration Details: Independent_hosts.tpl

Application Configuration

Specify parameters and content to create an application configuration.
* items are required.

Template: Independent_hosts.tpl
Filename:
Encoding:
Preserve Values:

Name	Value
/system/dns/host:	
/system/dns/host 1:	
ip:	
hostnames:	
hostnames 1:	

Inside the Value Set Editor, elements that are required will appear in bold.

- **Server Explorer or Server Groups Browser**: Value set elements you edit at this level replaced the actual values in the configuration files on the server when changes are applied (pushed) to a server.

Figure 16-3: Value Set Editor in the Server Explorer

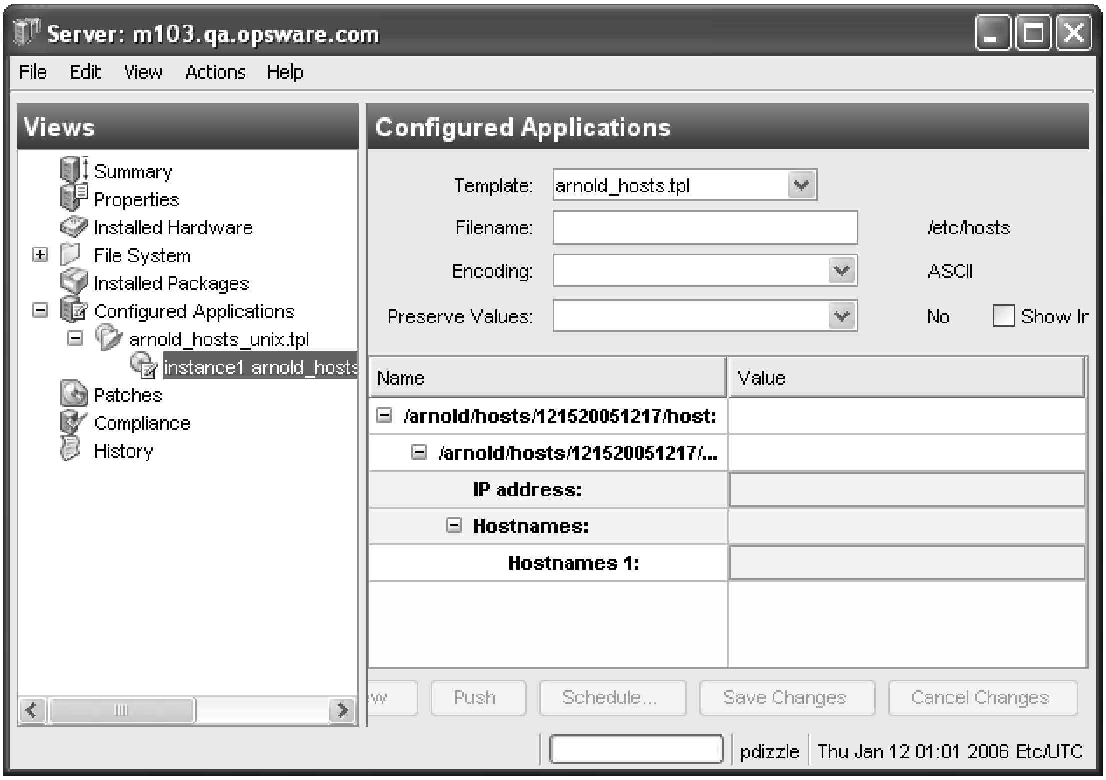

Server: m103.qa.opsware.com

Configured Applications

Template: arnold_hosts.tpl

Filename: /etc/hosts

Encoding: ASCII

Preserve Values: No

Name	Value
/arnold/hosts/121520051217/host:	
/arnold/hosts/121520051217/...	
IP address:	
Hostnames:	
Hostnames 1:	

pdizzle Thu Jan 12 01:01 2006 Etc/UTC

The left side of the Server Explorer's Configured Applications enables you to browse and select an Application Configuration to edit. If an application has more than one instance, then those instances are displayed as children of the main application. The values you edit at the parent application level apply to all instances of the application on the server. You can also edit the values of individual instances of the application.

Value Set Editor Fields

The Value Set Editor contains the following fields:

- **Template**: This enables you to choose the template you want to edit. (Some application configurations can contain multiple Application Configuration Templates.)
- **Filename**: The name of the configuration file on the managed server that is being managed by the Application Configuration Template. If no name is set, then the file

name is inherited from its parent in the inheritance hierarchy. If no file name is set anywhere in the application configuration hierarchy, then the file name listed in the Application Configuration Template is used. This field is set so the Application Configuration knows the configuration file it is supposed to manage. If you have multiple instances of an application on a server, then indicate the full pathname for each configuration file here.

- **Encoding**: Choose a character encoding for the source configuration file that the Application Configuration will be managing. The default encoding is used is the encoding used on the managed server.
- **Preserve Values**: Choose this option if you want to preserve the values contained in the actual configuration file on the server. With this option selected, the actual file's values will serve as default values for the template, and will be used unless overridden by values at some level of the inheritance hierarchy. In other words, if you would like to preserve a value of the configuration file on the server, then choose this option and leave the value blank in all scope levels. By default, this option is turned off.
- **Show Inherited Values**: Choose this option if you want to show what values are being inherited from higher levels of inheritance hierarchy. When turned off, you will only see the values set at the current level of the Application Configuration inheritance scope. When turned on, you will see both values set at the current level and those that are inherited. This view is read-only.

Value Set Editor Columns

- **Name**: This is the element name from the configuration file. A name can be a simple type, a list of simple types, or a multidimensional list. Multidimensional list names are displayed beneath their parent. Elements that are required appear in bold font. You can double-click to show or hide multidimensional lists. To add another entry to a list type value, right-click the parent and choose **Add Item**. Elements that are required will appear in bold.
- **Value**: Lists all values for each value set in the Application Configuration. You can either enter a literal value or choose an attribute from the Server's settings, such as customer name, customer ID, chassis ID, device ID, and so on. If you leave a setting blank, then the setting is inherited from its parent or ancestor (given that a parent or ancestor has settings configured). To use an Opsware SAS or custom attribute for the value, click the browse (...) button to access the Set Value dialog box.
- **Inherited From**: Indicates where the value is inherited from. The value is applied at the server instance level or inherited from its ancestors in ascending order. The order is

server instance, server, group instance, group, customers facility, and application default. However, if Preserve Values option is set in the Value Set Editor, then the configuration file on the server becomes the outermost level of the inheritance hierarchy.

Configuration Markup Language (CML)

To create a Application Configuration Template, you need to transform an application's configuration so that all its value sets become variables. See your Opsware Administrator for more information about using CML.

Application Configuration Inheritance

There are two means of controlling how an application configuration's values are applied and inherited:

- **Default Values Level**: Changes to an Application Configuration's values at this level apply to all instances of the application on all servers. You can, however, override the application configuration by customer or facility.
- **Application Level on a Managed Server**: Changes to an Application Configuration's values at this level apply to applications on a specific server, either globally to all instances of the application, or individually to specific instances of applications on the server.

Application Configuration Default Values

From the Configuration Details dialog box, you can set configuration values at the root level, and further control the scope of the configuration at the customer and facility level.

You can access this level of configuration by opening an application configuration from the OCC Client. Changes made here affect only the application configuration and do not affect the actual configuration file on the server until you push the changes onto a server using the Server Explorer. Figure 16-2 shows the application configuration hierarchy.

Application Defaults apply to all instances of all applications everywhere in the managed server environment, on all managed servers and server groups. These defaults are subdivided into the two following groups:

- **Facility**: This applies to all applications (and all instances) existing on servers that belong to a specific facility. Facility settings inherit the Application Configuration default values unless otherwise specified.
- **Customer**: This applies to all applications existing on servers that belong to a specific Customer. Customer settings inherit the facility and then the Application Configuration default values unless otherwise specified.

Application Instance Values

From the Server Explorer (or Server Groups Browser), you can manage configuration values for all or individual instances of an application on a specific managed server. Application configurations at this level inherits default values from the Application Configuration, unless you override them.

You can access this level of configuration by selecting the application or application instance from the Server Explorer ➤ Configured Applications. These Application Configurations represent actual instances of the application and its configuration on the server. Changes made here can be applied directly to the server when you click **Push**. Figure 16-4 shows the Application Configuration hierarchy at the server level.

Figure 16-4: Application Configuration Inheritance Hierarchy at the Server Level

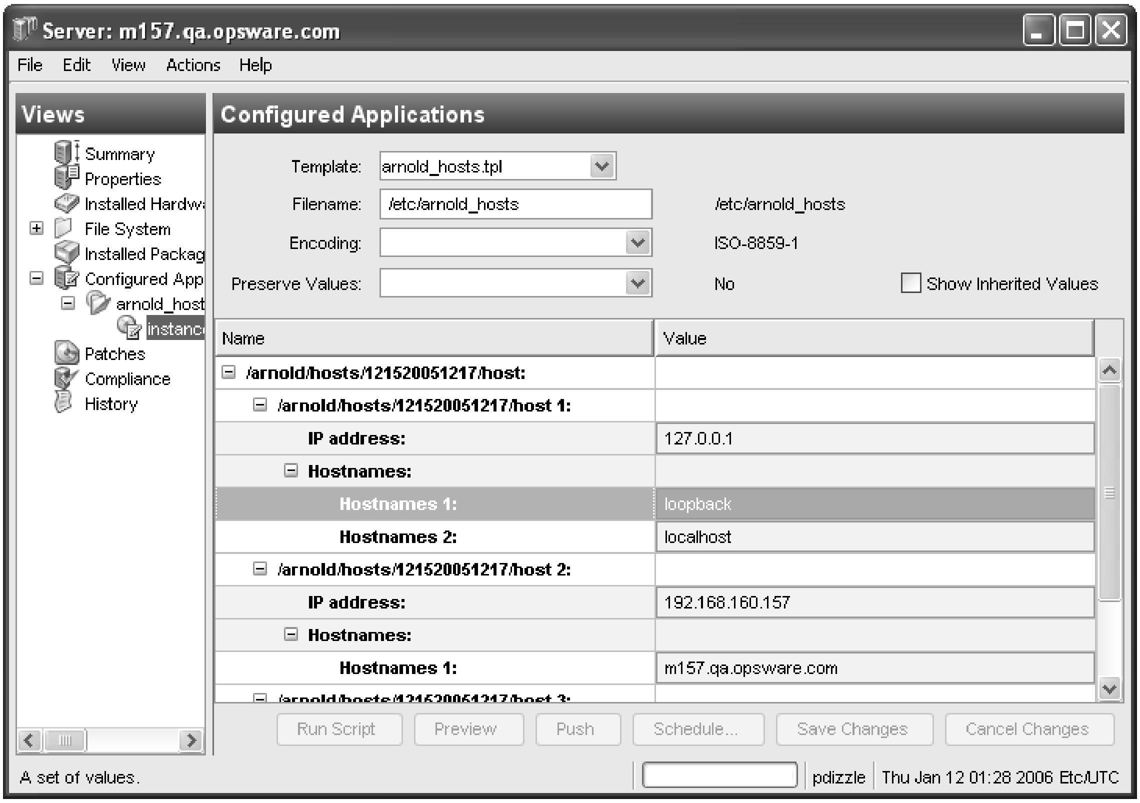

Template:	arnold_hosts.tpl	
Filename:	/etc/arnold_hosts	/etc/arnold_hosts
Encoding:		ISO-8859-1
Preserve Values:		No

□ Show Inherited Values

Name	Value
⊟ /arnold/hosts/121520051217/host:	
⊟ /arnold/hosts/121520051217/host 1:	
IP address:	127.0.0.1
⊟ Hostnames:	
Hostnames 1:	loopback
Hostnames 2:	localhost
⊟ /arnold/hosts/121520051217/host 2:	
IP address:	192.168.160.157
⊟ Hostnames:	
Hostnames 1:	m157.qa.opsware.com
⊟ /arnold/hosts/121520051217/host 3:	

A set of values.

pdizzle | Thu Jan 12 01:28 2006 Etc/UTC

Application configuration inheritance on a managed server adheres to the following hierarchy:

- **Server Group**: This applies to all applications on all servers within the specific server group. Configuration values are inherited from the Application Configuration default values unless otherwise specified. (For example, if this server group belongs to a specific customer, it inherits the values of that customer.)

  - **Server Group Application Instance**: This applies to a specific instance of an application on all servers in the specific server group. This instance inherits configuration values from the application defaults and any other application configuration default values, unless otherwise specified.

- **Server**: This applies to all applications on the server. The instance inherits configuration values from application defaults on the managed server from the server group it belongs to (if it belongs to a group), and any Application Configuration default values.

  - **Server Instance**: This applies only to the specific instance of the application on the specific server. This instance inherits configuration values from application defaults, from defaults server settings, from the server group the server belongs to (if it belongs to a group), and any other Application Configuration default values.

Application Configuration Inheritance Visualized

Figure 16-5 illustrates how Application Configuration values are inherited.

Figure 16-5: Application Configuration Inheritance

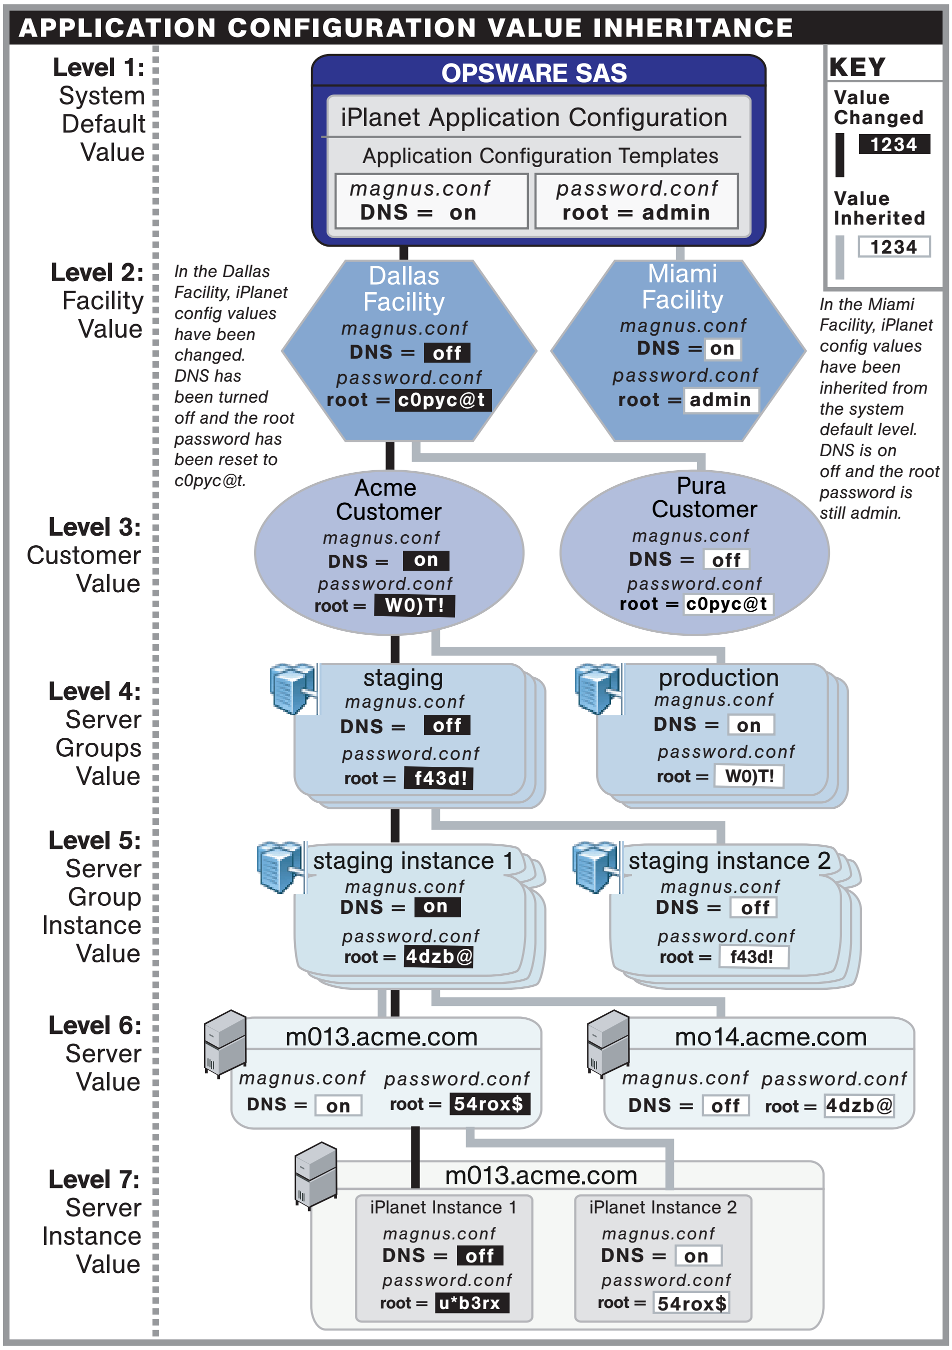

In the Dallas Facility, iPlanet config values have been changed. DNS has been turned off and the root password has been reset to c0pyc@t.

In the Miami Facility, iPlanet config values have been inherited from the system default level. DNS is on off and the root password is still admin.

Sequence Merging and Inheritance

Because Application Configuration values can be set across many different levels in the Application Configuration inheritance hierarchy (also referred to as the inheritance scope), it is important that you be able control the way multiple sequence values are merged together when you push an Application Configuration on to a server.

ACM allows you to control the way sequence values are merged across inheritance scopes. This means that you can, for example, add some values to a sequence in the Customer scope, Group scope, and the Server scope, and all the values will be merged together to form the final sequence.

The manner in which sequence values are merged is controlled by special tags in the CML template, using three different sequence merge modes:

- **Sequence Replace**: Sequence values from more specific scopes completely replace those from less specific scopes. This occurs for both sequences of sets and lists.
- **Sequence Append**: For lists, values at more general scopes are appended (placed after) to those at more specific scopes. Duplicates, if present, are not removed. For sets, the behavior is the same, except duplicates are merged. For lists, duplicates are identified according to child elements marked with the `primary-key` tag, and then merged. For scalars, this is done by simply removing duplicate values, leaving only the value from the most specific scope (the last occurrence is the merged sequence). This is the default mode, and will be used if nothing else is specified.
- **Sequence Prepend**: Works the same as append, but values at more general scopes are preprended (placed before) to those at more specific scopes.

For example, with these two sets:

- "a, b" – At a more specific (inner) level of the inheritance scope, for example, server instance level.
- "c, d" – At a more general (outer) of the inheritance scope, for example, the server group level.

When the application configuration template is pushed onto the server, the merging results would be:

- Sequence replace: "a, b"
- Sequence append: "a, b, c, d"

- Sequence prepend: “c, d, a, b”

Sequence aggregation occurs not only between scopes, but also within a scope itself. This is evident if there are duplicate values within a sequence of namespaces.

Sequence Replace

In the Replace merge mode (CML tag “`sequence-replace`”), the contents of a sequence defined at a particular scope replace those of less specific scopes, and no merging is performed on the individual elements of the sequence.

For example, if the `sequence-replace` tag has been set for a list in an Application Configuration Template CML source, then values set for that list at the server instance level will override, or replace, those set at the group level and at the Application Configuration default values level.

For example, if a list in an `etc/hosts` file was defined at the group level (outer) as the following:

```
/system/dns/host/1/ip           127.0.0.1
/system/dns/host/1/hostnames/1  localhost
/system/dns/host/1/hostnames/2  mymachine
/system/dns/host/2/ip           10.10.10.10
/system/dns/host/2/hostnames/1  loghost
```

And the same list was defined at the device scope (inner), as the following:

```
/system/dns/host/1/ip           127.0.0.1
/system/dns/host/1/hostnames/1  localhost
/system/dns/host/1/hostnames/2  mymachine.mydomain.net
/system/dns/host/2/ip           10.10.10.100
/system/dns/host/2/hostnames/1  mailserver
```

If template had defined the `/system/dns/host` element with the `sequence-replace` tag, the final results of the configuration file on the server after the push would be:

```
127.0.0.1 localhost mymachine.mydomain.net
10.10.10.100 mailserver
```

Sequence Append

When the append list merge mode (CML tag “`sequence-append`”) is used for sequences, the values at more general scopes are appended (placed after) those of more specific scopes. Sequence append mode is the default mode for merging list values. If nothing is specified in the CML of the template, the sequence append will be used.

If a list in an `etc/hosts` file was defined at the group level (outer) as the following:

```
/system/dns/host/1/ip           127.0.0.1
/system/dns/host/1/hostnames/1  localhost
/system/dns/host/1/hostnames/2  mymachine
/system/dns/host/2/ip           10.10.10.10
/system/dns/host/2/hostnames/1  loghost
```

And the same list was defined at the device scope (inner), as the following:

```
/system/dns/host/1/ip           127.0.0.1
/system/dns/host/1/hostnames/1  localhost
/system/dns/host/1/hostnames/2  mymachine.mydomain.net
/system/dns/host/2/ip           10.10.10.100
/system/dns/host/2/hostnames/1  mailserver
```

Using the value sets from the above example, if the `/system/dns/host` element was a list with the `sequence-append` tag set in the Application Configuration Template, the final results of the configuration file on the server after the push would be:

```
127.0.0.1 localhost mymachine.mydomain.net
10.10.10.100 mailserver
127.0.0.1 localhost mymachine
10.10.10.10 loghost
```

But since it is not allowable for a hosts file to contain duplicate entries, the`/system/dns/host` element will have to be flagged in the Application Configuration Template as a set rather than a list, because sets do not allow duplicates. To avoid duplication of the list values in the example, the Application Configuration Template author would use the Primary Key option.

Primary Key Option in Sequence Merging

When operating in append mode on sets, new values in more specific scopes are appended to those of less specific ones, and duplicate values are merged with the resulting value placed in the resulting sequence according to its position in the more specific scope.

How this affects merged sequence values depends on what kind of data is contained in the sequence:

- For elements in a sequence which are scalars, the value from the most specific scope is used. In other words, values at the server instance level would replace the values at the group level.

- For elements which are namespace sequences, the value is obtained by applying the merge mode specified for that element (in this example, append) based upon matching up the primary fields.

To avoid the duplication of the /system/dns/host/.ip value, the Application Configuration Template author would use the CML primary-key option. With this option set, ACM will treat entries with the same value for /system/dns/host/.ip as the same and merge their contents.

In the example above, the final results of the configuration file on the server after the push would be:

```
127.0.0.1 localhost mymachine.mydomain.net mymachine
10.10.10.100 mailserver
10.10.10.10 loghost
```

Since it is possible to have a set without primary keys, if there are scalars in the sequence, then an aggregation of all scalar values will be used as the primary key. If there are no scalars, then the aggregation of all values in the first sequence will be used as the primary key. Although this is an estimate, in most cases the values will be merged effectively. To ensure that the correct values are used as primary keys, we recommend that you always explicitly set the primary key in a sequence.

Sequence Prepend

When the append list merge mode (CML tag “sequence-prepend”) is used for sequences, the values at more general scopes are prepended (placed before) those those of more specific scopes.

For example, if a sequence in an etc/hosts file was defined at the group level (outer) as the following:

```
/system/dns/host/1/ip            127.0.0.1
/system/dns/host/1/hostnames/1   localhost
/system/dns/host/1/hostnames/2   mymachine
/system/dns/host/2/ip            10.10.10.10
/system/dns/host/2/hostnames/1   loghost
```

And the same sequence was defined at the device scope (inner), as the following:

```
/system/dns/host/1/ip            127.0.0.1
/system/dns/host/1/hostnames/1   localhost
/system/dns/host/1/hostnames/2   mymachine.mydomain.net
```

```
/system/dns/host/2/ip           10.10.10.100
/system/dns/host/2/hostnames/1  mailserver
```

If the /system/dns/host element was a set with the sequence-prepend tag set in the Application Configuration Template, the final results of the configuration file on the server after the push would be:

```
10.10.10.10 loghost
127.0.0.1 mymachine localhost mymachine.mydomain.net
10.10.10.100 mailserver
```

To find out how sequences are handled in an Application Configuration Template before you push, you need to look at the contents of the CML template source. For information on how to examine the CML contents of an Application Configuration Template, see "Viewing Application Configuration Template Sources" on page 620.

If you would like to see preview the results of sequence merging before you push, see "Comparing a Template Against an Actual Configuration File" on page 638.

Using ACM

This section contains the following tasks:

- Creating an Application Configuration
- Creating a Configuration Template
- Viewing Application Configuration Template Sources
- Adding or Removing Configuration Templates
- Deleting Application Configurations
- Loading a Template File
- Setting a Configuration Template to Run as a Script
- Specifying Template Order
- Editing Default Values for an Application Configuration
- Attaching an Application Configuration to a Server or Group
- Setting Application Configuration Values on a Server or Group
- Loading Existing Values into a Configuration Template
- Pushing Changes to a Server or Group

- Scheduling an Application Configuration Push
- Comparing Two Configuration Templates
- Comparing a Template Against an Actual Configuration File
- Auditing an Application Configuration
- Scheduling an Application Configuration Audit
- Rolling Back to a Previous State

Creating an Application Configuration

An application configuration can contain one or more Application Configuration Templates (and scripts). Because an application is likely to have more than one configuration file and thus necessitate multiple Application Configuration Templates, you need to create an application configuration to organize and manage your templates from a single location.

If you only want to manage a single configuration file with a single Application Configuration Template, you still need to create an Application Configuration to deploy the template on a server.

To create an application configuration, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library.
2. Select Application Configuration, and then select the Application Configurations tab.
3. From the **Action** menu, select **New**.
4. In the Properties tab of the Configuration Detail dialog box, specify the following properties:
   - **Name**: This field enables you to name the Application Configuration. (This is required.)
   - **Description**: This field enables you to describe the Application Configuration.
   - **Version**: This section enables you to give a version number to the Application Configuration. This value is set by the person who creates and modifies the Application Configuration. (This version number is not incremented automatically.)
   - **OS**: This allows you to limit the use of the Application Configuration to specific operating systems. The Available list indicates the operating systems you can associate with the Application Configuration. The Selected list shows the operating systems currently associated with the Application Configuration. Click the arrow to

add or remove an operating system to the Application Configuration. Once you add an operating system, then only servers using those operating systems will be able to use the Application Configuration. If you do not want this Application Configuration to be associated with an operating system, select OS Independent.

- **Customers**: This option enables you to limit the use of the application configuration to a specific customer. The Available list of platforms indicates the customers currently supported for the Application Configuration. The Selected list shows the customers associated with the Application Configuration. Click the arrow to add or remove customers from the Application Configuration. If you do not want this Application Configuration to be associated with a customer, select Customer Independent.
- **Notes**: This section allows you to add notes to the Application Configuration.
- **Created**: The date that the Application Configuration was created.
- **Created By**: The user who created the Application Configuration.
- **Last Modified**: The date that the Application Configuration was last modified.
- **Modified By**: The user who last modified the Application Configuration.
- **Tested**: This option allows you to indicate that the Application Configuration has successfully been pushed to a server and that it works.

5. Select the Content tab.
6. To add an application configuration template, click **Add**.
7. In the Select Configuration File dialog box, select an Application Configuration template, and then click **OK**.
8. If the Application Configuration is run as a script, select the Application Configuration, right-click, and select one of the following menu items: **None** (will not run as script), or **Data-manipulation**, **Pre-install**, **Post-install**, **Post-error**.
9. Click **OK** to create the new Application Configuration.

Creating a Configuration Template

An Application Configuration Template is similar to an actual native application configuration file, but one that has had its variable portions marked up with Opsware's Configuration Markup Language (CML). (CML is a markup language used for managing configuration files.)

To manage a configuration file with ACM, create an Application Configuration Template. Before a Application Configuration Template can be applied to a server, it needs to be added to an Application Configuration.

An Application Configuration Template can be configured to run as a script, either before all the configurations are made or after. Also, you can set a script to run as a post-error script to rollback all changes if the configuration push fails. See "Setting a Configuration Template to Run as a Script" on page 624 in this chapter for more information.

To create a Application Configuration Template, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library.
2. Select Application Configuration, and then select the Configuration Templates tab.
3. From the **Action** menu, select **New**.
4. In the Properties tab of the Template Detail dialog box, enter the following information:
   - **Name**: This allows you to enter a name for the Application Configuration or Application Configuration Template. (This is required.)
   - **Description**: This enables you to enter a description.
   - **Version**: This value is set by the person who creates and modifies the Application Configuration/Application Configuration Template. (The version number is not incremented automatically.)
   - **OS**: This allows you to limit the use of the Application Configuration Template to a specific operating system. The Available list of operating systems indicates the operating systems you can associate with the Application Configuration or Application Configuration Template. The Selected list shows the operating systems currently associated with the Application Configuration/Application Configuration Template. Click the arrow to add or remove an operating system to the Application Configuration Template. Once you add an operating system, then only servers using those operating systems can use the Application Configuration Template. If you do not want this Application Configuration/Application Configuration Template to be associated with an operating system, select the OS Independent option.
   - **Customers**: This option allows you to limit the use of the Application Configuration/Application Configuration Template to a specific customer. The Available list of platforms indicates the customers that are currently supported for the Application Configuration or Application Configuration Template. The Selected list shows the customers associated with the Application Configuration/Application

Configuration Template. Click the arrow to add or remove customers from the Application Configuration or Application Configuration Template. If you do not want an Application Configuration or Application Configuration Template to be associated with customer, select the OS Independent option.

- **Script Type**: This allows you to set the Application Configuration Template to function as a template, localization file, or script. If the file is a script, you can specify the script language, such as WIndows BAT, JS, VBS, CMD, WSF, and PY; and Unix SH or Other script.
- **Created**: This shows the date that the Application Configuration Template was created.
- **Created By**: This shows the user who created the Application Configuration Template.
- **Last Modified**: This shows the date that the Application Configuration Template was last modified.
- **Modified By**: This shows the user who last modified the Application Configuration Template.
- **Tested**: This option allows you to indicate that the Application Configuration Template has successfully been pushed to a server and that it works.

5. Select the Content tab.
6. Copy the contents of your CML file here.
7. Click **Validate** to validate the CML syntax.
8. When you are finished, click **OK**.

Viewing Application Configuration Template Sources

In some cases, you will need to examine the contents of your Application Configuration Template and view its CML source, especially if you need to understand which list merging modes have been set in the template before you push the Application Configuration to a server.

For information on Application Configuration sequence merge modes, see “Sequence Merging and Inheritance” on page 612.

To view Application Configuration Template CML source, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library.
2. Select Application Configuration, and then select the Configuration Templates tab.

3. To open an Application Configuration Template in the list, double-click it. (Or right-click the template and choose **Open**.)
4. Select the Content tab, and you see the CML contents of the Application Configuration Template.

Adding or Removing Configuration Templates

You can add as many Application Configuration Templates to an Application Configuration as you like. If an Application Configuration Template doesn't belong or you no longer need it in an Application Configuration, you can remove it.

To add an Application Configuration Template to an Application Configuration, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library.
2. Select Application Configuration, and then select the Application Configurations tab.
3. To open an Application Configuration in the list, double-click it.
4. Select the Content tab.
5. To add an Application Configuration Template, click **Add**.
6. From the Select Configuration dialog box, select the Application Configuration Template, and then click **OK**.

Deleting Application Configurations

If you no longer need an Application Configuration, you can delete it. Once you delete an Application Configuration, you cannot recover it.

To delete an Application Configuration, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library.
2. Select Application Configuration, and then select the Application Configurations tab.
3. Select an Application Configuration, right-click, and choose **Delete**. (This will not delete any Application Configuration Templates that belong to the Application Configuration.)
4. To delete a Application Configuration Template, select the Configuration Templates tab.
5. Select an Application Configuration Template, right-click, and choose **Delete**.

Loading a Template File

If a CML template is already created for use in an Application Configuration, you can upload the template from a local or remote file system.

For configuration files on Windows servers which are encoded in UTF-8, the first three characters of the configuration file might contain a Byte Order Mark (BOM). If you import this file into an Application Configuration Template, the BOM will appear in the template after the file is loaded. If you do not want this BOM to be included in the Application Configuration Template, remove it after you upload the configuration file into the template.

To load a template file, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library.
2. From the **Action** menu, select **Upload Template**.
3. In the Open dialog box, browse to locate the template file (a CML file should have the TPL file extension, but this is not mandatory). If the character encoding of the template file is different than the default encoding of your desktop, select an item from the Encoding drop-down list.
4. Click **Open**.
5. In the Configuration File Upload dialog box, fill out the following information:
   - **Name**: This allows you to enter a name for the Application Configuration or Application Configuration Template. (This is required.)
   - **Description**: This enables you to enter a description.
   - **Version**: This value is set by the person who creates and modifies the Application Configuration/Application Configuration Template. (The version number is not incremented automatically.)
   - **OS**: This allows you to limit the use of the Application Configuration Template to a specific operating system. The Available list of operating systems indicates the operating systems you can associate with the Application Configuration or Application Configuration Template. The Selected list shows the operating systems currently associated with the Application Configuration/Application Configuration Template. Click the arrow to add or remove an operating system to the Application Configuration Template. Once you add an operating system, then only servers using those operating systems can use the Application Configuration Template. If

you do not want this Application Configuration/Application Configuration Template to be associated with an operating system, select the OS Independent option.

- **Customers**: This option allows you to limit the use of the Application Configuration/Application Configuration Template to a specific customer. The Available list of platforms indicates the customers that are currently supported for the Application Configuration or Application Configuration Template. The Selected list shows the customers associated with the Application Configuration/Application Configuration Template. Click the arrow to add or remove customers from the Application Configuration or Application Configuration Template. If you do not want an Application Configuration or Application Configuration Template to be associated with customer, select the OS Independent option.
- **Script Type**: This allows you to set the Application Configuration Template to function as a template, localization file, or script. If the file is a script, you can specify the script language, such as WIndows BAT, JS, VBS, CMD, WSF, and PY; and Unix SH or Other script.
- **Created**: This shows the date that the Application Configuration Template was created.
- **Created By**: This shows the user who created the Application Configuration Template.
- **Last Modified**: This shows the date that the Application Configuration Template was last modified.
- **Modified By**: This shows the user who last modified the Application Configuration Template.
- **Tested**: This option allows you to indicate that the Application Configuration Template has successfully been pushed to a server and that it works.

6. Next, select the Content tab.
7. You should see the CML template. Click **Validate** to validate the CML syntax.
8. When you are finished, click **OK**. This will create both the Application Configuration Template and an Application Configuration to house the template.

Setting a Configuration Template to Run as a Script

In addition to using Application Configuration Templates to replace values of actual configuration files, you can also add scripts to an Application Configuration.

For example, you might want to add a post-install script that reboots the server after configuration changes have been made. Or, you might want to use a data-manipulation script to handle certain configuration files which contain unreadable or otherwise unmanageable data before you perform an import, preview, or push the Application Configuration.

If you are configuring an IIS server, you can use a data-manipulation script to read the metabase information into a flat file. When this information gets parsed with the Application Configuration Template, you can run a data-manipulation script to implement the changes in the flat file.

To set an Application Configuration Template as a script, you need to set the Application Configuration Template script type and then specify the type of script execution.

To set a template to run as a script, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Software Library.
2. Select Application Configuration, and then select the Application Configurations tab.
3. In the Content pane, double-click the Application Configuration that contains the Application Configuration Template that you want to run as a script.
4. In the Configuration Details window, select the Content tab.
5. Select the Application Configuration Template in the list, right-click, and choose **Data-manipulation**, **Pre-install**, **Post-install**, or **Post-error** to set the script execution type.

If you would like to change the order in which the Application Configuration Template is run inside the Application Configuration, select the Application Configuration Template, right-click, and select **Move Up** or **Move Down**.

6. Select the Application Configuration Template again, right-click, and select **Open Template**.
7. In the Template Details window, choose a script type from Type drop-down list. Click **OK**.

8 Click **OK** to close the Configuration Details window.

When pushing an application configuration that contains a JScript or VBScript pre- or post-install and post-error scripts, the push succeeds even though the scripts fail. In these cases, the push ignores the scripts errors altogether. The application configuration does not catch the failure of the scripts and allows the push to complete without errors.

If you plan to use these types of scripts, you must make sure that the scripts are free of errors to detect possible failures, and have the script forcibly return a non-zero exit status by invoking WScript.Quit(<status>).

Specifying Template Order

An Application Configuration can contain one or several Application Configuration Templates and scripts. However, you may want to control templates application and script execution order.

For example, you may want to apply changes to certain configuration files before others. Or, you may have a script in the Application Configuration that restarts the server after all the Application Configuration changes have been applied to the application on the server.

To specify template order, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Software Library.

2 Select Application Configuration, and then select the Application Configurations tab.

3 To open an Application Configuration in the list, double-click it.

4 In the Configuration Detail dialog box, select the Content tab.

5 All the Application Configuration Templates and scripts (if there are any) contained within the Application Configuration are displayed. Notice that each Application Configuration Template has a number next to it that indicates the order.

6 To reorder the Application Configuration Templates, select one and then click **Move Item Up** or **Move Item Down**.

For better organization, it is useful to position at any pre-install scripts at the top of the list, and position post-install or post-error scripts at the bottom of the list.

7 When you are finished, click **OK**.

Editing Default Values for an Application Configuration

Once you have created an Application Configuration, you can edit its default configuration values. An Application Configuration's default values apply to all instances of the application on all attached servers. (An Application Configuration only affects attached servers.)

However, you can override the scope of an application configuration's default values by customer or facility. You can also edit specific instances of the application configuration to override the scope of an application configuration's default values. All elements that are required appear in bold font.

To set default values for an application configuration, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Software Library.

2 Select Application Configuration, and then select the Application Configurations tab.

3 In the Content pane, double-click the Application Configuration.

4 In the Configuration Details dialog box, select the Default Values tab.

5 The left side of the dialog box shows the Application Configuration hierarchy; this allows you to set default values at the application defaults (root) level, the customer level, and the facility level.

6 To set default values, select a server in the hierarchy and double-click it. The default values will display.

Figure 16-6 shows an example of the Application Defaults node selected. Any changes to value sets at this level will apply to all facilities and customers – including all applications on all attached servers.

Figure 16-6: Application Configuration Default Values Hierarchy

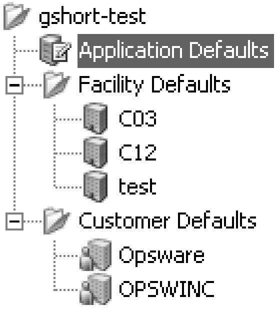

7. Edit the default values for each value set in the Application Configuration Template. The following settings will be displayed:

   - **Template**: This enables you to choose the template you want to edit. (Some application configurations can contain multiple Application Configuration Templates.)
   - **Filename**: The name of the configuration file on the managed server that is being managed by the Application Configuration Template. If no name is set, then the file name is inherited from its parent in the inheritance hierarchy. If no file name is set anywhere in the application configuration hierarchy, then the file name listed in the Application Configuration Template is used. This field is set so the Application Configuration knows the configuration file it is supposed to manage. If you have multiple instances of an application on a server, then indicate the full pathname for each configuration file here.
   - **Encoding**: This enables you to choose a character encoding for the source configuration file that the Application Configuration will be managing. The default encoding is the encoding used on the managed server.
   - **Preserve Values**: To preserve the values contained in the actual configuration file on the server, choose **Yes** for this option and leave the value blank in all scope levels. With this option selected, the actual file's values will serve as default values for the template, and will be used unless overridden by values at some level of the inheritance hierarchy. By default, this option is turned off.
   - **Show Inherited Values**: This appears only on an Application Configuration

instance attached to a server or server group, not at the Application Configuration default values level. Choose this option if you want to show at what values are being inherited from higher levels of inheritance hierarchy. When turned off, you will only see the values set at the current level of the Application Configuration inheritance scope. When turned on, you will see both values set at the current level and those that are inherited. This view is read-only.

– **Name column**: This is the value set element name from the configuration file. A value set name can consist of a simple type, a list of simple types, or a multidimensional list. Elements that are required appear in bold font. Multidimensional list key names are displayed beneath their parent. Double-click to show or hide multidimensional lists. To add another key name, right-click the parent and select **Add Item**. You can also use the right-click menu to search for other values or keys, copy values, or clear values.

– **Value column**: This allows you to enter a literal value or choose an attribute from the Server's settings, such as customer name, customer ID, chassis ID, device ID, and so on. If you leave a setting blank, then the setting is inherited from its parent or ancestor (if a parent or ancestor has settings configured). To use an Opsware or custom attribute for the value, click the browse (...) button to access the Set Value dialog box.

8 To edit or change a value, either type a string value directly into the field, or click once in the value field, then click the browse (...) button to access the Set Value window. Choose one of the following options:

– **No Value**: Choose this to set no value to the value set key.

– **Block Inheritance**: Choose this option if you do not want to inherit any values from values set at higher levels of the Application Configuration inheritance hierarchy. The effect this has when you push the template to the server depends on if the value is a scalar or a list value:

- **Scalars**: If the value is a scalar, this key's value will be removed from the configuration file when pushed to the server. Thus, this option is a means of removing a scalar value from a configuration file.

- **Lists**: If the value is list, then any values from higher levels of the inheritance hierarchy will be blocked, but the current level and any lower levels of the scope will be pushed to the configuration file on the server.

- **Note**: To block a namespace sequence from inheriting from other scopes, you should add a new namespace sequence that has the a single scalar value or the only entry in a sequence set to <Block Inheritance> with all other fields empty.

- **Any Value**: Enter a value here.
- **Opsware Attribute**: From the drop-down list, choose an Opsware attribute to use, such as customer name, customer ID, chassis ID, device ID, and so on.
- **Custom Attribute**: Enter your own custom attribute here.

9 (Optional) You can copy and paste one value set to another. To do this, select the value set name, right-click, and choose **Copy Values**. Then, paste this value by right-clicking the target value set and choosing **Paste Values**. Copying and pasting will copy the entire value set and will override the old value set.

10 (Optional) You can expand and retract the Application Configuration value set, by right-clicking and choosing **Collapse Subtree**. All name-value hierarchies will be closed. If you would like to find a value set name or value, select the value set, right-click and choose **Find Name** or **Find Value**.

11 When you have finished editing the value sets for the Application Configuration, click **Save Changes**.

Attaching an Application Configuration to a Server or Group

After you have created an Application Configuration and added all the necessary Application Configuration Templates and scripts and edited its default values, you can add the Application Configuration to a server or public server group.

For an Application Configuration to manage an application on a server, it must be added to a server or server group. Once you add an Application Configuration to a server or server group, the values of the Application Configuration are not applied to the configuration files on the server until you push them to the server. This enables you to add the Application Configuration, edit its values, and then wait until you are ready to apply the changes before pushing them to the server.

You can only add an Application Configuration to a public server group.

To attach an Application Configuration to a server or group, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Servers.

2 From the Navigation pane, select a server or server group.

3. Select a server (or group) from the Content pane, right-click and choose **Configure Applications**.

4. You now see the Server Explorer (or Server Groups Browser), with the Configured Applications folder selected. From the **Action** menu, select **Add Configuration**.

5. In the Select Application Configuration dialog box, select an Application Configuration.

   Use the search tool 🔍 in the upper right corner of the dialog box if the list is large and you want to search by a specific criteria (such as OS, last modified, and so on).

6. When you have selected an Application Configuration, click **OK**. The Application Configuration is added to the server or group.

Setting Application Configuration Values on a Server or Group

Once an Application Configuration has been attached to a server or server group, you can edit its values. You can also override the default values set at the Application Configuration level. If the server (or group) has multiple instances of an application installed, you can set values for all instances of the application or individual instances.

If you do not edit any values on the Application Configuration at the server or group level, then the values are inherited from the default values set at the Application Configuration level. See “Application Configuration Inheritance” on page 607 in this chapter for more information.

To set Application Configuration values on a server or group, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers.

2. Select a server or server group in the Navigation pane.

3. Select a server or group in the Content pane, right-click, and choose **Configure Applications**.

4. You now see the Server Explorer (or Server Groups Browser), with the Configured Applications folder selected. Select the Installed Configurations tab. All Application Configurations that have been attached to the server (or group) will be displayed.

If you do not see an Application Configuration, then none have been attached to the server or group. See “Attaching an Application Configuration to a Server or Group” on page 629 in this chapter for more information.

5 From the left side of the Application Configuration hierarchy, select either the top level application folder or an instance of the application, then edit the values of the Application Configuration. Before you start editing values, consider the following about Application Configuration inheritance:

- If you do not edit any values on the application or application instance level, then all values are inherited from the Application Configuration’s default values. (See “Editing Default Values for an Application Configuration” on page 626 in this chapter for more information.)
- If you want to see which values are being inherited from a higher level of the Application Configuration hierarchy, select the Show Inherited Values option. Selecting this option will show a read only view of all names and values in the Application Configuration, and the inherited from column shows where inherited values are derived from.

Once you have selected a level of the Application Configuration to edit, you can now start editing values. Because every configuration file is unique, what you actually see and are able to edit will be different for each Application Configuration.

6 Edit the default values for each value set in the Application Configuration Template. The following settings will be displayed:

- **Template**: This enables you to choose the template you want to edit. (Some application configurations can contain multiple Application Configuration Templates.)
- **Filename**: The name of the configuration file on the managed server that is being managed by the Application Configuration Template. If no name is set, then the file name is inherited from its parent in the inheritance hierarchy. If no file name is set anywhere in the application configuration hierarchy, then the file name listed in the Application Configuration Template is used. This field is set so the Application Configuration knows the configuration file it is supposed to manage. If you have multiple instances of an application on a server, then indicate the full pathname for each configuration file here.
- **Encoding**: This enables you to choose a character encoding for the source configuration file that the Application Configuration will be managing. The default

encoding is the encoding used on the managed server.

- **Preserve Values**: To preserve the values contained in the actual configuration file on the server, choose **Yes** for this option and leave the value blank in all scope levels. With this option selected, the actual file's values will serve as default values for the template, and will be used unless overridden by values at some level of the inheritance hierarchy. By default, this option is turned off.
- **Show Inherited Values**: This appears only on an Application Configuration instance attached to a server or server group, not at the Application Configuration default values level. Choose this option if you want to show at what values are being inherited from higher levels of inheritance hierarchy. When turned off, you will only see the values set at the current level of the Application Configuration inheritance scope. When turned on, you will see both values set at the current level and those that are inherited. This view is read-only.
- **Name column**: This is the value set element name from the configuration file. A value set name can consist of a simple type, a list of simple types, or a multidimensional list. Elements that are required appear in bold font. Multidimensional list key names are displayed beneath their parent. Double-click to show or hide multidimensional lists. To add another key name, right-click the parent and select **Add Item**. You can also use the right-click menu to search for other values or keys, copy values, or clear values.
- **Value column**: This allows you to enter a literal value or choose an attribute from the Server's settings, such as customer name, customer ID, chassis ID, device ID, and so on. If you leave a setting blank, then the setting is inherited from its parent or ancestor (if a parent or ancestor has settings configured). To use an Opsware or custom attribute for the value, click the browse (...) button to access the Set Value dialog box.

7. To edit or change a value, either type a string value directly into the field, or click once in the value field, then click the browse (...) button to access the Set Value window. Choose one of the following options:

   - **No Value**: Choose this to set no value to the value set key.
   - **Block Inheritance**: Choose this option if you do not want to inherit any values from values set at higher levels of the Application Configuration inheritance hierarchy. The effect this has when you push the template to the server depends on if the value is a scalar or a list value:

   - **Scalars**: If the value is a scalar, this key's value will be removed from the configuration file when pushed to the server. Thus, this option is a means of removing a scalar value from a configuration file.
   - **Lists**: If the value is list, then any values from higher levels of the inheritance hierarchy will be blocked, but the current level and any lower levels of the scope will be pushed to the configuration file on the server.
   - **Note**: To block a namespace sequence from inheriting from other scopes, you should add a new namespace sequence that has the a single scalar value or the only entry in a sequence set to <Block Inheritance> with all other fields empty.
- – **Any Value**: Enter a value here.
- – **Opsware Attribute**: From the drop-down list, choose an Opsware attribute to use, such as customer name, customer ID, chassis ID, device ID, and so on.
- – **Custom Attribute**: Enter your own custom attribute here.

8. (Optional) You can copy and paste one value set to another. To do this, select the value set name, right-click, and choose **Copy Values**. Then, paste this value by right-clicking the target value set and choosing **Paste Values**. Copying and pasting will copy the entire value set and will override the old value set.

9. (Optional) You can expand and retract the Application Configuration value set, by right-clicking and choosing **Collapse Subtree**. All name-value hierarchies will be closed. If you would like to find a value set name or value, select the value set, right-click and choose **Find Name** or **Find Value**.

10. When you have finished editing the Application Configuration values, click **Save Changes**. These changes won't be applied to the configuration files on the server or group until you push the changes. To preview what the changes will look like before you push them, click **Preview**. To push the changes, click **Push**.

Loading Existing Values into a Configuration Template

You might want to import values into the value set editor from a configuration file on a managed server. Selecting the **Import Values** menu item reads the actual existing configuration file on a server, parses the values, and applies them into the instance level value sets for the Application Configuration Template. This shows the values currently in the actual configuration. After you import the values, you can modify some of those values and then push the changes back onto the server.

To load existing values into the value set editor, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers.
2. Select a server or server group in the Navigation pane.
3. Select a server or group in the Content pane, right-click, and choose **Configure Applications**.
4. You now see the Server Explorer (or Server Groups Browser), with the Installed Configurations tab selected. All Application Configurations that have been attached to the server (or group) will be displayed.
5. From the View pane, click the plus (+) symbol to expand Application Configuration folder and select an Application Configuration instance to edit.
6. From the Content pane, choose an Application Configuration Template from the Template drop-down list.
7. In the File name field, enter the absolute file name of the configuration file that contains the values that you want to import.
8. Next, right-click in the Name column and choose **Import Values**. A confirmation message appears, warning you that proceeding with this operation will overwrite any current values. Click **Yes** to proceed.
9. All of the values for the Application Configuration Template are replaced with the values from the actual configuration file.
10. Click **Save Changes**.

Pushing Changes to a Server or Group

After you have edited Application Configuration values in the Value Set Editor, you must apply them to the application on the server. To do so, you need to perform a push operation. Performing a push operation applies modifications to the actual configuration files on the server (or group).

The way in which sequences (of lists and scalars) are merged when you push depends upon how values have been set in the Application Configuration inheritance hierarchy and what sequence merge modes have been configured in the CML template for the Application Configuration. For more information about sequence merging, see "Sequence Merging and Inheritance" on page 612.

To push Application Configuration changes to a server or group, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers.
2. From the left Navigation pane, select a server or server group.
3. Select a server (or group) from the Content pane, right-click and choose **Configure Applications**.
4. You now see the Server Explorer (or Server Groups Browser), with the Configured Applications folder selected. Select the Installed Configurations tab. All Application Configurations that have been attached to the server (or group) will be displayed.
5. From the Views pane of the Server Explorer (or Server Groups Browser), select an Application Configuration instance to edit.
6. If you wish, make edits to the Application Configuration. (See "Setting Application Configuration Values on a Server or Group" on page 630 in this chapter for more information.)
7. To preview the changes and see how they differ from the configuration file on the server, click **Preview**. The Comparison dialog box opens and shows any differences. Click **Close** when you are finished.
8. When are ready to apply the changes to the server, click **Push**.

Scheduling an Application Configuration Push

You can schedule an Application Configuration push to run a single time, or on a recurring schedule, such as daily, weekly, or monthly.

To schedule an Application Configuration push, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers.
2. Select a server or server group in the Navigation pane.
3. Select a server or group in the Content pane, right-click, and choose **Configure Applications**.
4. You now see the Server Explorer (or Server Groups Browser), with the Installed Configurations tab selected. All Application Configurations that have been attached to the server (or group) will be displayed.
5. From the View pane, click the plus (+) symbol to expand Application Configuration folder and select an Application Configuration instance.

6 Click **Schedule**.

7 In the Schedule Job dialog box, set the following parameters:

- **Schedule**: Choose to Run Once, Daily, Weekly, Monthly, or Custom. By default, the Schedule is set to Weekly.
- **Crontab String**: (This field appears only if you chose a custom schedule. If you did not choose Custom, then skip to the Start Time field below.) Enter a crontab string for date in this order:
  - Minute (0-59), Hour (0-23)
  - Day of the month (1-31)
  - Month of the year (1-12)
  - Day of the week (0-6 with 0=Sunday)

  Any of these fields can contain an asterisk * standing for all possible values. For example, the following crontab string runs the job at midnight every weekday:

  0 0 * * 1-5

  The crontab string can also handle serial (1,2,3,4) as well as range (1-5) values. For more information about using crontab strings, consult the crontab man pages on a Unix computer.
- **Start Time**: Select a time for the job to begin running. For one-time jobs, enter the full date and time. For weekly and monthly jobs, enter the time of day. You can enter the values by typing directly into the field using up or down arrows.
- **Time Zone**: Select a default time zone for the job execution time, or accept the default time zone. The default time shown converts the scheduled time to the time zone set in your user preferences. If you do not set a preferred time zone, the time zone is derived from the Opsware SAS core server (typically UTC).
- **Day** (Monthly only): Choose the day of the month to run this job.
- **Days To Run** (Weekly only): Choose the days of the week you want the job to run at the specified time.
- **Months to Run** (Monthly only): Choose the months during which you want the job to run.

8 In the Run Jobs between these Dates section, select a date range during which you would like the job to run.

- **Start**: Choose a start date for the date range.
- **End**: Choose an end date for the date range.
- **No End Date**: Choose if you want the job to run indefinitely.

9 In the Job Run Notification Email section, enter an email address to receive the results of the job. You can enter multiple email addresses separated by commas or spaces.

- **On Success**: Enter email addresses that will receive notifications of jobs that complete successully.
- **On Failure**: Email addresses that will receive notifications of jobs that failed to complete.

10 When you have finished setting the parameters, click **OK**.

Comparing Two Configuration Templates

To show the difference between two Application Configuration Templates, you can perform a compare operation between them.

To compare two Application Configuration Templates, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Software Library.

2 Select Application Configuration, and then select the Configuration Templates tab.

3 Hold down the CTRL key and select two Application Configuration Templates, right-click, and choose **Compare**.

4 The Comparison dialog box displays the difference between the two files. Use the arrows in the upper right of the dialog box to navigate through the two files. To indicate the differences, the Comparison feature uses the following colors:

- **Green**: This indicates that new information has been added.
- **Blue**: This indicates that information has been modified.
- **Red**: This indicates that information has been deleted.
- **Black**: This indicates no changes.

5 When you are finished viewing the differences, click **Close**.

Comparing a Template Against an Actual Configuration File

To show the difference between an Application Configuration Template and the actual file on the server (or group), perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers.
2. From the Navigation pane, select a server or server group.
3. In the Content pane, select a server or group, right-click, and choose **Configure Applications**.
4. The Installed Configurations tab will be selected. From the Views pane of the Server Explorer (or Server Groups Browser), select an Application Configuration instance.
5. If the Application Configuration contains more than one Application Configuration Template, then from the Template drop-down list in the Content pane, choose a Application Configuration Template to compare.
6. To preview the differences between the Application Configuration Template and the actual configuration file on the server, click **Preview**. The Comparison dialog box shows the differences between the Application Configuration Template and the actual configuration file. Use the arrow keys in the upper right of the dialog box to navigate through the two files. To illustrate the differences, the Comparison feature uses the following color scheme:
   - **Green**: This indicates that new information has been added.
   - **Blue**: This indicates that information has been modified.
   - **Red**: This indicates that information has been deleted.
   - **Black**: This indicates no changes.
7. When you are finished viewing the differences, click **Close** to close the Comparison dialog box.

Auditing an Application Configuration

After an Application Configuration has been pushed to a server, it is possible that the configuration file on the server becomes changed or altered, either intentionally or by accident. You can audit the Application Configurations on a server to determine if any of the configuration files on the server are out of sync with the values stored in your templates.

If an Application Configuration on a server is out of sync, the server that the Application Configuration is attached to will show the following icon in the server list inside the OCC Client:

For example, open the OCC Client and select the Servers feature icon. A list of all managed servers in your environment is displayed. If you scan the list of servers, you can see if any servers show the out of sync icon.

If a server shows this icon, run an Application Configuration audit to find out which configuration files on the server are out of sync with the Application Configuration.

To run an Application Configuration audit, perform the following steps:

1. Launch the OCC Client. From the Navigation pane, select Servers.
2. Select a Server or Server Group from the Navigation pane, and then select a server that shows the out of sync icon from the Content pane. From the **Actions** menu, select **Audit Application Configurations**. (You can also multiple-select and audit more than one out of sync server.)
3. You will be asked if you are sure you want to audit the Application Configuration on the selected managed server. Click **Yes** to run the audit.
4. The Job dialog box appears, showing the details of the audit. Make sure to deselect the Close when finished option at the bottom of the dialog box so the Job dialog box remains open after the audit job has run. Once the job has finished, look in the Completed Status section, and select the Success text. You see a list of servers in the Servers section to the right.
5. To view the audit details for a server, in the Servers section, select a server. Below in the Server Detail section, a list of all discrepancies shows which files are out of sync with Application Configuration Templates on the server. To view the Application Configuration, click **Configurations**. The Server Browser appears.
6. To troubleshoot the discrepancies, select the out of sync Application Configuration and its templates and click **Preview**. This will show you where the configuration file on the server differs from the values defined in the Application Configuration. Once

you have found the discrepancies, you can modify them as needed in the Value Set Editor, and then push the changes to the server. See "Comparing a Template Against an Actual Configuration File" on page 638 in this chapter for more information

Scheduling an Application Configuration Audit

You can schedule an Application Configuration audit to run a single time, or on a recurring schedule, such as daily, weekly, or monthly.

To schedule an Application Configuration audit, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Servers.

1 Select a server or server group from the Navigation pane, and then select a server from the Content pane.

2 From the **Actions** menu, select **Schedule Application Configuration Audit**.

3 In the Schedule Job dialog box, set the following parameters:

- **Schedule**: Choose to Run Once, Daily, Weekly, Monthly, or Custom. By default, the Schedule is set to Weekly.
- **Crontab String**: (This field appears only if you chose a custom schedule. If you did not choose Custom, then skip to the Start Time field below.) Enter a crontab string for date in this order:
  - Minute (0-59), Hour (0-23)
  - Day of the month (1-31)
  - Month of the year (1-12)
  - Day of the week (0-6 with 0=Sunday)

  Any of these fields can contain an asterisk * standing for all possible values. For example, the following crontab string runs the job at midnight every weekday:

  0 0 * * 1-5

  The crontab string can also handle serial (1,2,3,4) as well as range (1-5) values. For more information about using crontab strings, consult the crontab man pages on a Unix computer.
- **Start Time**: Select a time for the job to begin running. For one-time jobs, enter the full date and time. For weekly and monthly jobs, enter the time of day. You can enter the values by typing directly into the field using up or down arrows.

- **Time Zone**: Select a default time zone for the job execution time, or accept the default time zone. The default time shown converts the scheduled time to the time zone set in your user preferences. If you do not set a preferred time zone, the time zone is derived from the Opsware SAS core server (typically UTC).
- **Day** (Monthly only): Choose the day of the month to run this job.
- **Days To Run** (Weekly only): Choose the days of the week you want the job to run at the specified time.
- **Months to Run** (Monthly only): Choose the months during which you want the job to run.

4 In the Run Jobs between these Dates section, select a date range during which you would like the job to run.

- **Start**: Choose a start date for the date range.
- **End**: Choose an end date for the date range.
- **No End Date**: Choose if you want the job to run indefinitely.

5 In the Job Run Notification Email section, enter an email address to receive the results of the job. You can enter multiple email addresses separated by commas or spaces.

- **On Success**: Enter email addresses that will receive notifications of jobs that complete successully.
- **On Failure**: Email addresses that will receive notifications of jobs that failed to complete.

6 When you have finished setting the parameters, click **OK**.

Rolling Back to a Previous State

Every time you push an Application Configuration to a server, that push is saved in a configuration backup list. At any time, you can revert or rollback to a previous state of an Application Configuration in this list. This enables you to go back to a known configuration state for a specific Application Configuration.

To rollback an Application Configuration to a previous state, perform the following steps:

1 Launch the OCC Client. From the Navigation pane, select Servers.

2 From the Navigation pane, select a server or server group.

3 Select a server or server group from the Content pane, right-click and choose **Configure Applications**.

4 You now see the Server Explorer (or Server Groups Browser), with the Installed Configurations tab selected.

5 Select the Configuration History tab. A list of all Application Configuration pushes will display. You can sort this list by application name, configuration backup name, date created (when the Application Configuration was pushed), and by user.

If the list is empty, the Application Configuration has never been pushed to the server.

6 To rollback to a saved configuration, select a item in the list, and click **Revert**. This restores all configuration files to the state immediately after this backup was made. The original configuration files are also restored and suffixed with "_opsware_ backup".

Chapter 17: Configuration Tracking

IN THIS CHAPTER

This section discusses the following topics:

- Overview of Configuration Tracking
- Configuration Tracking Operations
- Customized Configuration Tracking Policies
- Manual Backups
- Backup History
- Restoration of Backed Up Files

Overview of Configuration Tracking

The Configuration Tracking feature of Opsware SAS allows you to monitor critical configuration files and configuration databases. When Opsware SAS detects a change in a tracked configuration file or configuration database, the system can perform a number of actions, including backing up the configuration file or sending an email to a designated individual or group. You use configuration tracking policies to identify the files to be tracked and actions to be taken when change is detected. A configuration tracking policy consists of one or more configuration tracking policy entries that specifies the configuration file, the directory of configuration files, or the configuration database that you want to track.

The Configuration Tracking feature is designed for flexibility. For example, you can set configuration tracking defaults for a node that contains a particular software application, and all servers attached to that node automatically get those defaults. You can also quickly deploy the common configuration tracking defaults to a large number of servers in your Opsware managed environment or create a specific policy for a single server.

Configuration Tracking allows you to recover from many problems caused by changes to configuration files. Using Automated Configuration Tracking, you can identify which tracked configuration files have changed, thus helping you identify the potential source of a problem. If you back up your configuration files with the Configuration Tracking feature, you can quickly restore the changed configuration files to a previous version.

You can also view a detailed history of all backup activity. This history includes a list of all tracked files that have been backed up and what types of backups occurred. If the backed up configuration files are text-based, you can download the files from the backup history and compare them to determine what specific changes have been made.

The Configuration Tracking feature is not a general-purpose backup solution. Configuration Tracking is designed to monitor text-based configuration files and specific types of configuration databases. The number and size of files that can be monitored on any managed server is limited.

File Types Supported

You can use the Configuration Tracking feature with the following types of files:

- Text-based configuration files
- The COM + Registration Database (Windows 2000)
- The IIS Metabase
- Windows Registry keys

Configuration Tracking Operations

This section discusses the following topics:

- Change Detection
- Types of Actions Performed
- Types of Backups Performed
- Email Automated Configuration Tracking and Logging Actions
- Creating the Email Notification List

Change Detection

All servers that Opsware SAS manages have an Opsware Agent installed on them. On servers that use Configuration Tracking, every four hours the Opsware Agent inspects the configuration files and databases that you select to track.

The Opsware Agent computes an MD5 checksum to determine if the contents of a tracked file have changed. (Any change to the contents of the file results in a change to the MD5 checksum.) If the contents of the file have changed, the action that you specify in the tracking policy is performed. For example, if you create an entry in your tracking policy for the `/etc/passwd` file and select backup as the action to be taken, the file is backed up when the Opsware Agent discovers a change in the `/etc/passwd` file.

The creation or deletion of a tracked file (or files inside a tracked directory) also counts as change and triggers a policy's action. (There are some exceptions; see the *Opsware® SAS Configuration Guide* for more information.)

Changes to the properties of a tracked file or directory (such as changes to permissions or time stamps) do not count as a change. When a file or directory is backed up, however, its properties are backed up as well.

The first time that a tracking policy is deployed, all targets are considered changed. The Opsware Agent is encountering the files for the first time, and all of the policy's actions are triggered.

Types of Actions Performed

Opsware SAS can perform the following actions when change is detected in a tracked configuration file or configuration database:

- Back up
- Send email to addresses specified in the policy entry
- Send email to a designated notification group specified by a custom attribute
- Create an entry in the server's standard system log, (the syslog on Unix servers and the event log on Windows servers)

Types of Backups Performed

If you selected backup as the action for a tracked configuration file, the two general types of backups that can occur are incremental backups and full backups.

Incremental Backups

During an incremental backup, only targets that have changed since the last backup (and that have been selected to be backed up) are backed up.

An incremental backup occurs automatically when the Opsware Agent detects change in a tracked file that is selected to be backed up. (The Opsware Agent checks for change every four hours.)

Incremental backups also occur before and after you restore a previous version of a backed up configuration file to a server. These backups allow you to rollback the restored files.

Full Backups

During a full backup, all tracked configuration files that were selected to be backed up are backed up, not just the files that have changed.

Once a week, the Opsware Agent on a server checks to see if any files have changed since the last full backup. If any files have changed, Opsware SAS performs a new full backup. If no files have changed, the full backup does not take place.

You can also force Opsware SAS to perform a full backup on a server by selecting the Perform Manual Backup option. (See "Manual Backups" on page 657 in this chapter for more information.)

See "Backup History" on page 658 in this chapter for information about Backup types.

Backups are stored in the Software Repository until you delete them. You should delete old backups periodically, especially if you are backing up a large number of files that change frequently. See "Deleting Backups" on page 665 in this chapter for information about the procedure for deleting backups.

Email Automated Configuration Tracking and Logging Actions

You can choose to have email sent when a monitored target changes. The following example shows the text of an email generated when a tracked file changed.

```
From: <configurationtracking@yourcompany.com>
Date: Thu Jan 16, 2003  5:40:11 PM US/Pacific
To: <joe@yourcompany.com>
Subject: athena.cust.com: Configuration Tracking CHANGE
notification
```

```
Configuration Tracking has detected a CHANGE event
Host: athena.cust.com
Object: /db/file111
```

The email specifies the name of the server and the name of the object that changed. The object can be a file, a directory, or a configuration database.

If you are monitoring a directory target, you receive email about the directory itself and about changes to the files in the directory (except when a file is deleted.) For example, if three new files are created in a directory, you would receive four emails, one for the directory and three for the new files.

If you selected the logging action, an entry is made to the server's standard system log when a change is detected. You select the type of log entry that you want to have written. Opsware SAS uses the following three standard entry types:

- Info
- Warning
- Error

How the entry types are identified is system-dependent. For example, on most systems, Warning entries are identified by the word warning. In some systems, however, a number is used to identify the log entry type.

The following example shows a warning log entry written on a Solaris Server:

```
Jan  8 00:05:25 athena.cust.com Configuration Tracking:
[ID702911
local0.warning] Configuration Tracking: /other/otherfile1 :
Event CHANGE occurred
Jan  8 00:05:25 athena.cust.com Configuration Tracking: [ID
702911
local0.warning] Configuration Tracking: /other/otherfile1 :
Event CHANGE occurred
```

Creating the Email Notification List

Sending email to a server’s backup notification list is one of the actions that you can select in a tracking policy entry. The email notification list is a list of email addresses that you define for the following custom attribute:

```
backup_notification_email
```

This attribute can be set on the server itself or on the customer to which the server is attached. Setting the attribute at the customer node level allows you to use the same email notification list for all servers that belong to the same customer (assuming that these servers have all been attached to the same customer and do not have the `backup_notification_list` set on the servers themselves). See the *Opsware® SAS Configuration Guide* for more information about setting custom attributes for customers.

Search Order for Email Notification List Attribute

On a server that has a policy that includes the Email Notification List for Server action selected, the server searches for the backup_notification_email attribute in the following order:

- Server
- Customer

After the custom attribute is found, its value is used (for example, the email address of the notification list) and the search is concluded. If, for example, the backup_notification_email attribute is set on a server, the server's email notification list is used, even if the server is assigned to a customer that has a different backup_notification_email attribute.

Format of Email Notification List

The notification list can contain multiple email addresses. The email address must be formatted as a comma-separated list.

Customized Configuration Tracking Policies

This section provides information on customizing configuration tracking policies within Opsware SAS and contains the following topics:

- Node-Based Entries and Server-Based Entries
- Policies for Customizing Multiple Servers
- Adding or Editing Customized Tracking Policy Entries
- Disabling Customized Tracking Policy Entries
- Enabling Customized Tracking Policy Entries
- Viewing a Server's Tracking Policy
- Reconciling Customized Tracking Policies

- Enabling and Disabling Configuration Tracking on Servers

You can customize the tracking policy for a server or selected group of servers. Ordinarily, a server gets its tracking policy from the nodes to which it is attached. In some cases, however, you might need to customize the tracking policy for a particular server or set of servers. If, for example, an application on one server is using an optional configuration file, you can customize the tracking policy for that server so that the optional configuration file is monitored.

Node-Based Entries and Server-Based Entries

A tracking policy entry created for a specific server or selected group of servers is called a server-based tracking policy entry. A tracking policy entry that a server obtains from a node that it is attached to is called a node-based tracking policy entry.

Table 17-1 shows the differences between the actions that you can perform on node-based tracking policy entries and server-based tracking policy entries.

Table 17-1: Entry Types

ENTRY TYPE	EDIT	DELETE	ENABLE/DISABLE
Server-based	Yes	Yes	No
Node-based	No	No	Yes

The restrictions for node-based policy entries hold true only when you are customizing the tracking policy for a server or set of servers. You cannot edit a policy entry on a server that is obtained from a node. You can, however, edit the policy of the node itself. (See the *Opsware® SAS Configuration Guide* for more information about selecting the node and editing the node’s tracking policy.)

To determine if a tracking policy entry is node-based or server-based, check the Source column in the Track Configurations: Customize Tracking Policies page. Figure 17-1 shows the policy entries for a single selected server.

Figure 17-1: Viewing Configuration Tracking Policy

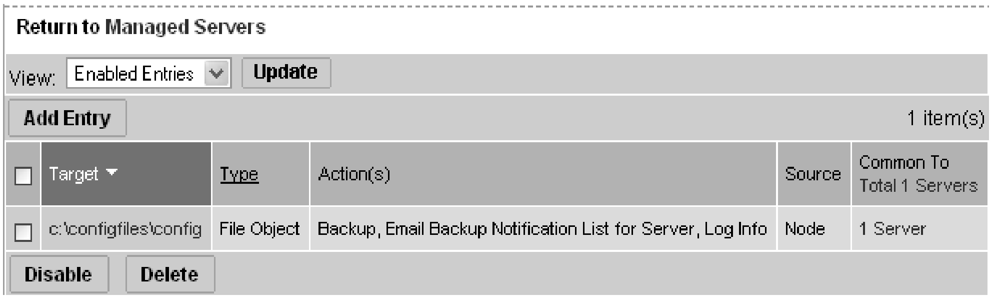

□	Target ▼	Type	Action(s)	Source	Common To Total 1 Servers
□	c:\configfiles\config	File Object	Backup, Email Backup Notification List for Server, Log Info	Node	1 Server

You can edit only tracking policy entries that have been created on the server level. You cannot edit tracking policy entries that have been obtained from a node.

Policies for Customizing Multiple Servers

When you customize a tracking policy, you can select one or multiple servers. Here are some considerations when you customize the policies of multiple servers.

- When you select a group of servers that contain both Unix-based (for example, Solaris, HP-UX, Linux, and so forth) and Windows-based servers, you cannot add tracking policy entries. (You can always add entries when you select a single server.)
- When you select a group of servers made up entirely of either Unix-based servers or entirely of Windows-based servers, you have the option of making new tracking policy entries. The policy entries are added to all servers that you select.
- When you select a group of servers made up entirely of either Unix-based servers or of Windows-based servers, you have the option of editing a tracking policy entry. Any policy entry that can be edited can also be added to all selected servers (if it is not already common to all of them).

It can take 20 seconds or longer to customize polices when you select six or more servers.

Adding or Editing Customized Tracking Policy Entries

Use the following procedure to add or edit tracking policy entries for a server or set of servers running the same general type of operating system. If you select a set of servers running both Windows and Unix operating systems, you cannot add tracking policy entries.

1. From the navigation panel in the Opsware Command Center, click Servers.

2. Click **Server Search** to search for the server whose policy you want to edit. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)

3. Select the server or set of servers that you want to add server-based tracking policy entries to, as Figure 17-2 shows.

Figure 17-2: Selecting Servers

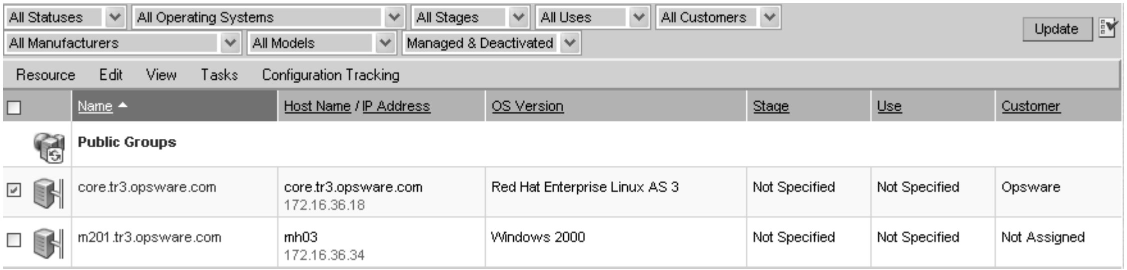

□	Name ▲	Host Name / IP Address	OS Version	Stage	Use	Customer
	Public Groups					
☑	core.tr3.opsware.com	core.tr3.opsware.com 172.16.36.18	Red Hat Enterprise Linux AS 3	Not Specified	Not Specified	Opsware
□	m201.tr3.opsware.com	mh03 172.16.36.34	Windows 2000	Not Specified	Not Specified	Not Assigned

4. From the Configuration Tracking drop-down menu, select Edit Tracking Policies. The Configuration Tracking: Edit Tracking Policies page appears, as Figure 17-3 shows.

Figure 17-3: Configuration Tracking: Edit Tracking Policies Page

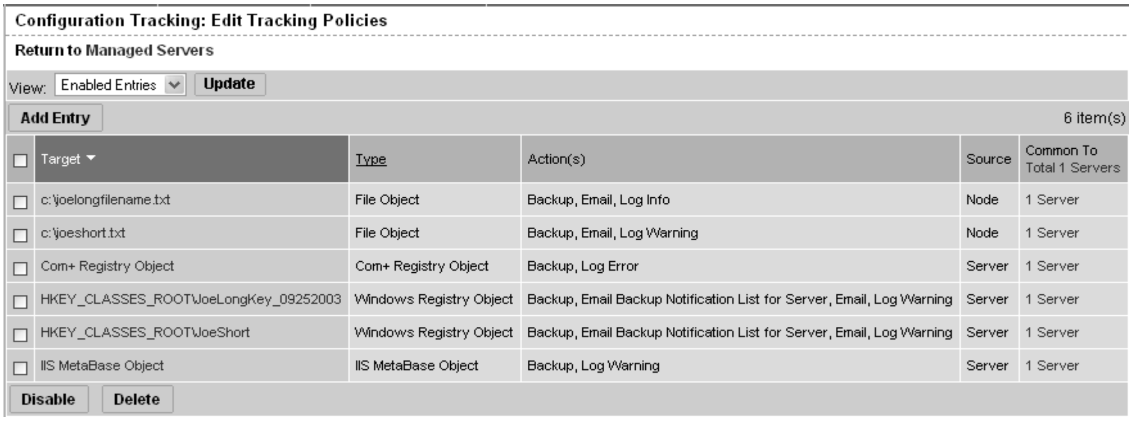
Configuration Tracking: Edit Tracking Policies

Return to Managed Servers

View: Enabled Entries

6 item(s)

□	Target ▼	Type	Action(s)	Source	Common To Total 1 Servers
□	c:\joelongfilename.txt	File Object	Backup, Email, Log Info	Node	1 Server
□	c:\joeshort.txt	File Object	Backup, Email, Log Warning	Node	1 Server
□	Com+ Registry Object	Com+ Registry Object	Backup, Log Error	Server	1 Server
□	HKEY_CLASSES_ROOT\JoeLongKey_09252003	Windows Registry Object	Backup, Email Backup Notification List for Server, Email, Log Warning	Server	1 Server
□	HKEY_CLASSES_ROOT\JoeShort	Windows Registry Object	Backup, Email Backup Notification List for Server, Email, Log Warning	Server	1 Server
□	IIS MetaBase Object	IIS MetaBase Object	Backup, Log Warning	Server	1 Server

5 Click **Add Entry** to create a new policy entry. To edit an existing entry, click the link for the entry in the Target field.

6 Define the target, select the type, and select the actions to be performed on the target.

Table 17-2 describes each of the selections that you must make.

Table 17-2: Editing Configuration Tracking Policies

FIELD	DESCRIPTION
Type	**File**: Monitor the file specified in the target field. **Directory**: Monitor all the files in the directory specified in the target field. The following types are available only for Windows servers: **Windows Registry**: Specify key in target field. **IIS Metabase**: Entire Metabase is monitored; do not specify target. **COM + Registration Database**: Entire Registry is monitored; do not specify target.
Target	If you selected the file type, specify the full path (including the drive letter on Windows servers) of the file that you want to monitor. If you selected the directory type, specify the full path of the directory (including the drive letter on Windows machines) that you want to monitor. You also have the option of monitoring subdirectories. (Select the include subdirectories check box.) If you select the file or directory type, you can use wildcards in the target. (See the *Opsware® SAS Configuration Guide* for more information about Configuration Tracking Policy Targets and Wildcards.) If you selected the Windows Registry type, specify the Windows registry key. This key and all its subkeys are backed up. Use standard syntax for Windows Registry keys (For example, `HKEY_LOCAL_MACHINE\SOFTWARE)` If you selected the IIS Metabase or COM + Registration Database type, you do not specify the target.

Table 17-2: Editing Configuration Tracking Policies (continued)

FIELD	DESCRIPTION
Actions (you can select multiple options)	**Backup**: Back up the specified file, directory, COM + Registration Database, Windows Registry keys, or IIS Metabase. **Email Backup Notification List for Server**: Send an email to the backup notification list for the selected server. (Not available for Windows Registry.) See “Creating the Email Notification List” on page 647 in this chapter for more information. **Email**: Send an email to the address or addresses specified in this field when a change is detected. Use a comma-separated list for multiple email addresses. (Not available for Windows Registry.) **Log**: Add an entry to the server’s system log when a change is detected. (Not available for Windows Registry.) You can choose to write the following types of log entries to the server’s system log: Info Warning Error

7 Click **Save** to add the entry to the tracking policy.

8 If you want continue to add entries to the tracking policy, click **Add Entry** and repeat this procedure.

Changes do not take effect until you perform a Configuration Tracking policy reconcile. See “Reconcile” on page 699 in Appendix A for more information.

Disabling Customized Tracking Policy Entries

You can only disable node-based tracking policy entries.

Perform the following steps to disable customized tracking policy entries:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Edit Tracking Policies.

   The Configuration Tracking: Edit Tracking Policies page appears.
4. Select the tracking policy entry or entries that you want to disable. (Make sure that all your selected tracking policy entries are node-based.) If all the tracking policy entries displayed are node-based and you want to disable all the displayed tracking policies, you can select the first check box to select all tracking policy entries.
5. Click **Disable** to disable the tracking policy entries that you selected.

Changes do not take effect until you perform a Configuration Tracking policy reconcile.

Enabling Customized Tracking Policy Entries

You can re-enable any previously disabled node-based tracking policy entry.

Perform the following steps to re-enable tracking policies:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Edit Tracking Policies.

   The Configuration Tracking: Edit Tracking Policies page appears.
4. From the **View** menu, select **Disabled Entries** and then click **Update**.
5. Select the tracking policy entry or entries that you want to disable. (Make sure that all your selected tracking policy entries are node-based.) If all the tracking policy entries displayed are node-based and you want to disable all the displayed tracking policies, you can select the first check box to select all tracking policy entries.
6. Click **Enable** to disable the tracking policy entries that you select.

Changes do not take effect until you perform a Configuration Tracking policy reconcile.

Viewing a Server's Tracking Policy

Perform the following steps to view a server's tracking policy:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Edit Tracking Policies. The tracking policy entries display.
4. Click the Tracking Policy link to display the server's tracking policy.

Reconciling Customized Tracking Policies

When you create or edit any customized tracking policy entries, the entries are not deployed to your servers until you perform a configuration tracking reconcile.

A configuration tracking reconcile is not the same as a standard Opsware SAS reconcile. Performing a standard reconcile does not deploy your tracking policies to your servers.

A configuration tracking reconcile deploys all configuration polices, not just the customized policies. Both node-based polices and customized policies are deployed when you perform a configuration tracking reconcile.

Perform the following steps to reconcile the tracking policy to the servers whose policies you customized:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Reconcile Tracking Policies.

   The Track Configurations: Preview Reconcile page appears and displays the progress of the test reconcile.

4. After the test reconcile completes successfully, click **Reconcile** to perform the actual reconcile operation.

   The Track Configurations: Reconcile page appears.

5. If you want to see more information about the changes made during the reconcile operation, click **View Details**. Otherwise, click **Done**.

Enabling and Disabling Configuration Tracking on Servers

You can enable or disable Configuration Tracking on any server or set of servers in your Opsware-managed environment.

Disabling Configuration Tracking is not the same as disabling an individual tracking policy entry (See "Disabling Customized Tracking Policy Entries" on page 654 in this chapter for more information"). Disabling Configuration Tracking stops all configuration tracking activity on the selected server.

Disabling configuration tracking, however, does not change a server's tracking policy in any way. If you later re-enabled configuration tracking on the server, the server still has the same tracking policy that it did before you disabled it.

By default, Configuration Tracking is disabled on all managed servers. It is not necessary, however, to manually enable configuration tracking in order to turn on the Configuration Tracking feature. Configuration Tracking is automatically enabled on a server when you reconcile its configuration tracking policies, and you must perform a reconcile in order to deploy tracking policies.

If you want configuration tracking to remain disabled on a server, be careful not to perform a configuration tracking reconcile on the server. (You can, however, still perform a regular reconcile.)

Perform the following steps to enable or disable configuration tracking:

1. From the navigation panel in the Opsware Command Center, click Servers.

2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)

3 Under the Configuration Tracking drop-down menu, select Enable/Disable. A list that displays the state (Enabled/Disabled) of the servers that you selected, as Figure 17-4 shows.

Figure 17-4: Enable/Disable Configuration Tracking Page

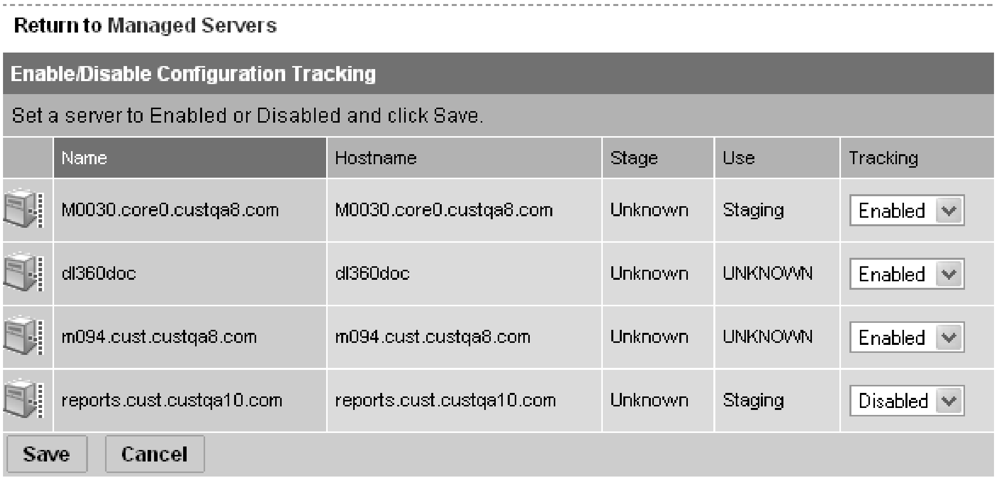

Enable/Disable Configuration Tracking

Set a server to Enabled or Disabled and click Save.

	Name	Hostname	Stage	Use	Tracking
	M0030.core0.custqa8.com	M0030.core0.custqa8.com	Unknown	Staging	Enabled
	dl360doc	dl360doc	Unknown	UNKNOWN	Enabled
	m094.cust.custqa8.com	m094.cust.custqa8.com	Unknown	UNKNOWN	Enabled
	reports.cust.custqa10.com	reports.cust.custqa10.com	Unknown	Staging	Disabled

4 Under the Tracking field, select Enabled or Disabled to enable or disable configuration tracking on the selected server.

5 Click **Save** to commit the changes.

Manual Backups

On a server or set of servers, you can perform a manual backup of all tracked configuration files and databases for which you have selected the backup action. Manual backups can be useful as a precaution before making changes to configuration files. If a problem arises, you can then immediately restore a backed-up configuration file or database to its previous state.

Manual backups are full backups. All tracked configuration files and databases for which the backup action has been selected are backed up, not just the files that have changed.

Performing Manual Backups

Perform the following steps to perform a manual backup:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select Perform Backup Now.
4. If you do not want to use the default backup name (Manual Backup), type a new name in the backup name field. If you provide a backup name, you can later perform a search for backup names to find this backup point. Backup names do not have to be unique.
5. Click **Start Backup**. The backup progress displays.
6. When the backup is completed, you can click **View Details** to review the list of configuration files or configuration databases that have been backed up.

Backup History

Opsware SAS provides a detailed history of backup activity as well as search capabilities to find backup points by backup name and to find backed up files by file name. This section contains the following topics:

- Viewing the Backup History
- Viewing the List of Backup Events
- Types of Backup Events
- Backup History Search Options
- Backup Info and Backup Manifest
- File Info and File Versions
- Deleting Backups

Viewing the Backup History

Perform the following steps to view the backup history:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. From the Configuration Tracking drop-down menu, select View Backup History.

   The Track Configurations: View Group Backup History page displays. This page displays the backup activity for the servers that you select. Figure 17-5 shows a sample backup history.

Figure 17-5: Backup History

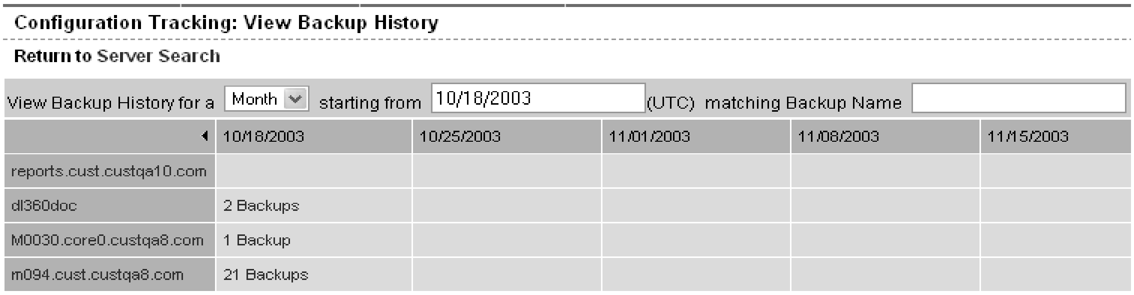

Configuration Tracking: View Backup History

Return to Server Search

View Backup History for a Month starting from 10/18/2003 (UTC) matching Backup Name

◀	10/18/2003	10/25/2003	11/01/2003	11/08/2003	11/15/2003
reports.cust.custqa10.com					
dl360doc	2 Backups				
M0030.core0.custqa8.com	1 Backup				
m094.cust.custqa8.com	21 Backups				

The number of backups that occurred on a particular date displays as a link in the server's date column.

The number of backups refers to the number of backup events (or backup points) when a particular type of backup took place. It does not refer to the number of files backed up. For example, if a server displays 3 backups in a date column, it could refer to the following three backup events:

- A scheduled incremental backup that backed up four files
- A second scheduled incremental backup that backed up 10 files
- A manual backup that backed up 30 files

Viewing the List of Backup Events

To view the list of individual backup events, click the link that displays the number of backups in the desired date column. Figure 17-6 shows a page with a list of backup events.

Figure 17-6: Backup Events

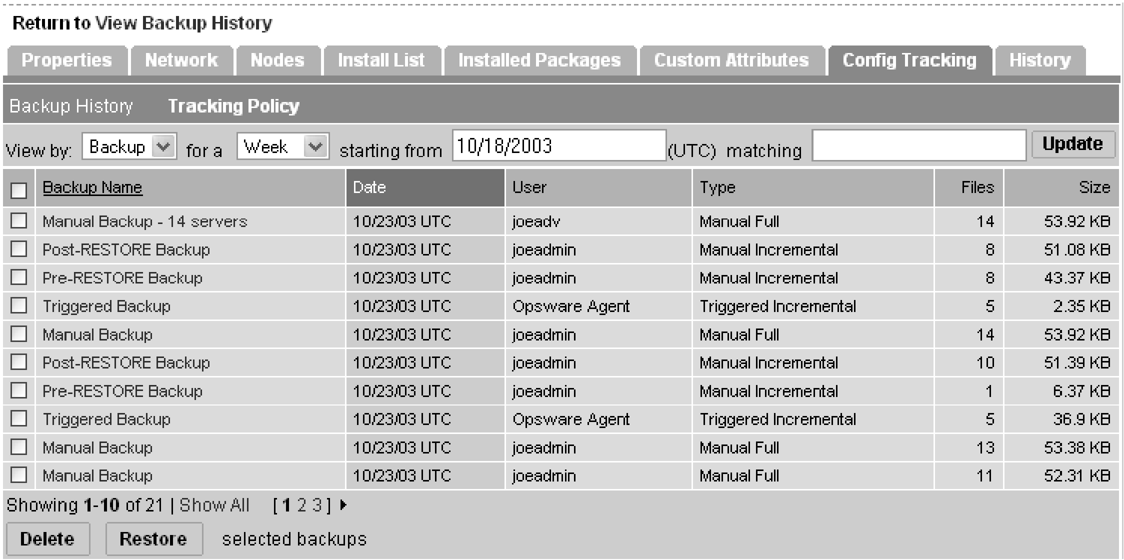

View by: Backup for a Week starting from 10/18/2003 (UTC) matching

□	Backup Name	Date	User	Type	Files	Size
□	Manual Backup - 14 servers	10/23/03 UTC	joeadv	Manual Full	14	53.92 KB
□	Post-RESTORE Backup	10/23/03 UTC	joeadmin	Manual Incremental	8	51.08 KB
□	Pre-RESTORE Backup	10/23/03 UTC	joeadmin	Manual Incremental	8	43.37 KB
□	Triggered Backup	10/23/03 UTC	Opsware Agent	Triggered Incremental	5	2.35 KB
□	Manual Backup	10/23/03 UTC	joeadmin	Manual Full	14	53.92 KB
□	Post-RESTORE Backup	10/23/03 UTC	joeadmin	Manual Incremental	10	51.39 KB
□	Pre-RESTORE Backup	10/23/03 UTC	joeadmin	Manual Incremental	1	6.37 KB
□	Triggered Backup	10/23/03 UTC	Opsware Agent	Triggered Incremental	5	36.9 KB
□	Manual Backup	10/23/03 UTC	joeadmin	Manual Full	13	53.38 KB
□	Manual Backup	10/23/03 UTC	joeadmin	Manual Full	11	52.31 KB

Showing **1-10** of 21 | Show All [**1** 2 3] ▸

Delete Restore selected backups

Types of Backup Events

In addition to information such as the date and time the backup occurred, the list of backup events indicates the type of backup. Table 17-3 describes the type of backup events that can occur.

Table 17-3: Types of Backup Events

TYPE	DESCRIPTION
Triggered Full	The automatic weekly backup of all tracked files for which the backup option was selected. This takes places only if any relevant files have changed since the last full backup.
Manual Incremental	A backup of all changed files (for which the backup option was selected) that occurs before and after a restoration or a rollback.

Table 17-3: Types of Backup Events (continued)

TYPE	DESCRIPTION
Triggered Incremental	An automatic backup that occurs when the Opsware Agent detects a change to a tracked file for which the backup option was selected. Only changed files are backed up.
Manual Full	A full backup initiated by the user of all tracked files for which the backup option was selected.

Backup History Search Options

By default, you see the backup history for the past week for the servers that you selected.

If you want to display a different date range, you have the following options:

- Display the backup history for different date range by selecting a different option in the "View Backup History for a" box.
- Use the "starting from" field to search for backup history from a past date until the current date.

You can use the matching field to search for a backup name within the selected date range. Wildcards are allowed (the * and ? characters). If you do not type the full name of the backup, you must use a wildcard for any missing parts of the name.

The results of this search show the date and number of backup names on that date that match the * pattern. When you click the link for the number of backups, a list of the matching backup names displays.

Backup Info and Backup Manifest

When you click a backup name (such as Scheduled Backup) anywhere in the backup history, the Backup Info and Backup Manifest tabs display.

Backup Info Tab

The Backup Info tab displays general information about a backup event. This information includes the name of the backup, the date and time of the backup, and the policies that triggered the backup event.

If the polices are node-based, they are identified by the name of the Node. All customized policies are identified as Server Policy, as Figure 17-7 shows.

Figure 17-7: Track Configurations: View Backup Triggered Backup Page

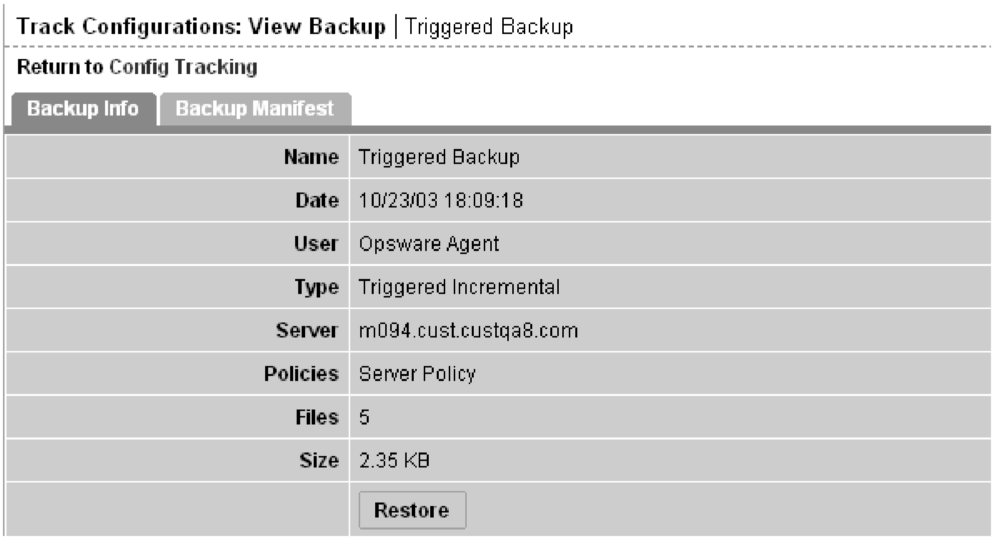

Track Configurations: View Backup | Triggered Backup

Return to Config Tracking

Name	Triggered Backup
Date	10/23/03 18:09:18
User	Opsware Agent
Type	Triggered Incremental
Server	m094.cust.custqa8.com
Policies	Server Policy
Files	5
Size	2.35 KB
	Restore

From the Backup Info tab, you can also place the files that were backed up into the Restore Queue. See "Restoring Backups" on page 669 in this chapter for information about how to use this feature.

Backup Manifest Tab

The Backup Manifest tab displays the list of files that were backed up during the selected backup event, as Figure 17-8 shows.

Figure 17-8: Backup Manifest Tab

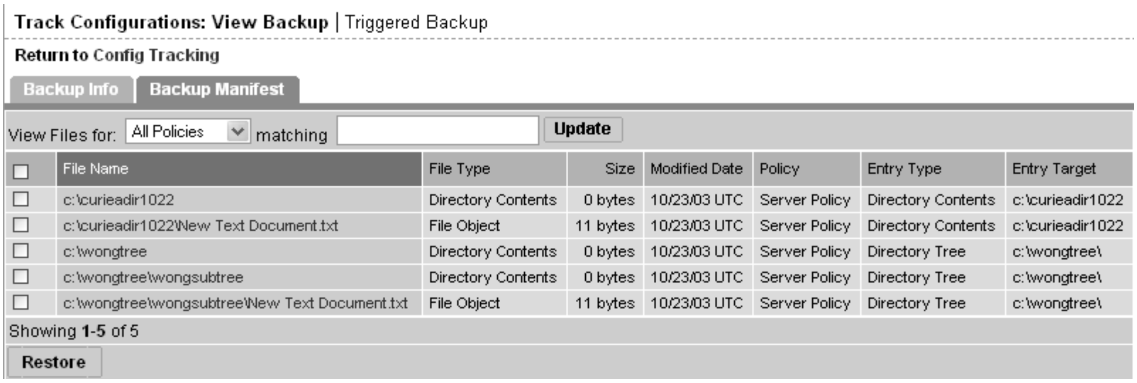

Track Configurations: View Backup | Triggered Backup

Return to Config Tracking

View Files for: All Policies matching Update

□	File Name	File Type	Size	Modified Date	Policy	Entry Type	Entry Target
□	c:\curieadir1022	Directory Contents	0 bytes	10/23/03 UTC	Server Policy	Directory Contents	c:\curieadir1022
□	c:\curieadir1022\New Text Document.txt	File Object	11 bytes	10/23/03 UTC	Server Policy	Directory Contents	c:\curieadir1022
□	c:\wongtree	Directory Contents	0 bytes	10/23/03 UTC	Server Policy	Directory Tree	c:\wongtree\
□	c:\wongtree\wongsubtree	Directory Contents	0 bytes	10/23/03 UTC	Server Policy	Directory Tree	c:\wongtree\
□	c:\wongtree\wongsubtree\New Text Document.txt	File Object	11 bytes	10/23/03 UTC	Server Policy	Directory Tree	c:\wongtree\

Showing **1-5** of 5

Restore

Two types of objects are backed up, as seen in the File Type field.

- File Object in the File Type field indicates that a file has been backed up. You can use the entry to restore the file.
- Directory Contents in the File Type field indicates that a directory object has been backed up. The directory object is the directory itself, and not the contents of the directory. You can use this entry to restore the directory, but you must select the files inside the directory to restore the contents of the directory.

The Entry Type field can have three possible values. The Entry Type identifies the target type in the policy entry that caused the file or directory object to be backed up.

- File Object in the Entry Type field indicates that the file was backed up as the result of a file target type.
- Directory Contents in the Entry Type field indicates that the file or directory object was backed up as result of a directory target type.
- Directory Tree in the Entry Type field indicates that the file or directory object was backed up as a result of a directory target type with the include subdirectories option selected.

File Info and File Versions

When you click a file or directory name anywhere in the backup history (such as in the Backup Info tab), the File Info and File Versions tabs display, as Figure 17-9 shows.

Figure 17-9: File Info

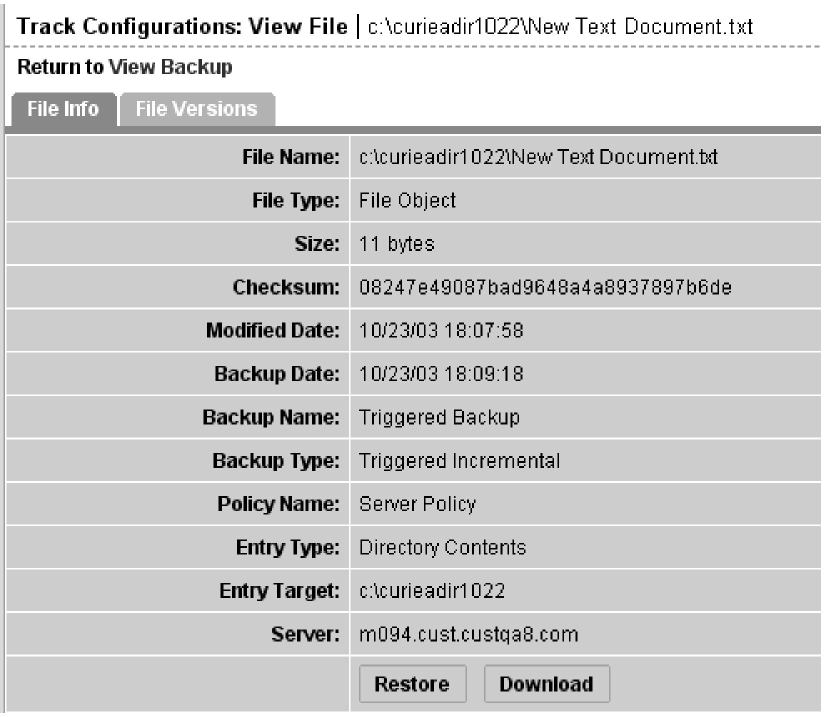

Track Configurations: View File | c:\curieadir1022\New Text Document.txt

Return to View Backup

File Name:	c:\curieadir1022\New Text Document.txt
File Type:	File Object
Size:	11 bytes
Checksum:	08247e49087bad9648a4a8937897b6de
Modified Date:	10/23/03 18:07:58
Backup Date:	10/23/03 18:09:18
Backup Name:	Triggered Backup
Backup Type:	Triggered Incremental
Policy Name:	Server Policy
Entry Type:	Directory Contents
Entry Target:	c:\curieadir1022
Server:	m094.cust.custqa8.com

The File Info tab displays information about the specific file that you selected. The Policy Name refers to the policy that caused the file to be backed up. The Policy Name is either the name of the node whose tracking policy triggered the backup, or the Server Policy if the server's customized tracking policy caused the file to be backed up.

You can place the file that you selected in the Restore Queue by clicking Restore. See "Restoration of Backed Up Files" on page 666 in this chapter for more information.

The File Versions tab displays a list of the backed up versions of the file, as Figure 17-10 shows.

Figure 17-10: File Versions

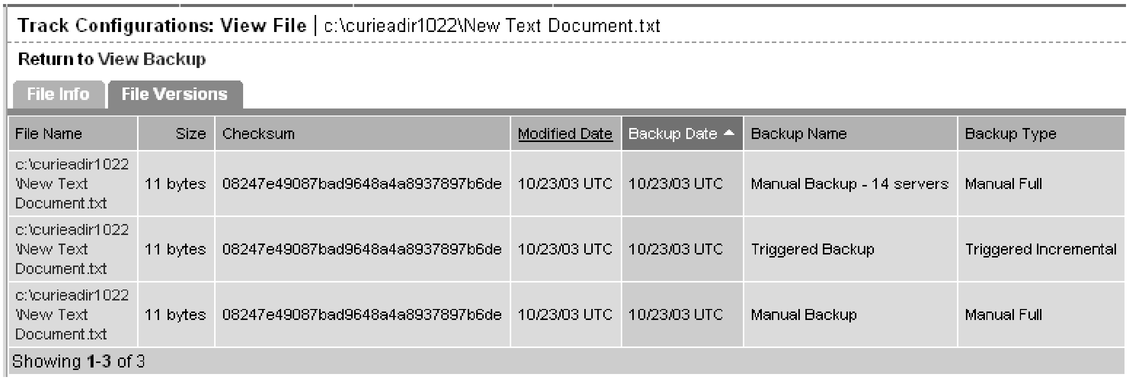

Track Configurations: View File | c:\curieadir1022\New Text Document.txt

Return to View Backup

File Name	Size	Checksum	Modified Date	Backup Date	Backup Name	Backup Type
c:\curieadir1022 \New Text Document.txt	11 bytes	08247e49087bad9648a4a8937897b6de	10/23/03 UTC	10/23/03 UTC	Manual Backup - 14 servers	Manual Full
c:\curieadir1022 \New Text Document.txt	11 bytes	08247e49087bad9648a4a8937897b6de	10/23/03 UTC	10/23/03 UTC	Triggered Backup	Triggered Incremental
c:\curieadir1022 \New Text Document.txt	11 bytes	08247e49087bad9648a4a8937897b6de	10/23/03 UTC	10/23/03 UTC	Manual Backup	Manual Full

Showing **1-3** of 3

It displays the backup name and backup type of the file. If you click any of the files, the File Info for the file is displayed. You can therefore use the File Versions tab to select a different version of the file, and then restore it or download it from the File History.

Deleting Backups

Backups remain in the backup history until you delete them or the server is deactivated. Backups are stored in the Software Repository. See “Software Package Management in the Model Repository” on page 550 in Chapter 14 for information about the Software Repository. You should delete old backup events periodically to reclaim disk storage from the Software Repository.

You can delete entire backup events (identified by a backup name); you cannot delete individual files from the backup history.

See the *Opsware® SAS Administration Guide* for more information about mass deletion of backup files.

Perform the following steps to delete backup points from your backup history:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. Select the server or set of servers that have backup points that you want to delete.
4. Search for a date that contains the backup points that you want to delete. Click the link with the number of backup events that occurred on that date.

5 Select the check boxes for the backup events that you want to delete from the backup history.

6 Click **Delete**.

The Delete Backup Confirmation page appears.

7 Click **Delete Backups**. The backups are now deleted.

Restoration of Backed Up Files

You can use the Configuration Tracking feature to restore configuration files or databases that have been backed up. You can restore all files that were backed up during a backup point, or you can select and restore individual files from a backup point. Opsware SAS allows you to rollback a restoration (for example, return your server's tracked files to the state immediately before the restoration).

Overview of Procedures Used to Restore Backed Up Files

This section presents an overview of the procedures you use to restore backed-up configuration files and databases. See "Restoring Backups" on page 669 in this chapter for information about the step-by-step procedure.

To restore backed-up configuration files and databases, view the backup history for a server or set of servers. The backup history displays entries for each date when a backup occurred. As Figure 17-11 shows, the entry for the date displays how many backup points occurred on that date.

Figure 17-11: Backup Points

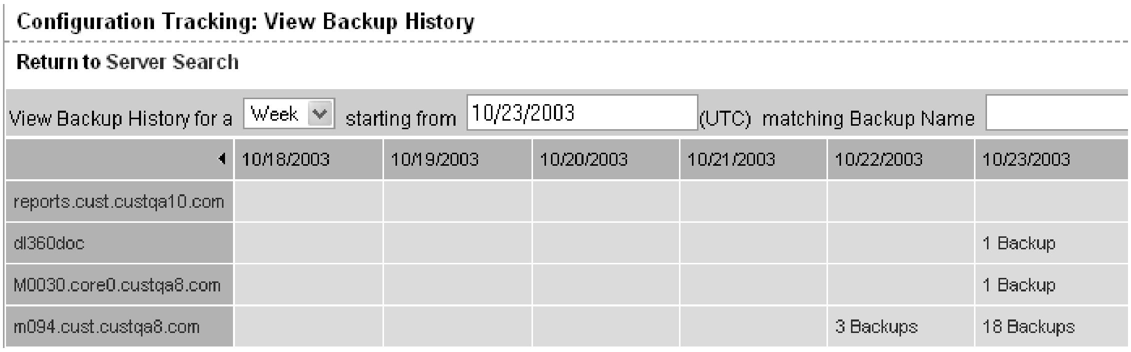

Configuration Tracking: View Backup History

Return to Server Search

View Backup History for a Week starting from 10/23/2003 (UTC) matching Backup Name

◂	10/18/2003	10/19/2003	10/20/2003	10/21/2003	10/22/2003	10/23/2003
reports.cust.custqa10.com						
dl360doc						1 Backup
M0030.core0.custqa8.com						1 Backup
m094.cust.custqa8.com					3 Backups	18 Backups

You select the desired backup date, and you can then select one or more backup points that occurred on that date. Before you restore the backups, you can review all files that were backed up at your selected backup points. You can either choose to restore all backed up files from your selected backup points, or you can select the files individually.

To select the backup, click on the link in the date field that displays the number of backups that were performed on that date.

Restore Queue

This section provides information on the restore queue within Opsware SAS and contains the following topics:

- Storing files in the Restore Queue
- Incremental Backups for Restoration and Rollbacks
- Entries for Directories in the Backup History
- File Not Found Entries
- Restoring Backups
- Rolling Back Restored Files

Storing files in the Restore Queue

When you click **Restore**, the files that you selected are placed in the Restore Queue and the View Restore Queue and Perform Restore page appears. You can then review and select files from all the backup points that you selected.

If you do not want to restore the files at this time, you can perform other actions and the files will remain in the Restore Queue as long as your session is active (even if you leave this page). You can also return to the backup history and select other files to put in the Restore Queue.

When you are ready to restore the backups, return to the Track Configurations: Select a Task page and click the **View Restore Queue** and the Perform Restore link.

Incremental Backups for Restoration and Rollbacks

To make rollbacks possible, Opsware SAS performs two automatic incremental backups. The first backup occurs immediately before the restoration, and the second occurs immediately after the restoration. Similarly, in order to undo a rollback, Opsware SAS performs two automatic incremental backups, one before and one after a rollback.

These backups do not occur if all the files you have selected to restore are identical to the files already on the server. In such a case, the restoration does not in fact change any files on the server, and the rollback option is not available (there are no changes to rollback).

Entries for Directories in the Backup History

As discussed in the "Special Considerations for Directory and Wildcard Targets" in the *Opsware® SAS Configuration Guide*, if you are tracking a directory and the contents of the directory change, the action that you selected is triggered for the directory object itself and for the files that changed. The only exception occurs when a file is deleted from the directory. In that case, the action is triggered for the directory object alone.

If you selected the backup action for the directory target, when the directory contents change, the directory object is backed up as well as the changed files. When a file is deleted, however, only the directory object is backed up.

If you restore files contained in a directory without selecting the entry for the directory object, the directory is re-created on the server if it does not already exist. If you do select the entry for the directory object, however, you can ensure that the directory object is restored with the same properties (such as permissions and time stamp) it had at the selected backup event.

File Not Found Entries

If you are monitoring specific files (as opposed to files monitored in tracked directories and files monitored as the result of wildcard targets), the backup history can contain entries noting "file not found" if the file does not exist on the server, or if the file is later deleted on the server. Similarly, if you are monitoring specific directories (opposed to monitoring directories through wildcard targets), the backup directory can contain "file not found" entries if the directory does not exist on the server, or if the directory is later deleted.

If you select and restore "file not found" entry and the file exists on the server, the file is deleted (it is reverted to the state of the backup event that you selected.)

Exercise caution in restoring entries for deleted directories. If the directory exists on your server and the directory contains files, both the directory and its contents will be deleted.

Restoring Backups

Perform the following steps to restore backed-up configuration files:

1. From the navigation panel in the Opsware Command Center, click Servers.
2. Click **Server Search** to search for the desired servers. (Alternatively, you can click **Manage Servers** and then select the server from the server list.)
3. Under the Configuration Tracking drop-down menu, select View Backup History.
4. Find the date and server that has the backup points that contain the files that you want to restore. See "Backup History" on page 658 in this chapter for information about finding backup points, viewing backup point details, and viewing lists of backed up files.
5. Select the check box for the backup points that have files that you want to restore.
6. Click **Restore**.

   The files or collection of files from the backup points you selected are placed in the Restore Queue. If you do not want to perform a restore now, you can click the Return to Select a Task link to leave this page. While your session is active, the files remain in the Restore Queue. You can also continue to add more files to the Restore Queue. To return later to the Restore Queue, click the View Restore Queue and Perform Restore from the Track Configurations: Select a Task page.
7. If you are ready to restore files, first review the files in the Restore Queue. If you want to restore all the files in the queue, you can select the check box at the top of the list. If you do not want to restore all of the files in the queue, select the checkboxes for the files that you want to restore.
8. Click **Restore** (or **Restore All** if you want to restore all of the files in the queue).
9. After the restoration is complete, you can click **View Details** for more information about the restored files.

Rolling Back Restored Files

You can rollback any files that you selected to restore. By rolling back the files, you revert the file to the version that was on the server before you performed the restoration. (In other words, you undo the restoration.)

The rollback option is not available if you restored any files that are no longer being monitored by the Configuration Tracking feature. You can choose to restore all changed files or select the files individually.

Perform the following steps to rollback restored files:

1. After you restore the backups, click **Rollback**.

   The Track Configurations: Restore Queue page appears, which contains the list of files that you need to revert in order to rollback the tracked configuration files to their previous state.

   This list might not contain all the files that you selected for your original restore. If any of the files that you selected during your original restore are identical to the files already on the server, they do not need to be rolled back. Only the files that changed display.

2. If you want to rollback all the changed files, click **Restore All**. Otherwise, you can select the individual files that you want to rollback and click **Restore**.

   The Track Configurations: Restore Progress page appears.

Exercise caution in restoring entries for deleted directories. If the directory exists on your server and the directory contains files, both the directory and its contents will be deleted.

Chapter 18: Code Deployment and Rollback

IN THIS CHAPTER

This section discusses the following topics:

- Opsware Code Deployment Process
- Services, Synchronizations, and Sequences

You must have specific permissions to deploy code and content by using the Opsware Command Center. Contact your Opsware administrator to obtain the necessary access rights. For more information, see the Permissions Reference appendix in the *Opsware® SAS Configuration Guide*.

Opsware Code Deployment Process

This section provides information on the code deployment process within Opsware SAS and contains the following topics:

- Deploying Code
- Uploading Code and Content to Staging
- CDR Operations and Directories
- CDR Features
- CDR Permissions
- Accessing CDR

Deploying Code

The Code Deployment & Rollback (CDR) feature in the Opsware Command Center provides tools for deploying new and updated code and content to your operational environment.

The following figure shows the architecture and process for updating a typical server hosted in an Opsware managed environment.

Figure 18-1: Typical Code and Content Update in the Opsware Managed Environment

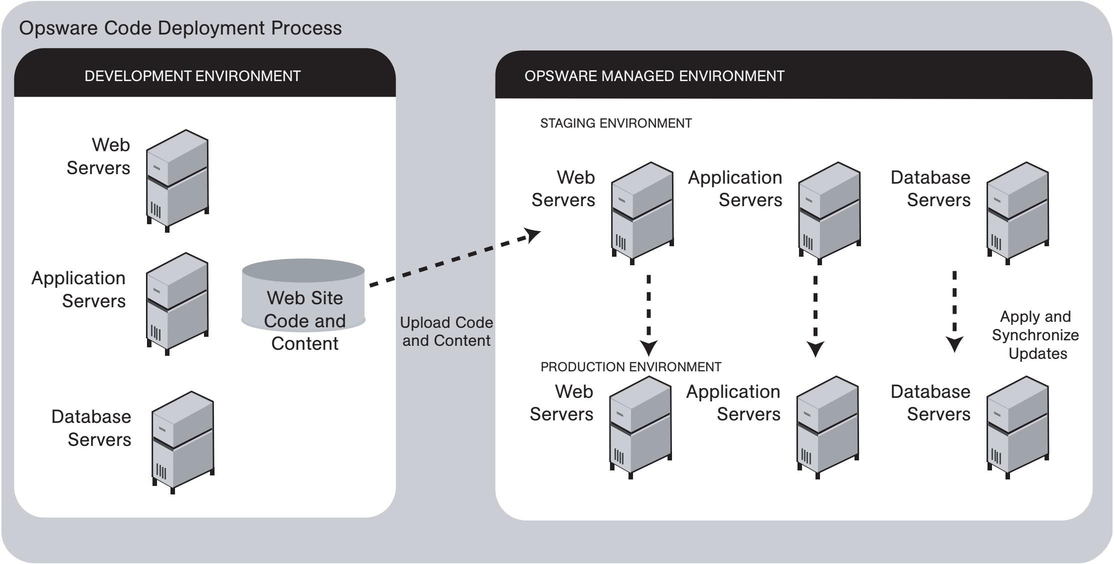

The deployment process involves performing the following high-level tasks:

1. Determining your application code and content deployment requirements and defining the CDR services, synchronizations, and sequences that you need to support them.

   - Services are defined for each different type of web server or application server applications (for example, WebLogic Server) that is installed on the staging and production hosts in your environment.

   - Synchronizations are defined for each service so that you can update files between the source location and one or more destination production hosts that are running the same service.

   - Sequences are optional but can simplify deployment by grouping a collection of service operations and synchronizations that can be performed as a single task.

2. Uploading new or updated code and content to your Opsware staging environment.

3. After performing any necessary testing, cutting over to the changed code and content on the staging environment.

4 As necessary, performing CDR service operations, such as backing up code and content from your live site.

5 Performing CDR operations available to synchronize the updated code and content to your production hosts in the Opsware managed environment.

6 To simplify subsequent deployments of new code and content, defining sequences that specify a series of service operations and synchronizations you want to perform as a single action.

The code and content deployment process that you follow might be different depending on the architecture of your operational environment and your deployment requirements.

Uploading Code and Content to Staging

Before you use CDR to push code and content, you must upload new or updated files to your Opsware staging environment. You can use content management tools, such as OpenDeploy, scp, or rsync over SSH, to do that.

The following figure shows an example of a typical development environment and how your uploaded code and content move to the staging environment.

Figure 18-2: How Code and Content Move to the Staging Environment

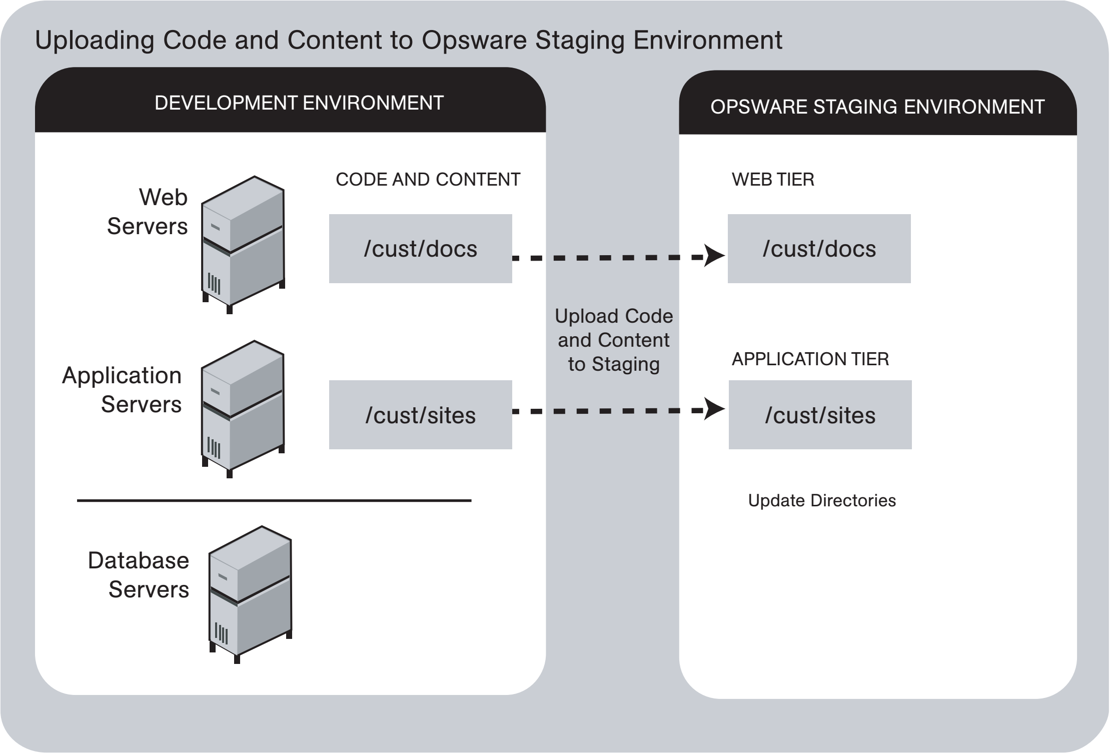

After you upload the files and test your changes, you can synchronize updates to the production hosts running your managed environment. You can run specific synchronizations and perform other service deployment operations by selecting CDR menu options available from the Opsware Command Center navigation panel.

CDR Operations and Directories

After you upload updated code and content to your Opsware-managed staging environment, you can use the CDR operations to cutover to new code and content, perform host synchronizations, and perform other service operations.

CDR uses the following directories to synchronize and cutover code and content for specified hosts:

- **Live directory**: The directory that stores the actual code and content required to run a live site.

- **Update directory**: The directory written to by CDR synchronizations. Stores only the files that changed between the source host Live directory and the Live directories of the destination hosts.
- **Site Previous directory**: This directory holds all the changes necessary to revert the Live directory back to the state it was in before the last cutover. Like the Update directory, the Site Previous directory only stores the files that changed between the current Live directory contents and its previous state.
- **Site Backup directory**: This directory stores a complete backup of the site. The directory is populated when the user issues a Backup service operation.

When you cutover to new code and content, CDR determines the differences between the new code and content in the current Update directory and the Live directory for your site. The files that are different are synchronized to the Live directory. When you synchronize source and destination hosts, CDR moves modified files from the Live directory on a source host to a directory on a destination host.

You cannot use CDR to automate database pushes. However, you can configure CDR so that you can synchronize modified database script files on different hosts.

CDR Features

CDR offers the following features:

- Provides a single tool for deploying code (such as ASP, JSP, and JAR files) and site content (such as HTML, JPEG, GIF, and PDF files). Using a single tool is helpful when the code and content for your site are intermingled.
- Provides direct control over code and content pushes by making it possible to decide what information to update and determine when and how to perform updates.
- Provides flexibility to accommodate frequent updates to staging and production hosts by enabling more frequent pushes in a shorter period of time.
- Allows verification of file changes between staging and production host directories by creating a manifest of updated files. You can verify changes before cutting over to new code and content.
- Provides administrative service operations, including starting and stopping services, and backing up, restoring, and rolling back code and content to return your site to the previous version.

- Lets you push incremental updates to your site so that only files that have changed are pushed to specified locations on staging or production hosts.
- CDR uses the same authentication and navigation that you use in accessing other information and performing other site operations from the Opsware Command Center.

CDR Permissions

As with all other features in Opsware SAS, the links that you see on the Opsware Command Center Home page and the links that you see in the navigation panel are based on the permissions that you have in combination with the customer you are associated with.

If you do not have permissions for CDR, you cannot see the Code Deployment links on the navigation panel, the link called Deploy Code in the Tasks panel of the Opsware Command Center home page appears in italics, and it is not an active link.

If you have CDR permissions to no more than one customer, when you expand the Code Deployment section in the navigation panel, you can see a link called Set Customer. Click that link to view the links to the specific Code Deployment functions that you have permissions for in combination with that single customer.

If you have CDR permissions to more than one customer, you can see a link called Select Customer. Click that link to display a page that shows the customers you are associated with. Select the customer you want to work with. The CDR Home Page appears, with links to the specific Code Deployment functions that you have permissions for. These links are the same functions that you can find in the navigation panel under Code Deployment.

The navigation instructions and screen captures in this chapter show what a user with permissions to all code deployment functions and access to only one customer can see. Consequently, because your permissions and customers might be different, the available menu selections and features that you see might likewise differ.

Accessing CDR

Perform the following steps to access CDR:

1. If necessary, click the Code Deployment link in the navigation panel to expand the list of CDR options.
2. Click the CDR Home link. The CDR Home Page for [*customer name*] appears, as the following figure shows.

Figure 18-3: Code Deployment Home Page

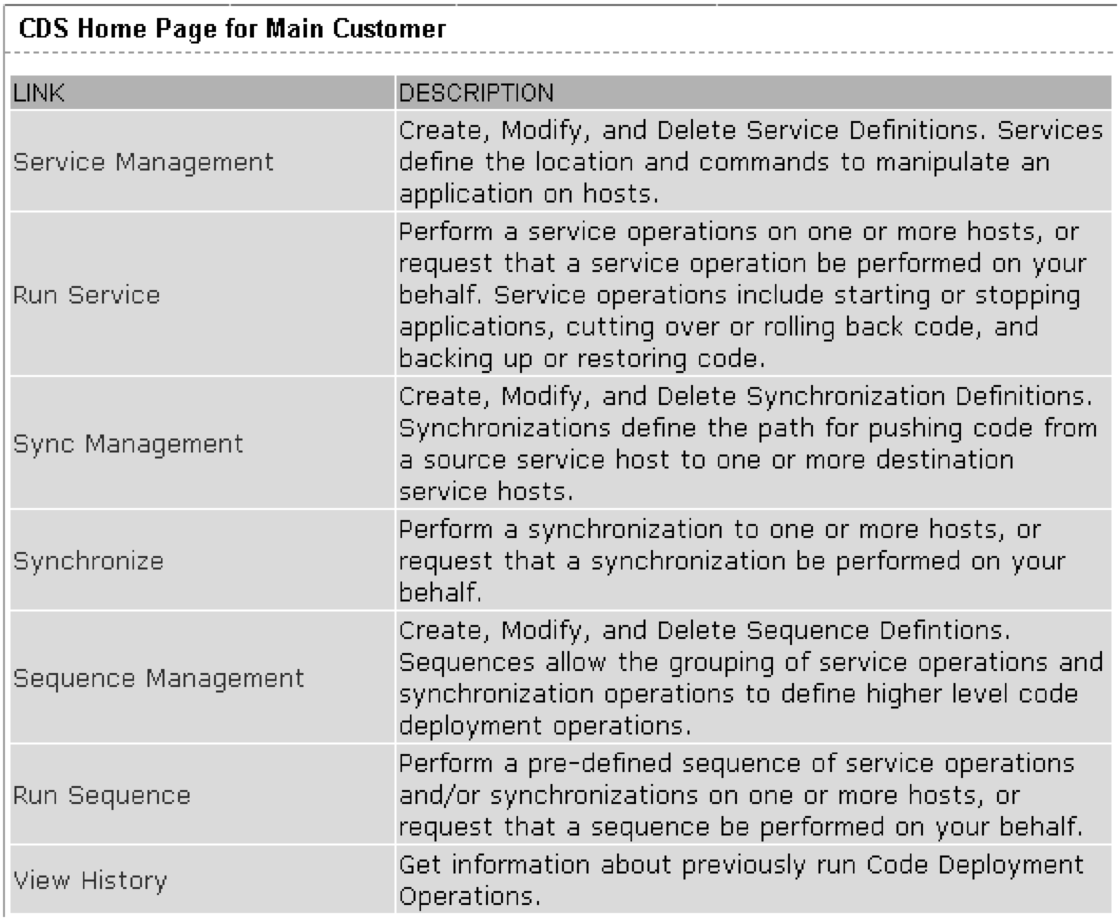

CDS Home Page for Main Customer

LINK	DESCRIPTION
Service Management	Create, Modify, and Delete Service Definitions. Services define the location and commands to manipulate an application on hosts.
Run Service	Perform a service operations on one or more hosts, or request that a service operation be performed on your behalf. Service operations include starting or stopping applications, cutting over or rolling back code, and backing up or restoring code.
Sync Management	Create, Modify, and Delete Synchronization Definitions. Synchronizations define the path for pushing code from a source service host to one or more destination service hosts.
Synchronize	Perform a synchronization to one or more hosts, or request that a synchronization be performed on your behalf.
Sequence Management	Create, Modify, and Delete Sequence Defintions. Sequences allow the grouping of service operations and synchronization operations to define higher level code deployment operations.
Run Sequence	Perform a pre-defined sequence of service operations and/or synchronizations on one or more hosts, or request that a sequence be performed on your behalf.
View History	Get information about previously run Code Deployment Operations.

Depending on your access permissions, the following CDR options appear:

- **Service Management**: Create, modify, or delete service definitions that define the location and commands to manipulate an application on hosts associated with each application instance running in your operational environment.

- **Run Service**: Perform a service operation or request that one be performed.
- **Sync Management**: Create, modify, or delete synchronization definitions associated with code pushes.
- **Synchronize**: Perform a synchronization or request that one be performed.
- **Sequence Management**: Create, modify, or delete sequences of operations.
- **Run Sequences**: Perform a selected sequence or request that one be performed.
- **View History**: View information stored in an operations log to determine the status of particular deployment operations, and whether they completed successfully.

3 Choose the CDR operations that you want to perform, selecting options from the navigation panel or from the CDR home page.

Services, Synchronizations, and Sequences

This section provides information on performing services, synchronizations, and sequences within Opsware SAS and contains the following topics:

- Overview of Performing Services, Synchronizations, and Sequences
- Synchronization of Site Code and Content
- Performing Synchronizations
- CDR Cutover Operation
- CDR Service Operations
- Starting and Stopping Host Services
- Backing Up Code and Content
- Restoring Code and Content from a Previous Version
- Rolling Back Code and Content to the Previous Version
- Accessing Service Operations in CDR
- Performing Service Operations by Service Name
- Performing Service Operations by Host Name
- Performing Sequences
- Processing Code Deployment Requests from Users

- Status View of Previous Operations

Overview of Performing Services, Synchronizations, and Sequences

After you upload updated code and content to your Opsware SAS staging environment, you use CDR to cutover to new code and content, perform host synchronizations, and perform other service operations.

When you cutover to new code and content, CDR determines the differences between the new code and content in the current Update directory and the Live directory. The files that are different are synchronized to the Live directory. When you synchronize source and destination hosts, CDR moves modified files from the Live directory on a source host to a directory on a destination host.

The code and content footprint that CDR maintains on a host is larger than the storage space for the code and content files of the live site. The amount of storage used beyond the actual size of your site increases with the number of files that you modify and back up.

Synchronization of Site Code and Content

CDR uses the following directories to synchronize and cutover code and content for specified hosts:

- **Live directory**: The directory that stores the actual code and content required to run a live site.
- **Update directory**: The directory written to by CDR synchronizations. Stores only the files that changed between the source host Live directory and the Live directories of the destination hosts.
- **Site Previous directory**: This directory holds all the changes necessary to revert the Live directory back to the state it was in before the last cutover. Like the Update directory, the Previous directory only stores the files that changed between the current Live directory contents and its previous state.
- **Site Backup directory**: This directory stores a complete backup of the site. The directory is populated when the user issues a Backup service operation.

CDR can synchronize updates from the Live directory of a source host to either the Update or Live directories on destination hosts. Decide whether you want to synchronize changed files on the source host to Update or Live directories on the destination hosts:

- **Synchronize to Update directories**: CDR determines files that have changed by comparing files in the destination host Live directory and the source host Live directory. Updated files are stored in the Update directories of destination hosts.
- **Synchronize to Live directories**: CDR updates changed files to the Live directory on destination hosts bypassing the Update directory.

⚠ If you choose to synchronize directly to Live directories, the Rollback operation does not function properly. Therefore, choose this option only for synchronizations that are not likely to impact site stability.

When you choose to synchronize to Live directories, back up the Live directories first and run the Restore operation to return your site to the previous version. See "Rolling Back Code and Content to the Previous Version" on page 687 in this chapter for more information.

Figure 18-4 shows an example of synchronization between hosts and how you synchronize updates for each service.

Figure 18-4: Service Synchronization of Hosts Using CDR

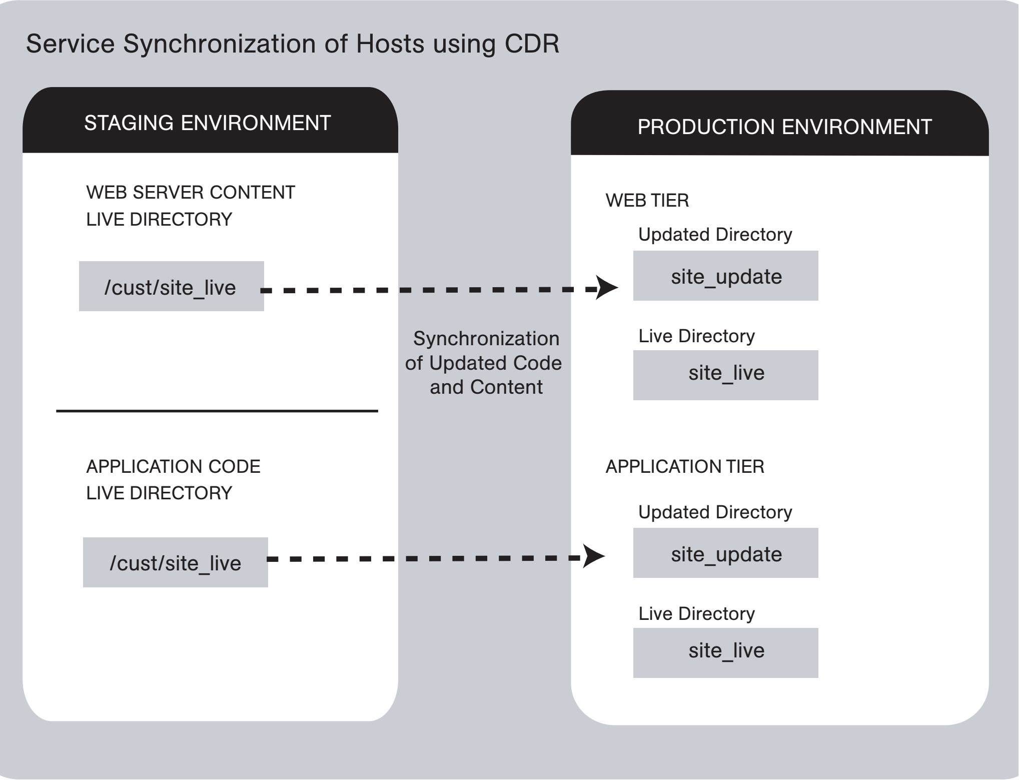

Contact your Opsware administrator for a description of services and synchronizations set up for your site and the specific operations that you need to run when updating code or content for a particular host service.

Performing Synchronizations

Perform the following steps to synchronize updated code and content from one source host to one or more destination hosts:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.

2. Click the Synchronize link.

   A page appears that displays the synchronizations that you can perform.

3. Select the synchronization that you want to perform by clicking the link. The CDR Synchronize for [customer name] page appears, as Figure 18-5 shows.

Figure 18-5: The CDR Synchronize Page

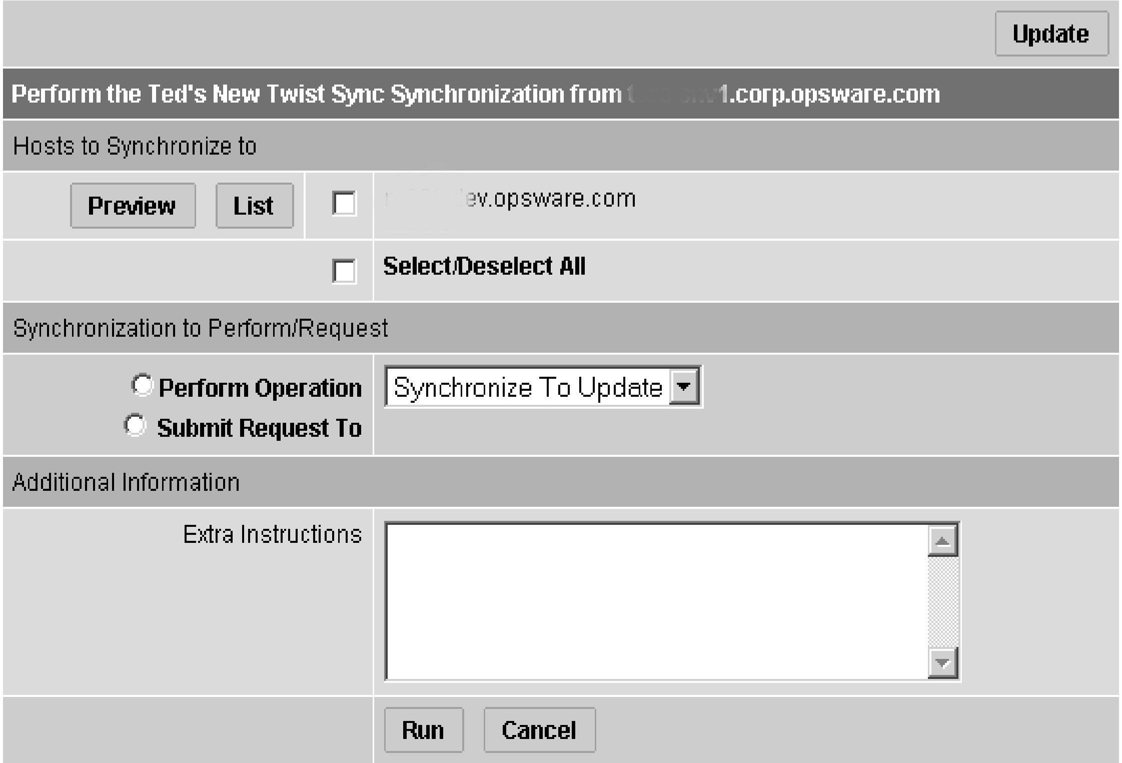

Perform the Ted's New Twist Sync Synchronization from [redacted]1.corp.opsware.com

Hosts to Synchronize to

☐ [redacted]ev.opsware.com

☐ Select/Deselect All

Synchronization to Perform/Request

○ Perform Operation
○ Submit Request To

Synchronize To Update ▾

Additional Information

Extra Instructions

4. Select one or more of the displayed host names on which you want to perform synchronization. The hosts that you select are the destination hosts.

Choose the Select/Deselect All option to select all or clear all host names.

5. (Optional) To view a list of files that will be created, modified, and deleted on the destination hosts, click **Preview**.

   Or

   (Optional) To view a list of all the files on the destination hosts, click **List**.

6. Select the Perform Operation option to directly perform a selected synchronization.

Or

Select the Submit Request To option to send a request to users specified to receive email notification for service and synchronization requests. When you submit a request, specify any additional instructions that you want to include for the requested synchronization. For example, these might be instructions such as the time that you want the synchronization performed, verification, or other related services to perform.

The Perform Operation option is only visible when you are a member of an Opsware Command Center user group allowed to directly perform a synchronization.

7. Choose the type of synchronization that you want performed from the drop-down list:

   - Synchronize To Update
   - Synchronize To Live

   See "Synchronization of Site Code and Content" on page 679 in this chapter for information about these options.

8. To initiate the synchronization or send the request, click **Run**.

CDR Cutover Operation

Perform the cutover operation to make the Update directory and the current live site identical.

When you cutover, CDR performs the following actions:

- Updates the Site Previous directory with files from the Live directory. CDR saves modified files and files-to-be-deleted to the Site Previous directory. The Site Previous directory contains the files necessary to restore the live site to the previous version.

- Determines the differences between the Update directory and the current Live directory. The files that are different are synchronized from the Update directory to the Live directory. See Figure 18-6.

Figure 18-6: Directory and File Updates for Cutover Operations Using CDR

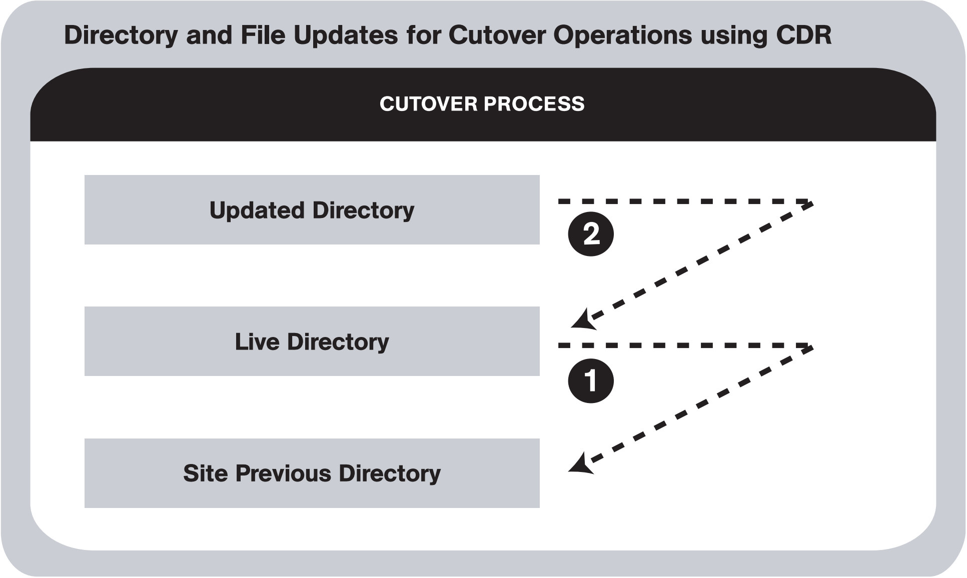

CDR determines file differences between source and destination directories based on file size, modification date and time, ownership, group attributes, and permissions attributes.

By using the cutover process, CDR ensures your ability to rollback to the previous version of your code and content if you experience a problem.

Your Opsware administrator can configure a CDR service to run scripts before and after cutting over to updated code and content. For example, before and after cutting over, you might distribute content on geographically disperse servers.

See "Synchronization of Site Code and Content" on page 679 in this chapter for information about a description of these directories. See "CDR Service Operations" on page 685 in this chapter for information about how to rollback code and content to the previous version.

CDR Service Operations

In addition to the Cutover operation, CDR provides a number of service operations:

- Starting and stopping host services
- Backing up code and content
- Restoring code and content from a previous version
- Rolling back code and content to the previous version

Performing these operations might be required depending on the type of code and content changes made or the host services that are affected.

The operations that you need to perform are specific to the service (for example, Web server or application server instance) for which you are updating code or content and the particular host.

Starting and Stopping Host Services

Stopping and starting services might be required depending on the type of code and content changes made or the host services that are affected. Discuss your requirements with your Opsware administrator.

Typically, you only stop and start host services for the hosts in your staging environment. Select members of your staff, or other individuals in your operations center, can stop and start host services for the hosts in your production environment.

Start: Launch a defined service; for example, starting a web or application server instance that is running on a specific host.

Stop: Shut down a defined service; for example, shutting down a web or application server instance before cutting over to new or changed code and content on a specific host.

Backing Up Code and Content

When you use CDR to back up your site, CDR saves the entire contents of the current Live directory for a specific service in the Backup directory. CDR saves the backup copy to the local disk for the host on which you ran the Backup operation. See Figure 18-7.

You can use CDR to keep only one backup copy at a time for a service.

Figure 18-7: Update Directory to Site Backup Directory for Backup Using CDR

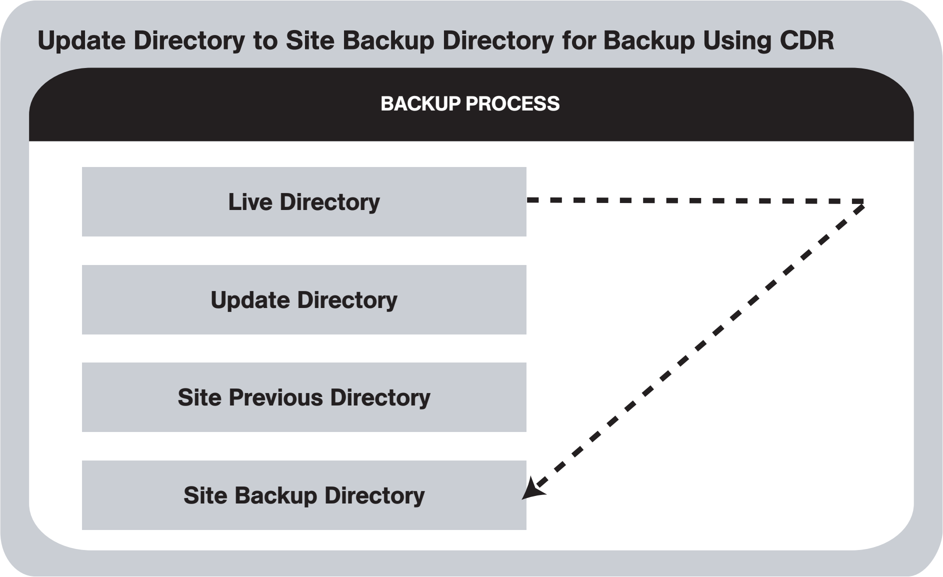

When you run the Restore operation, CDR replaces the Live directory contents with files stored in the Backup directory.

When you reach a high level of site stability, backing up your site is recommended, especially if you plan to make changes to site code and content.

Restoring Code and Content from a Previous Version

The Restore operation restores the previous Live directory by copying the contents of the Backup directory to the Live directory. Restoring code and content from the Backup directory does not change files that are stored in the Update directory. See Figure 18-8.

Figure 18-8: Site Backup Directory to Update Directory for Restoring Using CDR

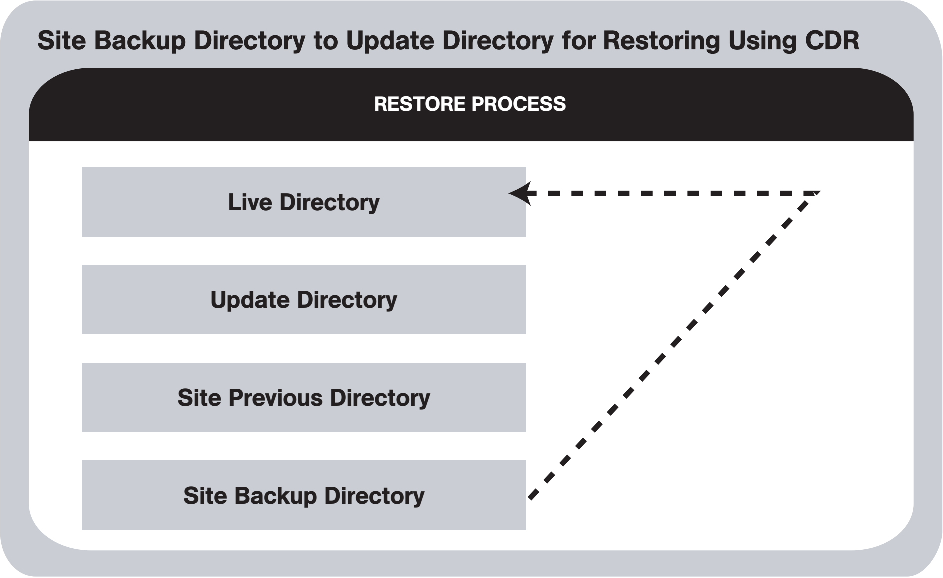

Before you restore code and content, you must have backed up the contents in the Live directory to the Backup directory by performing a Backup operation.

Rolling Back Code and Content to the Previous Version

If you experience a problem after cutting over updated code and content to your production site, you can rollback to the previous version.

Rolling back returns the site to the state it was in prior to the last cutover that you performed. See Figure 18-9.

Figure 18-9: Directory and File Updates for Rollback Operations Using CDR

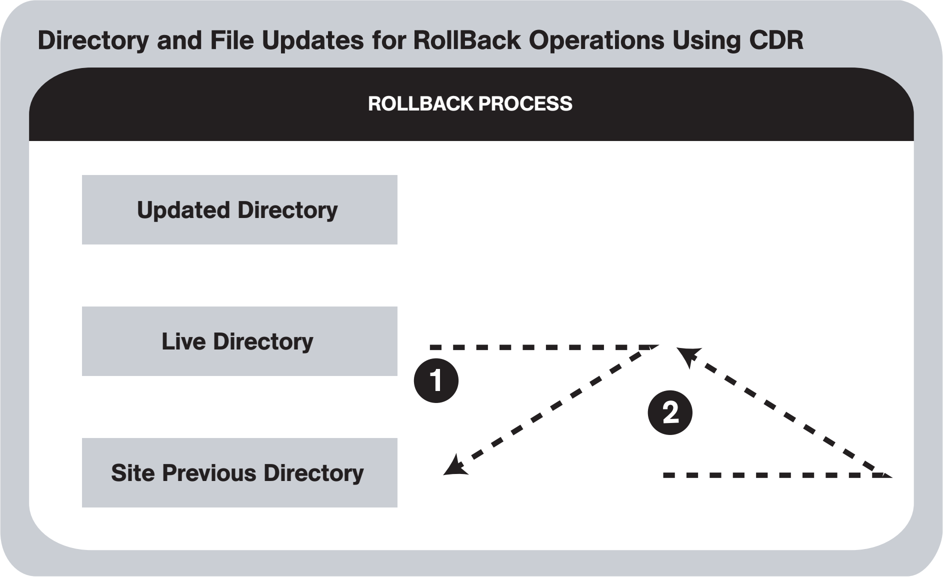

During cutover, CDR updates the Site Previous directory with files from the Live directory. CDR saves modified files and files-to-be-deleted to the Site Previous directory. The Site Previous directory contains the files necessary to restore the live site to the previous version.

During rollback, CDR restores the set of different files (modified files and files that were deleted during cutover) to the Live directory.

If you upload files directly to the Live directory or choose to synchronize directly to Live directories, the rollback operation does not function properly. Under these conditions, back up your Live directory and run the Restore operation to return your site to the previous version.

Accessing Service Operations in CDR

The Service Management option provides a number of service or administrative operations, including starting and stopping host services, backing up, restoring, or rolling back code and content, and cutting over to new site code and content.

To access CDR, your Opsware administrator must add you as a member of a user group authorized to use CDR.

You have the option of initiating a service either by selecting a service name or by selecting host names. You must select both a service to perform and the hosts on which to perform the service. See Figure 18-10.

Figure 18-10: Run Service Page

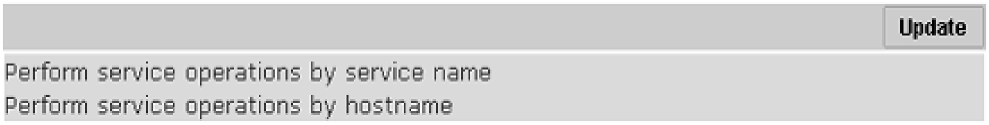

Perform service operations by service name
Perform service operations by hostname

- Initiate a service by selecting the “Perform service operations by service name” option first when you want to perform a service on multiple hosts at the same time. Selecting the service name first shows you all the hosts for which that service is defined.
- Initiate a service by selecting the “Perform service operations by host name” option first when you want to perform a service on a single host or on a specific host where you know the host name. Selecting the host name first shows you all the services that are defined for that host.

Performing Service Operations by Service Name

Perform the following steps to perform service operations by service name:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Select the Service Management option.
3. From the Service Management page, select the “Perform service operations by service name” link.

4. Select a service from a list of services defined for your site. A page that prompts you to select the hosts and the operation that you want to perform appears, as Figure 18-11 shows.

Figure 18-11: Perform Service Operations by Service Name Page

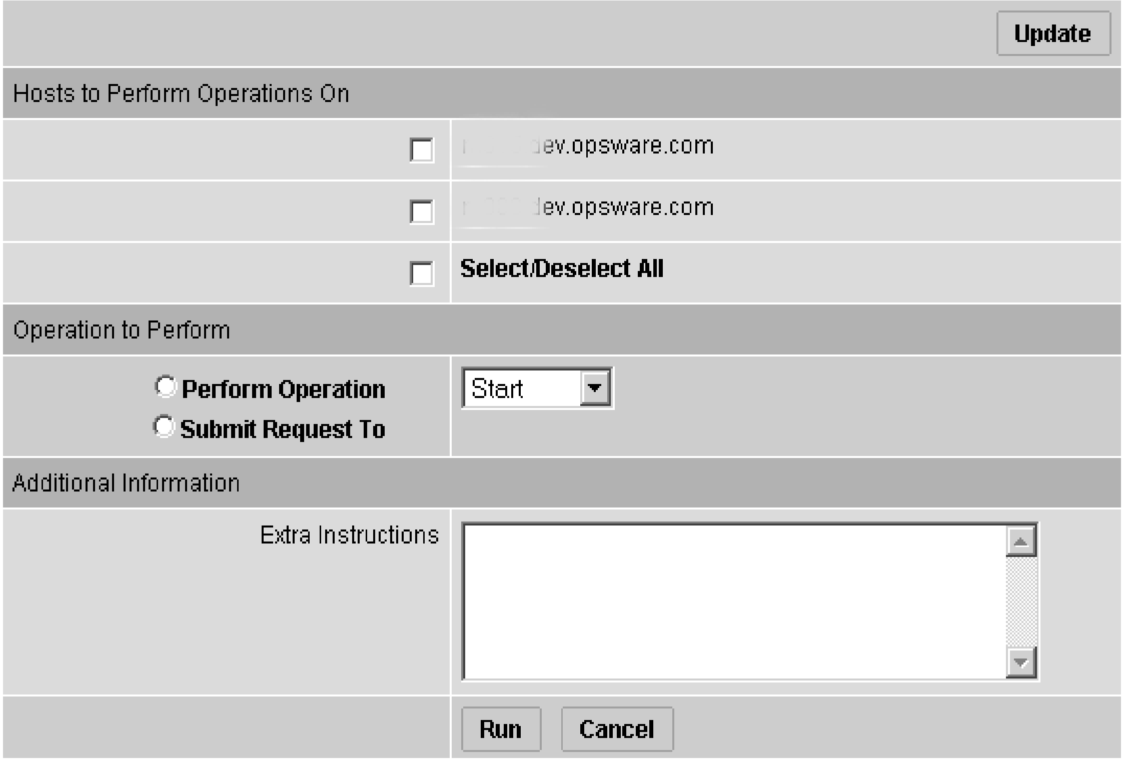

Update

Hosts to Perform Operations On

□ [redacted]ev.opsware.com

□ [redacted]ev.opsware.com

□ **Select/Deselect All**

Operation to Perform

○ **Perform Operation** Start ▾

○ **Submit Request To**

Additional Information

Extra Instructions

Run Cancel

5. Select one or more of the displayed host names. You can choose the Select/Deselect All option to select all or clear all host names for the operation that you want to perform.

6. Select the Perform Operation option to directly perform a selected operation.

The Perform Operation option is visible only when you are a member of the CDR user group that allows users to directly perform a service operation.

Or

Select the Submit Request To option to send a request for authorized individuals to perform the operation for you. When you submit a request, specify any additional instructions that might be required to perform the requested operation.

7 Choose the type of operation that you want performed from the drop-down list:

- Start
- Stop
- Cutover
- Rollback
- Backup
- Restore

See “CDR Service Operations” on page 685 in this chapter for information about these operations.

8 To initiate the operation or send the request, click **Run**.

Performing Service Operations by Host Name

Perform the following steps to perform service operations by host name:

1 Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.

2 Select the Service Management option.

3 From the Service Management page, select the “Perform service operations by hostname” link.

4 Select a host name from the list of staging and production hosts available for your site. A page that prompts you to select the service and the operation that you want to perform appears, as Figure 18-12 shows.

Figure 18-12: Perform Service Operations by Host Name Page

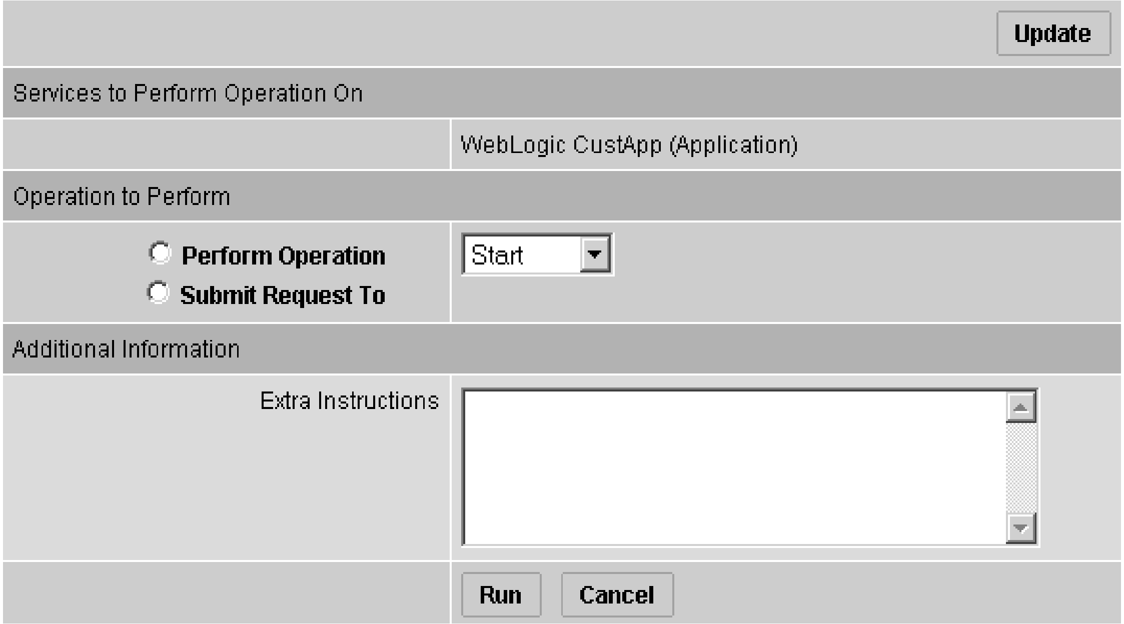

Update

Services to Perform Operation On

WebLogic CustApp (Application)

Operation to Perform

○ Perform Operation
○ Submit Request To

Start

Additional Information

Extra Instructions

Run Cancel

5 Select the service for which you want to perform an operation.

6 Select the Perform Operation option to directly perform a selected operation or select the Submit Request To option to send a request to have authorized individuals perform the operation for you. When you submit a request, specify any additional instructions that might be required to perform the requested operation.

The Perform Operation option is visible only when you are a member of the user group that allows users to directly perform a service operation.

7 Choose the type of operation that you want performed from the drop-down list:

- Start
- Stop
- Cutover
- Rollback
- Backup
- Restore

See "CDR Service Operations" on page 685 in this chapter for information about these operations.

8 To initiate the operation or send the request, click **Run**.

Performing Sequences

CDR also allows you to perform service operations and synchronizations that have been set up as a sequence of operations.

Perform the following steps to perform CDR sequences set up for your site:

1 Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.

2 Click the Run Sequence link.

The CDR Run Sequence for [customer name] page appears that shows the sequences that you can run. See Figure 18-13.

Figure 18-13: CDR Run Sequence Page

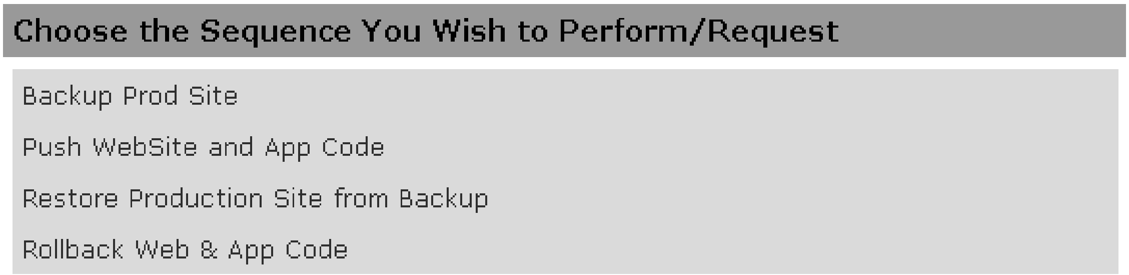

Choose the Sequence You Wish to Perform/Request

Backup Prod Site

Push WebSite and App Code

Restore Production Site from Backup

Rollback Web & App Code

3 Select the sequence that you want to perform by clicking the link. The Run Sequence page appears, as Figure 18-14 shows.

Figure 18-14: Run Sequence Page That Shows Details of a Specific Sequence

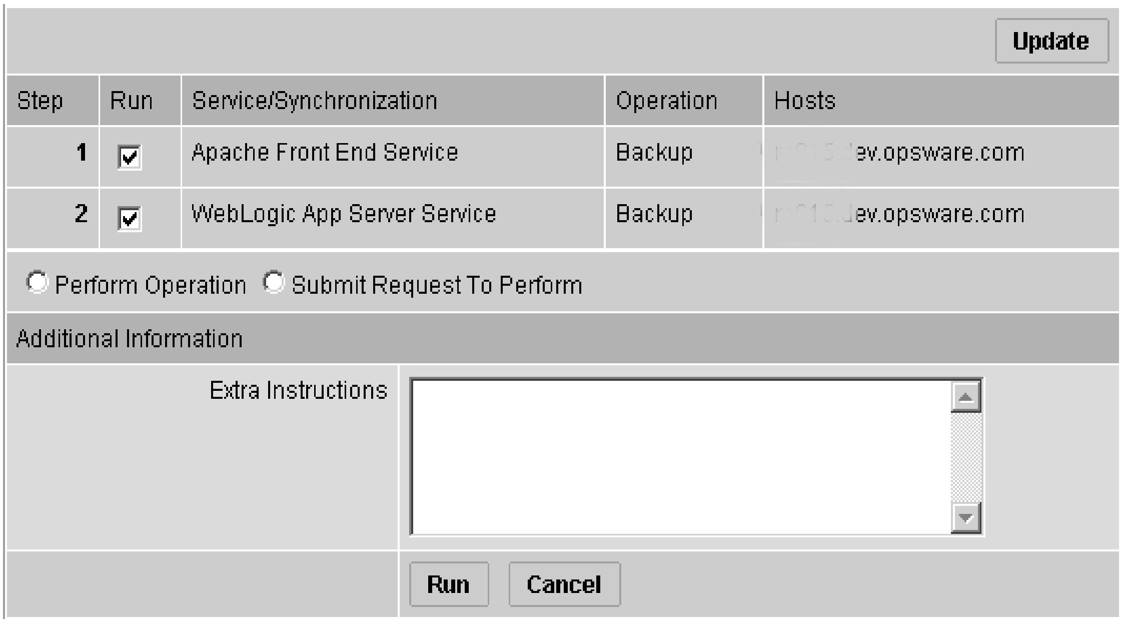

Step	Run	Service/Synchronization	Operation	Hosts
1	☑	Apache Front End Service	Backup	[redacted]ev.opsware.com
2	☑	WebLogic App Server Service	Backup	[redacted]ev.opsware.com

○ Perform Operation ○ Submit Request To Perform

Additional Information

Extra Instructions

4 Select the Perform Operation radio button to directly perform the selected sequence.

Or

Select the Submit Request To Perform radio button to send a request to the users specified to receive email notification for service, synchronization, and sequence requests. When you submit a request, specify any additional instructions that you want to include for the requested sequence. For example, you might want to include instructions such as the time that you want the sequence to run, verification, or other related services to perform.

The Perform Operation option is visible only when you are a member of a user group allowed to directly perform a sequence.

5 To initiate the sequence or send the request, click **Run**.

Processing Code Deployment Requests from Users

When CDR users request that a service operation or synchronization be performed on their behalf, an email notification is sent to the individuals assigned to perform the requested service operation or synchronization.

The following message is a typical example of an email notification request.

```
From:    CDR-tool@opsware.com
To:      opscenter@opsware.com
Date:    Tue, 10 Jul 2001 11:25:13 -0700
Subject: Request To Perform Start operation for opsware.com

Please perform the following request

Requestor: jhancock/opsware.com
Request Time: Jul 10, 2001 11:25:13 AM PDT
Requested Service: Demo Apache
Requested action: Start
Perform on the following hosts:
host1.opsware.com
host2.opsware.com

Extra Instructions:
```

In this case, the email message specifies a request from a user, jhancock, to perform a Start operation on two hosts running the Demo Apache service. The subject line provides a summary of the request. In addition, the email indicates the time that the request was sent.

Perform the following steps to process the request:

1. Identify any special instructions that might need to be carried out for the specific request.
2. Log into the Opsware Command Center, choose the CDR option, and then select the particular option within CDR to perform the requested operation.
3. When you successfully complete the CDR request, you might want to notify the individual user making the request, and all other involved parties, that the requested operation was completed.

If you encounter problems in completing a request and cannot resolve them, contact your Opsware administrator for any specific remedies and follow normal escalation procedures defined for your operational environment.

See the *Opsware® SAS Configuration Guide* for more information about troubleshooting tips.

Status View of Previous Operations

CDR maintains a log of operations (service operations, synchronizations and sequences) that were executed. You can view this information to determine the status of particular deployment operations, and whether they have completed successfully. You can also use the My Jobs task area of the Opsware Command Center Home page, or the My Jobs link to view this information.

Accessing the Log

Perform the following steps to access the log:

1. Click the Code Deployment link in the navigation panel to expand the list of CDR options and select a customer, if necessary.
2. Select the View History option.

CDR displays a page providing a list, most recent to oldest, of operations that are either in progress or those that have been completed. This information displays only for the past 60 days. Each page is limited to a list of 10 operations. A Next link displays at the end of the page if there are more than 10 operations to view. Click the Next link to view subsequent operations. A Previous link is available to return to previous pages.

You need to refresh the page to view the status of any operations initiated after you selected the View History link.

A similar page displays after you select the View History link. See Figure 18-15.

Figure 18-15: CDR View History Page

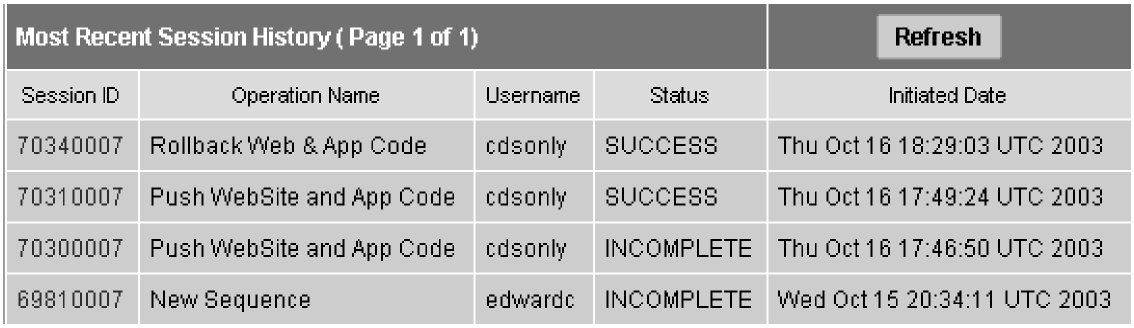

Most Recent Session History (Page 1 of 1)

Session ID	Operation Name	Username	Status	Initiated Date
70340007	Rollback Web & App Code	cdsonly	SUCCESS	Thu Oct 16 18:29:03 UTC 2003
70310007	Push WebSite and App Code	cdsonly	SUCCESS	Thu Oct 16 17:49:24 UTC 2003
70300007	Push WebSite and App Code	cdsonly	INCOMPLETE	Thu Oct 16 17:46:50 UTC 2003
69810007	New Sequence	edwardc	INCOMPLETE	Wed Oct 15 20:34:11 UTC 2003

The individual headings of column information included in the table are:

- **Session ID**: A session is created each time a CDR operation is performed. Click the Session ID to view detailed results of the operation.
- **Operation Name**: The name of the service, synchronization or sequence as specified when they were first defined.
- **User Name**: The user ID of the user that initiated the operation.
- **Status**: This describes the state of the operation at the end of the sequence. The status message varies depending on the type of operation. Single step operations always result in Success/Failure messages while multiple step operations (Sequence operations) could result in Complete with Error, Incomplete, or Success messages. Table 18-4 describes the possible status messages.

Table 18-4: CDR History Status Messages

STATUS TYPE	DESCRIPTION
Abort	This message displays only if a CDR specific script that is executed in order to complete service, synchronization or sequences fails. Such issues should be escalated to your Opsware Support Representative.
Active	This message displays if an operation is still in progress.
Complete with error	The sequence completed successfully but there were errors reported while the operation was in progress and the user opted to continue rather than cancel the operation.
Failure	The operation (service, sequence or synchronization) did not complete successfully.
Incomplete	The sequence resulted in an error and the user opted to cancel the sequence rather than continue.
Success	The operation (service, synchronization or sequence) was completed successfully and no errors were reported. In the case of a sequence, this message means that all the steps were completed successfully.
Initiated Date	The date on which the operation was initiated.

Appendix A: Reconcile

IN THIS APPENDIX

This section discusses the following topics:

- Overview of Reconcile
- Ways to Perform Reconcile
- Reconcile Operations and Guidelines
- Reconcile on Supported Operating Systems
- Specifying Scripts During Reconcile
- Reconcile Output
- Assigning to and Removing Servers from Nodes
- Reconcile Software Wizard

Overview of Reconcile

This section explains how reconcile and test reconcile work for software provisioning and installing Unix patches, and how different software types are treated during reconcile.

See “Reconciling Patches” on page 478 in Chapter 12 for information about reconciling windows patches.

It also provides information about the relationship of the installation and uninstallation wizards to the reconcile process, and explains the output from reconcile. The wizards are implemented using reconcile, and it is important to understand the reconcile process to understand what happens when a user invokes an installation or uninstallation wizard.

This section also explains how to reconcile selected servers through a specific reconcile wizard rather than through the wizards that are used to install and uninstall patches, software, or templates.

Opsware SAS uses a model-based approach that provides a high degree of change management control and detailed records of all changes made. These features in turn make it possible to apply and enforce policies even within large and heterogeneous environments.

One central mechanism that Opsware SAS uses to implement its model-based approach is called server reconcile. Server reconcile orchestrates the installation and uninstallation of all software, including applications and patches, on servers that Opsware SAS manages. See "Opsware SAS Model and Server Management" on page 70 in Chapter 1 for more information.

In Opsware SAS, the installation and uninstallation of software begins by first making a change to the model of the server contained in the Model Repository. When you use any of the install or uninstall software wizards, for example, you cause Opsware SAS to update the model of a server by adding or removing software nodes from the model. See the *Opsware® SAS Configuration Guide* for more information about software provisioning setup.

During the process of using a wizard, if you decide to install or uninstall software immediately, Opsware SAS begins a reconcile session. You can also schedule the installation or uninstallation for a later time, and Opsware SAS starts the reconcile session at the time that you request. The reconcile session determines what needs to be done to install or uninstall the requested software, tests the results of those actions, and then initiates the tasks necessary to install or uninstall the software, such as downloading packages and initiating system utilities on the servers.

This indirect method of installing and uninstalling software – first changing the model of the software, and then changing the server to match the model – allows you to install software more safely and consistently. See "Software Provisioning" on page 549 in Chapter 14 for information about how to use the install and uninstall software wizards.

Ways to Perform Reconcile

Most of the time, users do not actually encounter the term reconcile when they use the standard Opsware SAS installation and uninstallation wizards. The wizards are designed to automate and simplify the reconcile process. Reconcile occurs as a result of using any of the following wizards:

- Install OS (See "Operating System Provisioning" on page 567 in Chapter 15 for more information.)

- Install/Uninstall Unix Patch (See “Patch Management for Unix” on page 519 in Chapter 13 for more information.)
- Install/Uninstall Windows Patch (See “Patch Management for Windows” on page 453 in Chapter 12 for more information.)
- Install/Uninstall Software (See “Software Provisioning” on page 549 in Chapter 14 for more information.)
- Install Template (See “Software Provisioning” on page 549 in Chapter 14 for more information.)

In addition, you can invoke reconcile directly in order to install or uninstall software. You can do so by attaching or removing servers from nodes and then directly invoking a reconcile session on the servers by running the Reconcile Software Wizard.

Most often, users perform this direct reconcile when changes have been made to nodes, and the users want the servers that are already attached to these nodes to be reconciled to match the modified node.

Reconcile Operations and Guidelines

This section provides information on how reconcile works within Opsware SAS and contains the following topics:

- Reconcile Operations
- Reconcile and Package Metadata
- Installation and Uninstallation Order
- Software Installation Order for Adopted Software
- Patches and Reconcile
- Preview Reconcile
- Types of Reconcile

Reconcile Operations

Reconcile works by comparing what is actually installed on a server to the software that should be installed on the server according to the server’s model. Opsware SAS then determines what operations are required to make the server conform to its model.

To make this determination, Opsware SAS queries the Opsware Agent on the server and examines the server's model to assemble the following data:

- A list of all software (including patches) that is installed on the server. Although Opsware SAS records this information in the Model Repository, Opsware SAS uses the Opsware Agent on the server to compile a list of all installed software. Querying the Opsware Agent is necessary in case any software was installed manually, without using Opsware SAS.

Opsware, Inc. recommends that users install all software through Opsware SAS. Manually installed software, however, can be adopted by Opsware SAS. See "Software Installation Order for Adopted Software" on page 704 in this chapter for more information.

- A list of all the software that should be installed on the server, and the proper installation for the software. Opsware SAS obtains this information from the Model Repository.
- A list of software on the server that was installed through Opsware SAS. Opsware SAS requires this information because it cannot uninstall software that was not installed through Opsware SAS. This list is obtained from the Model Repository.

After this information is obtained, reconcile determines what actions are required to make the software match its model. Occasionally, the installation and uninstallation of software involves consequences that the user might not have anticipated. The installation of some software, for example, might require that other incompatible or outdated software be removed. In other cases, the installation of software might require the installation of other software components that were not explicitly requested, but are required by the packages that the user has selected.

Opsware SAS then calls the server's native utilities to carry out the installation or uninstallation of software. (For example, on a Solaris server, Opsware SAS uses the Solaris utility `pkgadd` to install software. On an AIX server, Opsware SAS uses `installp`. See "Reconcile on Supported Operating Systems" on page 706 in this chapter for more information. If the reconcile operation requires both uninstallation and installation of software, the uninstallation occurs before the installation.

Reconcile and Package Metadata

Reconcile relies on metadata about software that is going to be installed and removed, and this metadata is obtained when the software is first uploaded to the Software Repository (either though the command line interface, the Upload Patch Wizard, or though package management). Metadata includes information such as name and version, and this information varies depending on the package type. (The file name alone is not sufficient to identify the package in Opsware SAS.)

In most cases, this metadata is obtained automatically from the system utilities that run on the Software Repository server. The Software Repository server, however, is installed on a Unix-based machine. In the case of Windows packages the metadata must be entered manually, because the Windows utilities that obtain that type of information do not run on Unix.

Installation and Uninstallation Order

One of the most important determinations that the reconcile process makes is the order in which software should be installed or uninstalled.

During reconcile, Opsware SAS determines the correct installation or uninstallation order based on the following factors:

- If a node contains more than one package, the order in which the packages appear in the node affects the order in which they are installed
- The installation order dependencies between nodes, if any

  When users create nodes and associate the nodes with software, they can express installation order dependencies with other nodes (for example, the software in one node should be installed before the software contained in another node, according to how users define the installation order dependencies).
- The order in which software should be uninstalled

  When Opsware SAS installs software on a server, it records the order in which the software was installed (as determined by reconcile) in the Model Repository. The software is uninstalled in the opposite order in which it was installed. This record of software installation order is maintained as long as Opsware SAS manages the server. The record contains data about the order from all reconcile operations, not just individual reconcile operations.

Software Installation Order for Adopted Software

Opsware SAS does not uninstall software that was not installed through Opsware SAS. Software that was manually installed, however, can be adopted by Opsware SAS in a node. If a user attaches a package that was manually installed on a server, assigns the server to the node, and then performs a reconcile, the software becomes adopted by Opsware SAS. (The software is not actually reinstalled during this operation.)

Opsware SAS does *not* adopt Solaris patches. For example, if you uninstall a Solaris patch that was adopted into Opsware SAS, the patch will not be uninstalled.

When software is adopted by Opsware SAS, it is uninstalled according to when it was adopted, not when it was originally installed. The following example illustrates the order in which adopted software is uninstalled:

1. A server has a package that is installed on it before the server comes under management by Opsware SAS.
2. A user installs three packages on this server using Opsware SAS.
3. After these installations, a user decides to adopt the package by creating a node for the package, assigning that node to the server, and reconciling the server through the Reconcile Software Wizard. (The software is not reinstalled when the server is reconciled; instead, the package is simply adopted.)
4. If these packages are later uninstalled, the package that was adopted is uninstalled in the reverse order according to when it was adopted. The adopted package will, therefore, be uninstalled after the three packages, even though the package was in fact on the server before the three packages were installed.

Patches and Reconcile

All software in Opsware SAS is associated with nodes. Ordinarily, these nodes are created explicitly by users during software provisioning setup. Patches, however, are treated differently, in order to expedite the process of patch management. When a patch is first uploaded it is not immediately associated with a node. The first time that a patch is applied to a server, however, the node is created behind the scenes. This node is not part of the ordinary Software Tree and does not display in the tree.

If a user wants to add an installation order dependency to a patch, however, the user does in fact create a node for the patch. The node is used to express the installation order dependency.

See “Setting Patch Installation Order Dependencies for Unix” on page 533 in Chapter 13 for information about how to set installation order dependencies for Unix patches. See “Patch Dependencies and Supersedence” on page 467 in Chapter 12 for information about Windows patches.

Preview Reconcile

Before any changes are committed to a server, you have the option to performs a preview reconcile. The preview reconcile allows users to see exactly what happens to the server as a result of the software that they requested to be installed or uninstalled. (This information displays individually for each server that is selected for reconcile.)

Preview reconcile shows what packages will be installed and what packages will be removed. If a package is removed or installed as a result of another package being installed, the user is informed of the reason that the package must be removed or installed.

In some cases, installation and uninstallation require reboots. This information also displays during the preview reconcile.

When you use the install and uninstall patch wizards, Opsware SAS does not perform a preview reconcile.

Types of Reconcile

The two types of reconcile are partial and full.

During a partial reconcile, Opsware SAS only reconciles servers based on the nodes that the user has currently selected. For example, if a user has assigned a server to two nodes through the Install Software Wizard and then proceeds with the software installation, the server is reconciled only with those two nodes. If any other nodes have been assigned or removed through other means, such as nodes assigned through the Manage Servers list, these nodes are not reconciled.

During a full reconcile, a server is reconciled with all of the nodes that it has been assigned to. (If any nodes have been detached from the server, reconcile also uninstalls the software associated with those nodes.)

If any nodes were changed since they were attached to the server – for example, if patches were added to the nodes or if software was removed from the node, these changes are committed to the server during a full reconcile.

Reconcile on Supported Operating Systems

This section provides information on reconcile on supported operating systems and contains the following topics:

- AIX Reconcile
- HP-UX Reconcile
- Solaris Reconcile
- Linux Reconcile

After Opsware SAS determines what packages need to be installed or removed to complete the reconcile operation, reconcile uses a set of standard system utilities to complete the operation. The following table shows the utilities used during the reconcile session.

Table A-1: Utilities in Reconcile

SOLARIS	LINUX	AIX	HP-UX
patchadd (installs patches)	RPM (installs and removes software)	installp (installs software)	swinstall (installs software)
patchrm (removes patches)		installp -u (removes software)	swremove (removes software)
pkgadd (installs software)		inutoc (generates a table of contents of packages to be installed)	swlist (copies individual packages into one large depot)

Table A-1: Utilities in Reconcile (continued)

SOLARIS	LINUX	AIX	HP-UX
pkgrm (removes software)			swmodify (used to convert older format packages to a newer package format)
RPM (installs and removes software)		RPM (installs and removes software)	

See the *Opsware® SAS Configuration Guide* for more information about the package types that Opsware SAS supports.

AIX Reconcile

AIX software is delivered in LPPs, which are collections of filesets. When a server is reconciled and Opsware SAS determines that the reconcile requires filesets to be installed, Opsware SAS downloads the entire LPP that contains the filesets from the Software Repository to the server. If the filesets that the reconcile requires are contained in more than one LPP, the additional LPPs are also downloaded.

In AIX when you install a base fileset, it is always in the committed state and when you install an update fileset it is always in an applied state. When a fileset is an applied state, the previous version of the fileset is saved. When a fileset is a committed state, the previous version of the fileset is deleted. You can add the –c option in the install flag to install an update fileset in the committed state.

When you uninstall AIX filesets, the reconcile operation also uninstalls dependent filesets. The list of dependent filesets that are uninstalled appears in the reconcile status messages. The list of dependent filesets to be uninstalled does *not* appear in the Preview reconcile.

HP-UX Reconcile

HP-UX software is delivered in depots, which are collections of filesets. When a server is reconciled and Opsware SAS determines that the reconcile requires filesets to be installed, Opsware SAS downloads the entire depot that contains the filesets from the Software Repository.

If the filesets that the reconcile requires are contained in more than one depot, the additional depots are also downloaded. The depots are then combined into one large depot, from which the filesets will be installed.

HP-UX filesets often have dependencies on other filesets that the user has not specifically requested (or that have been included in a software node). These filesets can, however, be included in the HP-UX depot. By downloading the entire depot (instead of just the requested filesets), Opsware SAS is able to install any additional filesets that are required by the filesets associated with the nodes that are being reconciled. Combining the individual depots into one large depot allows the underlying installation utilities to locate all the filesets that require installation.

HP-UX 10.20 does not support the option `-x show_superseded_patches`; therefore, if you install a superseding patch, it removes the superseded patches. For example, if patch B supersedes patch A, the reconcile operation reports that patch A was removed when patch B is installed.

Solaris Reconcile

Solaris patches do not contain metadata that identify what cluster they belong to after the patch clusters are installed. During the reconcile process, however, Opsware SAS records the fact that the patches installed belong to a given patch cluster. This allows Opsware SAS to identify patch clusters that are installed on the servers. Opsware SAS can use this information to uninstall patch clusters.

Solaris patch clusters cannot be adopted by Opsware SAS. If the patch clusters were not installed through Opsware SAS, it is not possible to determine if a patch on a server originated from a patch cluster. Opsware SAS can, however, adopt individual patches.

Linux Reconcile

RPM is the only package type that Opsware SAS uses on the Linux operating system. When software is installed, the `-i` option is always used; when software is removed, the `-e` option is always used.

Specifying Scripts During Reconcile

When users upload software, they have the option of specifying scripts that should be run when software is installed or uninstalled. Reconcile executes these scripts on the servers local shell. For these scripts, users can elect to have reconcile react to a non-zero return code from the script by aborting reconcile operation when the non-zero return code is received. If a non-zero return code is encountered in a post-install script, Opsware SAS does not rollback or uninstall any software that has already been installed; again, the reconcile process simply halts when the non-zero return code is encountered (if this is the option that was selected for the script.)

These options are set when the software is uploaded to Opsware SAS and can be edited through the patch management and package management interfaces.

Reconcile Output

Reconcile provides detailed feedback about what occurs during the reconcile process, and what changes have been made on the servers selected for reconcile. Opsware SAS provides individual output for each server that has been selected for the reconcile operation. The output is the same for all wizards that use reconcile.

The output from a reconcile operation consists of the following types of data:

- A list of all software installed and uninstalled. If software is installed or uninstalled that was not specifically requested, but is required for the reconcile operation, the output specifies why the software was added or removed.
- If any pre-install or post-install scripts are executed, the first 1000 bytes of the scripts' `stdout` and `stderr` are displayed, as well as the return code for the script.
- Any reboots required by the reconcile operation.
- The output from the utilities that the reconcile operation used to install and uninstall the software. In some cases, these utilities might report errors. For example, a user might request AIX filesets to be installed that are dependent on other filesets that are not available to Opsware SAS. This error is reported as part of the reconcile output.

Assigning to and Removing Servers from Nodes

Servers are most often assigned and removed from software nodes by using the installation or uninstallation wizards for patches and software, or the installation wizard for templates. You can, however, manually assign or remove a server from a node.

If you manually assign or remove a server from a node, the software for that node is not installed or uninstalled until you use the Reconcile Wizard on that server.

Assigning Servers to Nodes

Perform the following steps to assign servers to nodes:

1. From the navigation panel in the Opsware Command Center, click Servers ➤ Server Search.
2. Use Server Search to find the server or servers that you want to assign to software nodes. The servers must all be running the same version of the same operating system.
3. From your search results list, select the server or servers that you want to assign to a software node.
4. From the **Tasks** menu, select **Assign Node**. The Assign Node Wizard appears.
5. Navigate to and select the node to which you want to assign the server.
6. Click **Assign**.

Removing Servers from Nodes

Perform the following steps to remove servers from nodes:

1. From the navigation panel in the Opsware Command Center, click Servers ➤ Server Search.
2. Use Server Search to find the server or servers that you want to remove from software nodes. If you are selecting multiple servers, the servers must have at least one node in common.
3. From your search results list, select the server or servers that you want to remove from a software node.
4. From the **Tasks** menu, select **Remove Node**.
5. Select the check boxes from the nodes that you want to remove and click **Remove**.

Reconcile Software Wizard

Servers are usually reconciled as a result of running any of the OS provisioning, patch management, or software provisioning wizards. Reconcile can, however, be directly invoked on a selected server or group of servers through the Reconcile Software Wizard. This wizard enables some of the "power user" flexibility that the other wizards hide.

Most often, users perform this direct reconcile when changes have been made to nodes, and the users want the servers that are already attached to these nodes to be reconciled to match the modified node.

Additionally, you can use the Reconcile Software Wizard to ensure that the server conforms exactly to its model. For example, if a node is deleted from the Software Tree and a server is attached to that node, the software for the node can only be uninstalled from the server by using the Reconcile Wizard.

You can select one or multiple servers to reconcile. If you select multiple servers, you can only reconcile software that is common to all of the servers you have selected (for example, you can only choose nodes that all servers have in common). You can, however, perform a full reconcile on a group of servers even if they do not have any software in common.

Directly Reconciling Servers

The Reconcile Software wizard is flexible when you run it on a server group. For example, when you select server groups, Opsware SAS will install the correct applications on the servers in the group, even if OS versions for the servers and the patches or applications do not all match. Opsware SAS also matches the customer association for applications, operating systems, and templates with the customer association of servers. If Opsware SAS cannot find a match, a "Valid devices not found" error message appears at the end of the wizard. Therefore, use caution when modifying these default values.

To directly reconcile servers, perform the following steps:

1 From the navigation panel in the Opsware Command Center, click Servers ➤ Server Search.

2 Use Server Search to find the server or servers that you want to reconcile.

3. Review the servers that your search returned and select the servers or groups that you want to reconcile.

4. From the **Tasks** menu, select **Reconcile**. The Reconcile Wizard appears.

5. If you have selected multiple servers, you can either select All Software to perform a full reconcile on all selected servers, or you can select Common Software to reconcile only the software common to all selected servers. If you selected a single server, you can either select to reconcile "Some Software" to choose the software (nodes) that you want to reconcile, or you can select "All Software" to perform a full reconcile.

6. Click **Next** to continue.

7. If you selected a partial reconcile (All Software or Common Software), you must now select the check boxes for the software that you want to reconcile your server with. If you selected a full reconcile, a list of all software to be reconciled displays and you can confirm your selection. Confirm your selections.

8. Click **Preview** to continue. A preview reconcile occurs. Review the results of the preview reconcile and click **Next**.

   Or

   Click **Skip Preview** to skip the preview process.

9. On the Schedule and Notify page, you have the following options:

   - **Notify**: Choose the Condition option and set the parameters to send an email when the operation is completed. You can choose to have an email sent on any result, if the operation fails, or if the operation is successful. To add another email recipient, click the plus (+) button next to the Recipients field.

   - **Schedule**: Choose either **Run Now** to execute the operation immediately, or choose **Specify Time** to schedule the operation for a later time.

     When you schedule a job for a server group, you can specify how the members of the group are determined. The membership of a dynamic server group changes based on the changes in your operational environment. If you have "Allow Run Refresh Jobs" permissions, you will see additional options. Select either of the following options:

- **Option 1**: Membership is determined based on the “Time of Confirm Selection.” Select this option to run the job on the servers that were in the group when you scheduled the job. Changes to the group membership do not affect the list of the servers that the job will run on.
- **Option 2**: Membership is updated when the job runs. Select this option to recalculate the group membership prior to running the job. Changes to group membership are reflected in the list of servers that the job will run on.

The time used for the scheduled job is specified in your preferred time zone which can be modified in My Profile. If you do not have the preferred time zone set, the time zone is derived from the Opsware SAS core server (usually UTC).

10 Click **Reconcile** to complete the process now.

Appendix B: Opsware Agent CLI Utilities

IN THIS APPENDIX

This section discusses the following topics:

- Agent Installation Using the CLI
- Opsware Agent Upgrade Tool

Agent Installation Using the CLI

This section provides information on Opsware Agent installation using the Agent Installer CLI and contains the following topics:

- Overview of Agent Installation Using the CLI
- Preparation for Opsware Agent Installation
- Preassimilation Checklist
- Installing an Opsware Agent by using the Agent Installer CLI
- Opsware Agent Installer Options
- Example: Opsware Agent Installer Command and Options
- Starting an Opsware Agent on a Server
- Verifying Opsware Agent Functionality
- Augmenting the Information for a Managed Server
- Uninstalling an Opsware Agent on Unix and Windows
- Uninstalling Earlier Versions of Opsware Agents on Unix
- Uninstalling Earlier Versions of Opsware Agents on Windows

Overview of Agent Installation Using the CLI

When you install Opsware Agents on existing operational servers, you should synchronize the local time on the servers with an external time-server that uses a network time protocol (NTP).

Installing Opsware Agents on servers makes existing operational servers known to Opsware SAS so that they can be managed. Assimilating servers into Opsware SAS is appropriate when many servers are already functioning in the operational environment and need to be managed (for example, when Opsware technology is initially deployed in a facility).

Installing an Opsware Agent on a server with a pre-built OS into Opsware SAS enables:

- Baseline discovery of the operating system on the server.
- Managing the baseline operating system, including patch management, when the operating system is defined in Opsware SAS with the Prepare Operating System Wizard.
- Full provisioning and management capabilities for any new applications deployed on the server.

When installed, the Opsware Agent registers the server with the Opsware Model Repository. Opsware SAS assigns the server to a generic operating system that corresponds to the operating system that the Opsware Agent discovered during the installation. The server is assigned to a placeholder OS node. For each operating system, the Opsware Command Center contains a node <operating_system_version>/Not Assigned, as Figure B-1 shows.

Figure B-1: Nodes Tab for an Assimilated Server

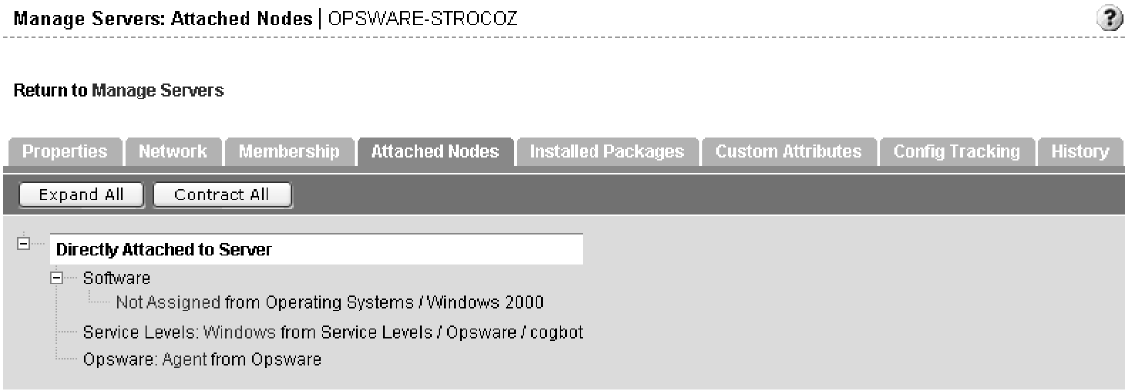

Manage Servers: Attached Nodes | OPSWARE-STROCOZ

Return to Manage Servers

Directly Attached to Server

Software

Not Assigned from Operating Systems / Windows 2000

Service Levels: Windows from Service Levels / Opsware / cogbot

Opsware: Agent from Opsware

The Opsware Agent Installer can install Agents when an Opsware core is not available to a server. If a newly-installed Agent cannot contact a core, the Agent runs in a dormant mode. While dormant, it periodically attempts to contact the core. When the core becomes available, the Agent performs the initialization tasks, such as hardware and software registration, that usually take place when the Agent is first installed.

The server is tracked in the Opsware Command Center. However, the server operating system cannot be managed while the server is assigned to the generic operating system node. You must reassign the server to the operating system that was defined with the OS Provisioning feature. (From the Manage Servers list, choose Servers ➤ Re-Assign Node.)

The server is associated with the default facility for the local instance of Opsware SAS.

If the managed server's IP address does not fall within a specified IP range, the server is associated with the default IP range group (Default). The default group is associated with the customer Not Assigned.

See the *Opsware® SAS Configuration Guide* for information about how servers are associated with customers.

Users install an Opsware Agent on each server. Running an Opsware Agent on a server allows Opsware SAS to manage the server. To install an Opsware Agent, run the Opsware Agent Installer.

The Opsware Agent Installer is an application that has the following features:

- Can be invoked from the command line or within a script.
- Installs an Opsware Agent.
- Logs its decisions and actions.
- Can be operated unattended because user interaction is not required.

The Opsware Agent Installer installs the Opsware Agent, retrieves cryptographic material, retrieves configuration information, and writes a configuration file.

Preparation for Opsware Agent Installation

Opsware Inc. recommends that you set up a Windows file share to make the Opsware Agent Installer for various operating systems available from one place. Setting up a file share allows you to install Opsware Agents on servers quickly and easily. If this is not possible, the Opsware Agent Installer needs to be moved by using an alternate file transfer mechanism, such as SFTP.

At the completion of the Opsware Agent installation process, a managed server is assimilated and the hardware and software data that the Opsware Agent discovered is stored in the Model Repository.

☑ To use the Patch Management features on Windows NT 4.0 and Windows 2000 servers, you must install Internet Explorer (IE) 6.0 or later on the server first because a patch utility depends on it. If you do not install IE 6.0 on the server first, the Opsware Agent Installer warns you that the Patch Management feature does not work as expected because it checks for IE 6.0 on all Window servers. This prerequisite is not required for Windows 2003, because IE 6.0 is pre-installed for this operating system.

Preassimilation Checklist

Prior to installation, you should perform the following tasks on the server where the Opsware Agent is to be installed. Performing these tasks is vital to installing the Opsware Agent quickly within maintenance windows.

1. For the Code Deployment & Rollback feature, verify that the following port is accessible between the server you will push code from and the server where you will push code to:

   - `telnet` <staging_server> `1002`

   - `telnet` <production_server> `1002`

2. Because the Opsware Agent runs on port 1002, verify that no other applications are using this port.

   - On a Unix server, enter this command from a terminal window:

     ```
     netstat -an | grep 1002 | grep LISTEN
     ```

   - On a Windows server, enter this command from a terminal window:

     ```
     netstat -an | find "1002" | find "LISTEN"
     ```

3. Check for sufficient disk space for Opsware Agent installation on the server.

   The Opsware Installer checks for the following amounts of free disk space in these directories:

   - 30MB in `/opt/OPSW/installdir` (Unix)

   - 100MB in `/var/lc/vardir` (Unix)

   - 30MB in `%SystemDrive%\\Program Files\\Loudcloud\\installdir` (Windows)

   - 100MB in `%SystemDrive%\\Program Files\\Common Files\\Loudcloud\\vardir` (Windows)

   (These default directories can be overridden with parameters at installation time.)

   These space requirements might not be enough. The `vardir` directory is used for dynamic content like logs and downloaded packages. If there is not enough disk space for the packages during a reconcile, reconcile will fail.

4. On the Solaris operating system, check for legacy sun4m architecture. Currently, the Opsware Agent works only for sun4u architecture.

5. For Windows, check the following items:

   - At a minimum, NT 4.0 Service Pack 6a must be installed on the server.

   - Verify that the Windows Registry has the correct settings:

   1. Start regedit and locate the following registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control
\FileSystem
```

2. Select the `NtfsDisable8dot3NameCreation` entry.
3. On the **Edit** menu, click **DWORD** and verify that the value is set to 0. The value must be set to 0. If necessary (because the value is set to 1), change the value by following your organization's IT policies and reboot the server.

6 To install an Opsware Agent on a server running Solaris, you must also install the following Solaris packages:

For Python:

```
SUNWtoo
SUNWtoox
```

For showrev:

```
SUNCadm
SUNWlibC
SUNWlibCx
SUNWadmfw
```

7 Before you install an Agent on a server, the server must meet certain patch requirements that vary by operating system, as Table B-1 shows.

Table B-1: Required Patches for Opsware Agent Installation

SERVER OPERATING SYSTEM	REQUIRED PATCHES
AIX 4.3	APAR IY39444
AIX 5.1	APAR IY39429 NOTE: If AIX 4.3.3.388, 4.4.4.89, or 5.1.0.3 is installed, the Opsware Agent Installer displays an error message that indicates the correct APAR to install on the server.
HP-UX (10.20, 11.00, 11.11/11i)	For HP-UX 10.20, PHCO_21018 Additionally, SW-DIST should be upgraded to the HP recommended patch level. You should continue to upgrade this package when HP recommends new versions.

Table B-1: Required Patches for Opsware Agent Installation (continued)

SERVER OPERATING SYSTEM	REQUIRED PATCHES
Linux AS 3.0	openssl096b NOTE: This package must be included in the `%packages` section of the Kickstart configuration.
Solaris 9, 8, 7, and 6	SUNWADMC SUNWcsl SUNWcsu SUNWesu SUNWlibms SUNWswmt
Windows 2000, NT 4.0	Service Pack 6a Internet Explorer 6.0 or later

Installing an Opsware Agent by using the Agent Installer CLI

The Opsware Agent needs administrator-level privileges (root on Unix servers and Local System on Windows servers) to manage a server. Therefore, Opsware Agent installation needs to be performed as root on Unix operating systems and as administrator on Windows operating systems.

You can install an Opsware Agent on any server listed in “Supported Operating Systems for Managed Servers” on page 115.

Perform the following steps to install an Opsware Agent on a server:

1 Log into the server that you want to assimilate by using a remote shell.

2 For Unix operating systems, change the user login to root (`su - root`) and for Windows operating systems, log in as administrator.

3 From the Opsware Command Center, download the package that contains the Opsware Agent Installer to a directory on the server you want to assimilate:

   1. Search for the package `opsware-agent`. From the navigation panel, enter `opsware-agent` in the Search box, select the Packages option, and click **Go**. The Manage Packages: Search Packages page appears.

Each operating system and operating system version has different packages for the Opsware Agent Installer.

Unix:

```
opsware-agent-<version>-<system_name>-<system_version>
```

Windows:

```
opsware-agent-<version>-<system_name>-<system_
version>.exe
```

2. Click the package name for the Opsware Agent Installer that you want to download. The Packages: Edit Properties page appears.
3. Click **Download** to save the package locally.

4 From the directory where the Opsware Agent Installer was copied, run the Installer by entering the correct executable and options for the installation environment.

See "Example: Opsware Agent Installer Command and Options" on page 726 in this appendix for more information.

Opsware Agent Installer Options

When you use the Opsware Agent Installer CLI, you can include the options that the following table shows to control the way that the Opsware Agent is installed on a server.

Table B-2: Agent Installer Options

OPTION	DESCRIPTION

`--clean`
`(-c)`
	Removes any machine-specific identifying material from the server. Specifically, removes the machine ID file (MID), and all machine-specific cryptographic material. Use this option when a server is deactivated and deleted from the Opsware Command Center and needs to be returned to service at a later time.

`-f`	Forces Opsware Agent installation and removes the target installation directory if it exists. REQUIREMENT: When using the `-f` option, you must run the Opsware Agent Installer as root on Unix operating systems and as the administrator on Windows operating systems.

Table B-2: Agent Installer Options (continued)

OPTION	DESCRIPTION
`--logfile`	Specifies the path to the Opsware Agent Installer log file. By default, the current directory is set as the path. By default, the log file has the following filename: `opsware-agent-installer-<`*date*`>.log`
`--loglevel <`*level*`>`	Sets the log level for log messages. With this option, specify one of the following levels: `error`, `warn`, `info`, `trace`, or `none`. The level `error` logs the least detail. The level `trace` logs all messages. By default, the log level is set to the log level `info`.
`-o`	Logs all output to `stdout` instead of a log file. This option is invoked automatically if the default log file or the log file passed with the `--logfile` option cannot be created, for example, when running the Opsware Agent Installer from non-writeable media, such as a DVD.
`--reconcile <`*type*`>`	Reconciles the server against any nodes assigned to the server. The <*type*> can be `full` or `addonly`. `full` – All nodes in a category are selected and reconcile removes software that Opsware SAS did not install. `addonly` – Software installed outside of Opsware SAS is not removed. WARNING: When assimilating a server that is already functioning in the operational environment, use caution when specifying the option `--reconcile`. If you specify this option, you might inadvertently uninstall software from the server.

Table B-2: Agent Installer Options (continued)

OPTION	DESCRIPTION
`--rpmbin <path>`	Specifies the path to the RPM binary to use for RPM operations. Use this option, when RPM is already installed on the server, to point the Opsware Agent at the RPM binary. Use the `--withrpm` option to install RPM if a usable instance of RPM is *not* already installed. NOTE: It is unnecessary to use this option with the `--withrpm` option.
`-s`	Starts the Opsware Agent after installing it. By default, the Opsware Agent Installer does not start the Opsware Agent.
`--template <ID>`	Assigns the nodes contained in the template to the server. `<ID>` can be an ID or a full name of a template. If this option is specified with the `--reconcile` option, Opsware SAS assigns the nodes in the template to the server before reconciling the server. WARNING: When assimilating a server that is already functioning in the operational environment, use caution when you specify the option `--template`. If you specify this option, you might inadvertently uninstall software from the server.
`--withmsi`	Installs MSI 2.0 along with the Opsware Agent. If MSI 2.0 is already installed, this option has no effect. Works with Windows NT 4.0 Service Pack 6a, Windows 2000, and Windows 2003.
`--withwmi`	Installs WMI 1.5 along with the Opsware Agent. If WMI 1.5 is already installed, this option has no effect. Works with Windows NT 4.0 Service Pack 6a.

Table B-2: Agent Installer Options (continued)

OPTION	DESCRIPTION
`--withrpm`	Installs the RPM handler with the Opsware Agent. By default, an Opsware Agent is not installed with this option. Opsware Inc. recommends that you always include the `--withrpm` option when you install Opsware Agents on Solaris servers. NOTE: Use the `--withrpm` option only with the Opsware Agent Installers for these operating systems: Solaris 5.6, 5.7, 5.8, and 5.9, and AIX 4.3 and AIX 5.1. On Solaris, RPM 3.0.6 is installed in the directory `/opt/OPSWrpm` and the RPM database is installed in the directory `/var/opt/OPSWrpm/lib/rpm`. On AIX, RPM 3.0.5 is installed in the directory `/opt/freeware` and the RPM database is installed in the directory `/var/opt/freeware/lib/rpm`.
`--workdir` <*path*>	Specifies the path to the Opsware Agent Installer temporary working directory. Use this option if the default working directory causes problems with installation.
`--reboot`	During Opsware Agent Installation on a Windows server, the Agent Installer copies the ogshcap.dll file to the following location: `%SystemRoot%\system32\ogshcap.dll` If the file is open or is in use, the Agent Installer is unable to copy the ogshcap.dll file. The Agent Installer then informs the user whether to restart the machine and copies the file after restart. You can specify the `--reboot` Installer option in the Opsware Command Line to initiate the reboot at the end of the Agent installation.
`--resetconf(-r)`	Resets Opsware Agent configuration file to default the settings.

Table B-2: Agent Installer Options (continued)

OPTION	DESCRIPTION
`--no_anonymous_ssl (-A)`	Disables anonymous SSL. This option applies to dormant Opsware Agents only. This option configures the Opsware Agent so that browsers cannot connect without a valid certificate.
`--settime (-t)`	Synchronizes the time on the server on which the Opsware Agent is installed with that of the Opsware core. NOTE: If the server on which the Opsware Agent is being installed is significantly ahead of the clock on the Opsware core, then the clock on the managed server is set back in time. Since this can cause problems, do not use the --settime option unless you are sure that this scenario is not a problem in your environment. If a managed server's clock is significantly behind of the clock on the Opsware core, the Opsware Agent installation might fail. To install an Opsware Agent successfully, use the `--settime` option or manually set the time and date on the managed server before retrying the Opsware Agent installation.

Example: Opsware Agent Installer Command and Options

Enter the following command and options to install the Opsware Agent for Solaris 5.7 in the default directories and log the results of the installation in the log file:

```
% opsware-agent-14.2.12.5-solaris-5.7 --logfile opsware-agent-
installer.log --loglevel info
```

Enter the following command and options to install the Opsware Agent for Windows NT 4.0 in the default directories and log the results of the installation in the log file:

```
% opsware-agent-14.2.12.5-win32-4.0.exe --logfile opsware-
agent-installer.log --loglevel info
```

Starting an Opsware Agent on a Server

If you do *not* include the `-s` option on the command line when you install an Opsware Agent, you have to start the Opsware Agent on the server manually.

Enter the following command, which varies with the operating system:

Solaris:

```
/etc/init.d/cogbot start
```

Linux:

```
/etc/rc.d/init.d/cogbot start
```

AIX:

```
/etc/rc.d/init.d/cogbot start
```

HP-UX:

```
/sbin/init.d/cogbot start
```

Windows:

```
net start shadowbot
```

Verifying Opsware Agent Functionality

Perform the following steps to verify Opsware Agent functionality:

1. From the navigation panel in the Opsware Command Center, click Servers ➤ Manage Servers. The Manage Servers page appears. Browse the list to find the server whose Opsware Agent installation you want to verify. If necessary, select the correct customer and facility for the server and click **Update**.

   Or

   Search for the server whose Opsware Agent installation you want to verify.

   Or

2. Verify that the server appears in the Manage Servers list and has the correct properties.

   See “Using the Search Feature” on page 125 in Chapter 4 for more information. See “Server Searching by IP Address” on page 141 in Chapter 4 for more information.

3. If you want to discover reasons why a server is unreachable, you can run a Communication Test. See "Agent Reachability Communication Tests" on page 204 in Chapter 6 for more information.

Augmenting the Information for a Managed Server

Use caution when you augment the discovery process for an Opsware-managed server that is functioning in the operational environment. You might inadvertently install or uninstall software from the server. During the test reconcile, verify what software will be uninstalled from the server before you perform the actual reconcile.

Perform the following steps to augment the information for an Opsware-managed server:

1. Model the OS and other applications running on the server in Opsware SAS by defining the OS with the Prepare Operating System Wizard and by creating nodes and templates for applications running on the assimilated server.

   See the *Opsware® SAS Configuration Guide* for more information about Operating System Definitions.

2. Move the server to the appropriate nodes for the OS and installed applications.

   The server is tracked in the Opsware Command Center; however, the server operating system *cannot* be managed while the server is assigned to the generic operating system node. You must reassign the server to the operating system that was defined with the OS Provisioning feature. (From the Manage Servers list, choose **Servers ➤ Re-Assign Node**.)

3. Reconcile the server. See "Directly Reconciling Servers" on page 711.

4. If an IP range group was set up, servers are automatically associated with customers when users install an Opsware Agent on the servers. Otherwise, the servers are associated with the Not Assigned customer. To change the customer associated with a server, See "Editing the Properties of a Server" on page 265 in Chapter 7 for more information.

5. To specify the server's use, stage, and state, edit the server's properties. See "Editing the Properties of a Server" on page 265 in Chapter 7 for more information.

Discovery is complete. Opsware SAS assumes that the server should always be running the specific OS build it has been associated with. Any changes to the OS outside of Opsware SAS are not captured in the model.

Users can deploy and manage new applications on the server, just as if Opsware SAS initially provisioned the server. Users can also deploy OS level patches on the server, or rebuild the OS by using the OS build with which the server was associated.

Uninstalling an Opsware Agent on Unix and Windows

To uninstall Opsware Agents on Windows NT, Windows Scripting Host 5.1 or Internet Explorer 5.5 must be installed on Windows NT.

Perform the following steps to uninstall an Opsware Agent on Unix or Windows:

1. Log into Unix as root user. Log into Windows as Administrator.
2. Change directories to any directory other than the Opsware Agent's installation directory.
3. On Unix, enter the following command:

   ```
   <installation_directory>/bin/agent_uninstall.sh
   ```

   By default, for Solaris and AIX, the Opsware Agent Uninstaller will not remove the Opsware RPM package. For command line options for the agent uninstaller, including how to activate removal of the Opsware RPM package, See “Opsware Agent Uninstaller Options” on page 730 in this appendix for more information.
4. On Windows, enter the following command:

   ```
   msiexec /x <installation_directory>\bin\agent_uninstall.msi
   ```
5. As the uninstall proceeds, the Unix platform `stdout` shows the uninstallation progress. The Windows uninstall does not show uninstallation progress.

Opsware Agent Uninstaller Options

When you use the Opsware Agent Uninstaller, you can include the options that Table B-3 and Table B-4 show.

Table B-3: Opsware Agent Uninstallation Unix Options

OPTION	DESCRIPTION
`--uninstallerVersion`	Show the uninstaller version.
`--help`	Show this help.
`--no_deactivate`	Do not deactivate the server; by default, the server is deactivated.
`--force`	Do not prompt for confirmation before deactivating the server.
`--delete_opsw_rpm`	Remove the OPSW RPM package (AIX, Solaris only). Use the following commands to remove the RPM package: Solaris: `pkgrm -n OPSWrpm` AIX: `installp -u rpm.rte`

Table B-4: Opsware Agent Uninstallation Windows Options

OPTION	DESCRIPTION
`NO_DEACTIVATE="1"`	Do not deactivate the server; by default the server is deactivated.
`FORCE="1"`	Do not prompt for confirmation before deactivating the server.

During Opsware Agent Uninstallation on a Windows server, the Agent Installer removes the ogshcap.dll file from the following location:

```
%SystemRoot%\system32\ogshcap.dll
```

If the file is open or is in use, the Agent Installer is unable to remove the ogshcap.dll file. The Agent Installer then prompts the user to restart the machine and removes the file after restart.

Uninstalling Earlier Versions of Opsware Agents on Unix

Perform the following steps to uninstall Opsware Agents versions 5.1 and earlier:

1. Stop the Opsware Agent on the server by running the following command as root:

Linux:

```
/etc/rc.d/init.d/cogbot stop
```

Solaris:

```
/etc/init.d/cogbot stop
```

HP-UX:

```
/sbin/init.d/cogbot stop
```

AIX:

```
/etc/rc.d/init.d/cogbot stop
```

2. Deactivate or delete the server by using the **Opsware Command Center Server** menu.

3. For Linux servers only, run `chkconfig` to de-register the Opsware Agent initialization script:

```
% /sbin/chkconfig -del cogbot
```

4. As root, delete the following files and directories to remove the Opsware Agent files from the server:

Linux:

```
/etc/rc.d/init.d/cogbot
```

Solaris:

```
/etc/init.d/cogbot

/etc/rc2.d/S79cogbot

/etc/rc0.d/K44cogbot
```

HP-UX:

```
/sbin/init.d/cogbot

/sbin/rc2.d/cogbot
```

AIX:

```
/etc/rc2.d/init.d/cogbot

/etc/rc.d/S79cogbot
```

All Unix:

```
/opt/OPSW

/var/lc
```

Uninstalling Earlier Versions of Opsware Agents on Windows

Perform the following steps to uninstall earlier versions of Opsware Agents on Windows:

1. Stop the Opsware Agent by running the following command as administrator:

   ```
   C:\> net stop shadowbot
   ```

2. Deactivate or delete the server by using the **Opsware Command Center Server** menu.

3. Deregister the Opsware Agent service by running the following command as administrator:

   ```
   C:\> "%SystemDrive%\Program
   Files\Loudcloud\blackshadow\watchdog\watchdog.exe" -x
   ```

4. As administrator, delete the following directories to remove the Opsware Agent:

   ```
   "%SystemDrive%\Program Files\Loudcloud"

   "%SystemDrive%\Program Files\Common Files\Loudcloud"
   ```

Opsware Agent Upgrade Tool

This section contains the following topics:

- Ways to Upgrade Opsware Agents
- Prerequisites for Using the Opsware Agent Upgrade Tool
- Upgrading the Opsware Agent on Managed Servers
- Commands for the Opsware Agent Upgrade Tool
- Options for the Opsware Agent Upgrade Tool

- Example: Options for the Opsware Agent Upgrade Tool
- Example: Commands and Output for Agent Upgrade Tool

Ways to Upgrade Opsware Agents

After you upgrade Opsware SAS running in a facility, you should upgrade the Opsware Agents on every managed server to the new version, so that you can utilize the new features in the newly-upgraded core.

Opsware SAS features continue to work on a managed server even when it is running an older Opsware Agent. However, new features in the new versions might not be available for that server.

Refer to the Release Notes for the new version for information about the compatibility of new features with older agents.

You can upgrade the Opsware Agents on managed servers in the following ways:

- Use the Opsware Agent Installer (a command line interface) to install a new Opsware Agent on one server at a time.

  See “Agent Installation Using the CLI” on page 715 in this appendix for information about how to use the Opsware Agent Installer.

- Use the Opsware Agent Upgrade Tool to upgrade Opsware Agents on groups of servers. Running the tool upgrades deployed Opsware Agents on managed servers. You can run the script simultaneously on many servers to upgrade large groups of Opsware Agents.

The Opsware Agent Upgrade Tool has the following characteristics:

- It is a command line interface that provides a flexible mechanism for selecting servers to upgrade, and for monitoring and reviewing upgrade operations.
- You can use it to upgrade many Opsware Agents on managed servers simultaneously.
- It runs within your preferred Unix shell, allowing it to leverage the power of standard Unix shells and text processing tools.
- You can use it to upgrade a server in any facility running Opsware SAS. You can run it from an Opsware shell attached to any Opsware SAS in any facility.

The Opsware shell is a program that authenticates users in Opsware SAS, starts the user’s normal Unix shell (as specified in the standard password database). Using the Opsware shell allows the user to run the Opsware Agent Upgrade Tool in this facility.

Prerequisites for Using the Opsware Agent Upgrade Tool

- On a core server of the facility being upgraded, install the Opsware Shell RPM by downloading the `opsh` package from the Opsware Command Center. See the *Opsware® SAS Configuration Guide* for instructions on downloading a package.

  Installing the `opsh` RPM places the Opsware shell and the Opsware Agent Upgrade Tool in the directory `/opt/OPSWopsh/bin`.

- You need the correct permissions to upgrade Opsware Agents. Run the Opsware shell by specifying the Opsware `admin` user and password to ensure that you have the appropriate permissions. (Contact your Opsware administrator to obtain the password.)

  When you start an Opsware shell to run the Opsware Agent Upgrade Tool, the user name and password are authenticated by Opsware SAS.

- The server where you install the `opsh` RPM must be able to resolve the name `way.<facility-domain>` to the host running the Command Engine in the facility's core. For example, if the Command Engine runs on a host in the `prod.opsware.com` domain, the servers must resolve the name `way.prod.opsware.com`. You specified the `facility-domain` when you installed the core.

Upgrading the Opsware Agent on Managed Servers

To upgrade the Opsware Agent, perform the following steps:

1. After you install the `opsh` RPM on a core server, enter the following command as root to start the Opsware shell:

   ```
   opsh [username@]facility-domain
   ```

   For example:

   ```
   opsh admin@prod.opsware.com
   ```

   See the preceding section for the name resolution requirement for `facility-domain`. See "Commands for the Opsware Agent Upgrade Tool" on page 735 for more information on `opsh`.

2. (Optional) To obtain information about the current Opsware Agents running on the managed servers before you upgrade them, enter any of the following commands and options:

   ```
   opsh_agent query server-options
   ```

(Enter this command if you want to view a report of the Opsware Agent versions running on the servers before you upgrade them.)

```
opsh_agent verify server-options schedule-options \
```

agent-version

(Enter this command if you want to verify the versions of the Opsware Agents running on the managed servers before you upgrade them.)

3 To upgrade Agents on specified servers, enter the following Opsware Agent Upgrade Tool commands and options:

```
opsh_agent stage server-options schedule-options \
[--always] agent-version
```

(Enter this command if you want to download the package for the Opsware Agent to the managed server before you run the upgrade.)

```
opsh_agent upgrade server-options schedule-options \
[--always] agent-version
```

4 (Optional) To review the status of the Opsware Agent upgrade, enter the following command and option:

```
opsh_agent review session-id
```

Commands for the Opsware Agent Upgrade Tool

- `opsh [username@]` facility-domain

  This command starts an Opsware shell and authenticates the user name against the Opsware facility running at the specified domain.

  If you do not specify a user name, the currently logged in user name is used. The Opsware shell prompts for a password.

  A new Unix shell (which is attached to the specified Opsware core-domain) is started. (The password database for the user specifies which Unix shell to use.)

- `opsh_agent query` server-options

  This command must be run from an Opsware shell started with the `opsh` command.

  This command queries the reported version of Opsware Agents and any staging status for the specified servers by examining data in the Model Repository.

One line is printed to stdout for each server that shows device ID, IP address, current Opsware Agent version, and any staging status.

You can specify the servers by using the `--device, --customer, --facility, and --os` options.

- `opsh_agent stage server-options schedule-options \`

  `[--always] agent-version`

  You must run this command from an Opsware shell started with the `opsh` command.

  This command contacts the Opsware Agent on each specified server and instructs it to download the package for the specified version of the Opsware Agent from the Software Repository.

  If the download is successful, the staging status is written to the Model Repository for the server.

  To download the package to the server even when this command was entered previously (recorded in the Model Repository), specify the `--always` option.

  One line is printed to stdout for each server that shows the device ID, IP address, and a success or failure indicator.

  You can specify the servers by using the `--server, --customer, --facility and --os` options.

  A session is started and the session ID displays for later review. After the session ID displays, you can type `CTRL-C` and review the session later using the `opsh_agent review` command.

- `opsh_agent upgrade server-options schedule-options \`

  `[--always] agent-version`

  You must run this command from an Opsware shell started with the `opsh` command.

  This command contacts the Opsware Agent on each specified server and instructs it to upgrade to the specified version. If the necessary package has not been downloaded on the server already (the `opsh_agent stage` command was entered), the package is downloaded from the Software Repository.

  If the upgrade is successful, the package is removed from the server and the staging status is deleted from the Model Repository.

To upgrade the Opsware Agent even when the specified version of the Opsware Agent was already installed on the managed servers, enter the `--always` option. (The Model Repository records when Opsware Agents are upgraded on servers.)

One line is printed to stdout for each server that shows the device ID, IP address, and a success or failure indicator.

You can specify the servers by using the `--server`, `--customer`, `--facility`, and `--os` options.

A session is started and the session ID displays for later review. After the session ID displays, you can type CTRL-C and review the session later by using the `opsh_agent review` command.

- `opsh_agent verify server-options schedule-options \`

  `agent-version`

  You must run this command from an Opsware shell started with the `opsh` command.

  This command contacts the Opsware Agent on each specified server to verify that it is running the specified version.

  One line is printed to stdout for each server that shows the device ID, IP address, the word OLD, NEW, or CURRENT and the actual Opsware Agent version running on the server.

  You can specify the servers by using the `--server`, `--customer`, `--facility`, and `--os` options.

  A session is started and the session ID displays for later review. After the session ID displays, you can enter CTRL-C and review the session later by using the `opsh_agent review` command.

- `opsh_agent review session-id`

  You must run this command from an Opsware shell started with the `opsh` command; although, not necessarily the same Opsware shell from which the original command was started.

  This command attaches to a running `opsh_agent stage, opsh_agent upgrade` or `opsh_agent verify` session running on the Command Engine. It prints the same output to stdout that the original command would have printed if the user had not typed CTRL-C and terminated the command. If the session is complete, it shows the same results that were shown when the session completed.

Options for the Opsware Agent Upgrade Tool

Server-options: `--server|-S <svr-spec> --customer|-C <cust-spec> --facility|-F <fac-spec> --os|-O <os-spec>`

If more than one of the `--customer`, `--facility`, or `--os` options is specified, only servers that match all options are selected. Any servers specified by using the `--server` option are added to (or subtracted from) the list specified by combining the `--customer`, `--facility`, and `--os` options.

Table B-5: Options for the Opsware Agent Upgrade Tool

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--server`	`-S`	`<svr-spec>`	Server by device ID, IP address, or system name
`--customer`	`-C`	`<cust-spec>`	All servers associated with the customer specified by the customer ID or name
`--facility`	`-F`	`<fac-spec>`	All servers in the facility specified by the facility ID or name
`--os`	`-O`	`<os-spec>`	All servers running the operating system specified by the OS name

Schedule-options: `--when|-W when-time --until|-U until-time`

Table B-6: Schedule Options

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--when`	`-W`	`<when-time>`	Start time for a stage, upgrade, verify, or test operation in the format: MM/DD/YYY-HH:MM If the `--when` option is used, the operation starts at the specified time, but the command displays a session ID and returns immediately. When you schedule an operation, you use the `review` command to review the results after the operation has run. If the `--when` option is not specified, the operation starts immediately and the command displays the output of the command.
`--until`	`-U`	`<until-time>`	End time for a stage, upgrade, verify or test operation in the format: MM/DD/YYY-HH:MM If the `--until` option is specified, the operation stops processing servers at the specified time. Any servers that are not complete are left in a consistent state; this might require that the session run past the specified time.

Miscellaneous options:

`--ip|-I --always|-A --parallel|-P --theword|-T`

Table B-7: Miscellaneous Options

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--ip`	`-I`	(N/A)	Display IP addresses instead of host names.

Table B-7: Miscellaneous Options (continued)

LONG OPTION	SHORT OPTION	VALUE	MEANING
`--always`	`-A`	(N/A)	Always stage or upgrade servers even if the current version is staged or upgraded.
`--parallel`	`-P`	`<Concurrency>`	The maximum of concurrent commands. (Recommended default = 10)
`--theword`	`-T`	`<hostname>`	The host name or IP address to use when contacting the Software Repository from a server.

Example: Options for the Opsware Agent Upgrade Tool

The following table provides examples for running the Opsware Agent Upgrade Tool.

Table B-8: Examples of Options for the Opsware Agent Upgrade Option

EXAMPLE	DESCRIPTION
`--server 1,2`	Selects servers 1 and 2.
`--facility Y,Z`	Selects all servers (for all customers) in facilities Y and Z.
`--customer -A,-B --facility Z`	Selects all servers in facility Z except those owned by customers A and B.
`--server 1,2,-3,-4 --customer A,B --facility Y,Z`	Select servers 1 and 2 as well as all servers owned by customers A or B which are in facilities Y or Z except servers 3 and 4.
`--server 1,-2 --customer A,B --facility -Y,-Z --os SunOS 5.8`	Select server 1 and all servers owned by customers A or B, except those in facilities Y or Z and which are Solaris 5.8 machines excluding server 2.

Example: Commands and Output for Agent Upgrade Tool

```
# cd /opt/OPSWopsh/bin
# ./opsh admin@prod.opsware.com
admin@prod.opsware.com's password:
#
# ./opsh_agent verify --os "SunOS*" 14a.2.12.18
Session 37802500101L
Device ID Name/IP address Version Result Status Reason
410101L core2-1.prod.opsware.com  14a.2.12.18 CURRENT SUCCESS
^C
Interrupted review of running session 37802500101L
Use review 37802500101L command anytime to review session status
#
# ./opsh_agent review 37802500101L
Session 37802500101L
Device ID Name/IP address Version Result Status Reason
410101L d033.prod.opsware.com 14a.2.12.18 CURRENT SUCCESS
670101L dhcp-174.prod.opsware.com 14a.2.12.16 OLDER SUCCESS
1460100L emb218-37.manu.opsware.com 14a.2.12.18 CURRENT SUCCESS
20100L f001.manu.opsware.com 14a.2.12.18 CURRENT SUCCESS
10100L f002.manu.opsware.com 14a.2.12.21 NEWER SUCCESS
210100L m022.manu.opsware.com 14a.2.12.18 CURRENT SUCCESS
Session 37802500101L completed.
```

Appendix C: Communication Test Troubleshooting

IN THIS APPENDIX

This section discusses the following topics:

- Overview of Agent Communication Tests
- Command Engine to Agent (AGT)
- Crypto Match (CRP)
- Agent to Command Engine (CE)
- Agent to Data Access Engine (DAE)
- Agent to Software Repository (SWR)
- Machine ID Match (MID)
- Common Troubleshooting Tasks

Overview of Agent Communication Tests

The Communication Test performs the following diagnostic tests to determine if an Agent is reachable:

- **Command Engine to Agent (AGT)**: Determines if the Command Engine can communicate with the Agent. The Command Engine is the Opsware SAS component that enables distributed programs to run across many servers. The Command Engine handles the entry of scripts into the Opsware Model Repository (the script storage location in Opsware SAS) and the versioning of stored scripts.
- **Crypto Match (CRP)**: Checks that the SSL cryptographic files that the Agent uses are valid.
- **Agent to Command Engine (CE)**: Verifies that the Agent can connect to the Command Engine and retrieve a command for execution.

- **Agent to Data Access Engine (DAE)**: Checks whether or not the Agent can connect to the Data Access Engine and retrieve its device record. The Data Access Engine provides an XML-RPC interface to the Model Repository that simplifies interaction with various clients such as the Opsware Command Center, system data collection, and monitoring agents on servers.
- **Agent to Software Repository (SWR)**: Determines if the Agent can establish an SSL connection to the Software Repository. The Software Repository is the central repository for all software that the Opsware system technology manages. It contains software packages for operating systems, applications, databases, customer code, and software configuration information.
- **Machine ID Match (MID)**: Checks that the Machine ID (MID) on the server matches the MID registered in the Model Repository for the Agent.

When the test run finishes, it returns results that show either success or failure for each test run on each server. For each failed test, the nature of failure is listed by error type in the error details column of the Communication Test window. In some cases, the failure of one test might prevent other tests from being executed.

See "Agent Reachability Communication Tests" on page 204 in Chapter 6 for information about how to run Communication Test.

Command Engine to Agent (AGT)

The Command Engine to Agent (AGT) communications test system checks that the Command Engine can initiate an SSL connection to the Agent and execute an XML/RPC request.

The thirteen possible results are:

- AGT – OK
- AGT – Untested
- AGT – Unexpected error
- AGT – Connection refused
- AGT – Connection time-out
- AGT – Request time-out
- AGT – Server never registered

- AGT – Realm is unreachable
- AGT – Tunnel setup error
- AGT – Gateway denied access
- AGT – Internal Gateway error
- AGT – Gateway could not connect to server
- AGT – Gateway time-out

AGT – OK

No troubleshooting necessary.

AGT – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the Command Engine cannot contact the Agent, then no other tests are possible.

What Can I Do If a Test Is Not Run During an AGT Test?

First resolve all tests that failed, and then run the Communication Test again.

AGT – Unexpected error

This result indicates that the test encountered an unexpected error.

What Can I Do If I Get an Unexpected Error?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware, Inc. Customer Support.

AGT – Connection refused

This result indicates that the Command Engine is receiving a TCP reset packet when it attempts to connect to the Agent on port 1002. The likely cause is that the Agent is not running. A firewall might also be blocking the connection.

What Can I Do If the Connection is Refused During an AGT Test?

1. Log into the server and confirm that the Agent is running. See “Verifying that an Agent is Running” on page 768 in this appendix for more information.

2. If the Agent is not running, restart the Agent. See "Restarting an Opsware Agent" on page 770 in this appendix for more information.
3. From the managed server, use netstat to confirm that a socket is in listen mode on port 1002. If not, stop and restart the Agent.
4. From the server itself, use Telnet to connect to the IP address of the server where the Agent is installed and port (1002) that the Agent is listening on. If this does not succeed, stop and restart the Agent.
5. Verify that the Management IP address that Opsware SAS is using to reach the server is the correct address. See "Checking Management IP of a Managed Server" on page 770 in this appendix for more information. If the IP addresses do not match, stop and restart the Agent, then rerun the test.
6. If the previous steps are performed and the test still fails, the problem is likely caused by either a software-based firewall on the server itself or an external firewall blocking the connection.

AGT – Connection time-out

This result indicates that the Command Engine is not receiving any reply packets when it attempts to initiate a TCP connection to the Agent on port 1002. The likely cause is that the server is not running, or that the IP address that Opsware SAS is using to reach the Agent is incorrect. (A firewall might also be blocking the connection.) To check the IP address that Opsware SAS is using to reach the Agent, see "Checking Management IP of a Managed Server" on page 770.

What Can I Do If the Connection Times Out During an AGT Test?

Follow the same steps used to resolve this issue specified in "What Can I Do If the Connection is Refused During an AGT Test?" on page 745.

AGT – Request time-out

This result indicates that the Command Engine is able to successfully complete a TCP connection to the Agent on port 1002, but no response is received from the Agent in response to the XML-RPC request. The likely cause is that the Agent is hung.

What Can I Do If the Request Times Out During an AGT Test?

1. Log into the server and restart the Agent. See "Restarting an Opsware Agent" on page 770 in this appendix for more information.

2. Check to see whether or not some other process is consistently utilizing an excessive amount of the CPU on the server where the Agent is installed. Also check to see if the system is performing slowly due to a lack of available memory and/or excessive file IO. In any of these cases, the system might be performing too slowly to permit the Agent to respond to the test in a timely manner.

AGT – Server never registered

This test indicates that the server being tested has neither been registered with the Command Engine, nor can it communicate with the Command Engine. The cause of this could be any number of reasons similar to those in the Agent to Command Engine (CE) test. It is also possible (but unlikely) that the Agent was installed but never started.

What Can I Do If the Server Has Not Been Registered with the Command Engine During an AGT Test?

To troubleshoot this error, use the following procedures:

1. Ensure that the Agent is running. For these instructions, See “Verifying that an Agent is Running” on page 768 in this appendix for more information.
2. Ensure that the Agent can contact the Command Engine.
3. If the Agent is in a Satellite facility, ensure that its Gateways are properly configured and that it is properly configured to use those Gateways. See “Checking Network Gateway Configuration” on page 770 in this appendix for more information.
4. If the Agent is not in a Satellite:
    - Ensure the host name “way” (no quotes) resolves to its valid IP address. See “Resolving Host Name” on page 771 in this appendix for more information.
    - Verify that a connection can be established to port 1018 of way. Use the command “telnet way 1018" (or equivalent).

One (or more) of the above checks will fail. To solve that failure, refer to the corresponding error code for the Agent to Command Engine (CE) test on page 751, or to the realm connectivity and configuration test.

AGT – Realm is unreachable

The Satellite realm where the managed server is located is unreachable. This means that a path of tunnels between the Gateways in the Opsware core and the realm of the managed server cannot be established.

What Can I Do If the Realm Is Unreachable During an AGT Test?

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact Opsware, Inc. Customer Support for assistance in troubleshooting the Gateway network.

AGT – Tunnel setup error

The Command Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do If I Get a Tunnel Setup Error During an AGT Test?

Contact your Opsware administrator.

AGT – Gateway denied access

The Gateway is working but refused to proxy the connection on behalf of the Agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the Command Engine access to the Agent.

What Can I Do If the Gateway is Denied Access During an AGT Test?

Contact your Opsware administrator.

AGT – Internal Gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do If There is an Internal Gateway Error During an AGT Test?

Contact your Opsware administrator.

AGT – Gateway could not connect to server

The Gateway could not establish a connection to the Agent. This might be because the Agent is not running, or because a firewall might be blocking the connection.

What Can I Do If the Gateway Couldn't Connect to the Server During an AGT Test?

If you suspect the Agent is not running, see "Verifying that an Agent is Running" on page 768. To make sure that the Gateway can establish a connection to the IP address of the server where the Agent is installed, try to ping the IP address of the server where the Agent is installed.

AGT – Gateway time-out

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do If the Gateway Times Out During an AGT Test

Ensure that network connectivity is available between the Gateways in the path between the realm of the managed server and the Opsware core.

Crypto Match (CRP)

This test checks that the SSL cryptographic files that the Agent uses are valid.

The five possible results are:

- CRP – OK
- CRP – Untested
- CRP – Unexpected error
- CRP – Agent certificate mismatch
- CRP – SSL negotiation failure

CRP – OK

No troubleshooting necessary.

CRP – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the Agent cannot be reached, then no other tests are possible.

What Can I Do If a Test Is Not Run During a CRP Test?

First resolve all tests that failed, and then run the Communication Test again.

CRP – Unexpected error

This result indicates that the test encountered an unexpected error.

What Can I Do If I Get an Unexpected Error During a CRP Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware, Inc. Customer Support.

CRP – Agent certificate mismatch

This result indicates that the SSL certificate that the Agent is using (cogbot.srv) does not match the SSL certificate that is registered with Opsware SAS for that Agent.

What Can I Do If I Get a Certificate CN Mismatch During a CRP Test?

Use the Recert Agent Custom Extension to issue a new certificate to the Agent. See "Opsware SAS Custom Extensions" on page 440 in Chapter 11 for more information.

CRP – SSL negotiation failure

This result indicates that the Agent is not accepting SSL connections for the Opsware core. (The Opsware core is the entire collection of servers and services that provide Opsware SAS services.) The likely cause of this error is that one or more files in the Agent crypto directory are missing or are invalid.

What Can I Do If I Get an SSL Negotiation Failure During an CRP Test?

Run the Server Recert custom extension in the "set allow recert flag only" mode on the server, and then Run the Opsware Agent Installer with the "-c" switch.

Reinstalling the Agent with the "-c" option ("c" stands for "clean") removes all certs on the server and also removes the MID file, which forces the Agent to retrieve a new MID from the Data Access Engine.

- See "Opsware SAS Custom Extensions" on page 440 in Chapter 11 for more information how to run the Recert Agent Custom Extension to issue a new certificate to the Agent.
- See "Agent Reachability Communication Tests" on page 204 in Chapter 6 for information abouthow to install an Opsware Agent using the "-c" switch.

After you reinstall the Agent, run the test again to check if the Agent is now reachable.

Agent to Command Engine (CE)

This test checks that the Agent can connect to the Command Engine and retrieve a command for execution.

The sixteen possible results are:

- CE – OK
- CE – Untested
- CE – Unexpected error
- CE – Connection refused
- CE – Connection time-out
- CE – DNS does not resolve
- CE – Old Agent version
- CE – Realm is unreachable
- CE – No Gateway defined
- CE – Tunnel setup error
- CE – Gateway denied access
- CE – Gateway name resolution error
- CE – Internal Gateway error
- CE – Gateway could not connect to server
- CE – Gateway time-out
- CE – No callback from Agent

CE – OK

No troubleshooting necessary.

CE – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the Agent cannot reach the Command Engine, then no other tests are possible.

What Can I Do If a Test Is Not Run During a CE Test?

First resolve all tests that failed, and then run the Communication Test again.

CE – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During a CE Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware, Inc. Customer Support.

CE – Connection refused

This result indicates that the Agent is receiving a TCP reset packet when attempting to connect to the Command Engine on port 1018. The likely cause is that the Agent is connecting to the wrong IP address. In other words, the Agent does not know the correct IP address of the Command Engine. It is also possible that a firewall might be blocking the connection.

What Can I Do If the Connection is Refused During a CE Test?

1. Check that the name "way" resolves to its correct IP address. For instructions on how to do this, see "Resolving Host Name" on page 771.

2. Check to make sure there isn't a firewall refusing the connection to this IP address.

CE – Connection time-out

This result indicates that the Agent is not receiving any reply packets when it attempts to initiate a TCP connection to the Command Engine on port 1018. The likely cause is that the Agent is connecting to the "wrong" IP address. In other words, the Agent doesn't know the correct IP address of the Command Engine. A firewall might also be blocking the connection.

What Can I Do If the Connection Times Out During a CE Test?

Follow the same steps specified in What Can I Do If the Connection is Refused During a CE Test?.

CE – DNS does not resolve

This result indicates that the Agent cannot resolve the host name “way” to a valid IP address. In other words, the Agent does not know the correct IP address of the Command Engine.

What Can I Do If the Command Engine Name Does Not Resolve During a CE Test?

Log into the server and use a command such as Telnet to confirm that the host name “way” can resolve (for example: “telnet way 1018”). If not, check the DNS configuration of the server to make sure that the host name “way” is configured to its correct IP address. See “Resolving Host Name” on page 771 in this appendix for more information.

CE – Old Agent version

This result indicates that the Agent was unable to contact the Command Engine, but the test was unable to determine the exact cause because the Agent is out of date.

What Can I Do If the Agent is Out of Date During a CE Test?

If this error occurs, it will likely be for one of two reasons: either the host name of the Command Engine (“way”) did not resolve, or the connection was refused.

- If you believe that the host name of the Command Engine (“way”) did not resolve, then See “CE – DNS does not resolve” on page 753 in this appendix for more information.
- If you determine that the connection was refused, See “CE – Connection refused” on page 752 in this appendix for more information.

Alternatively, you can upgrade the Agent to the latest version (contact Opsware, Inc. Customer Support) and re-run the test. See “Agent Reachability Communication Tests” on page 204 in Chapter 6 for information about how to install an Agent.

CE – Realm is unreachable

The Satellite realm where the managed server is located is unreachable. This error means that a path of tunnels between the Gateways in the Opsware core and the realm of the managed server cannot be established.

What Can I Do if the Realm is Unreachable During a CE Test?

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact your Opsware administrator for assistance in troubleshooting the Gateway network.

CE – No Gateway defined

The managed server is in a Satellite realm, but its Agent is not properly configured to use a Gateway. Agents located in satellites must use a Gateway to contact the core.

What Can I Do If No Gateway is Defined During a CE Test?

To troubleshoot this error, try the following:

1. Create or open the opswgw.args file on the managed server. The opswgw.args file is located on the managed server at:

   - **Unix/Linux**: /var/lc/cogbot/etc
   - **Windows**: %SystemDrive%\Program Files\Common Files\Loudcloud\cogbot\etc

2. Make sure that this file contains a single line as shown:

   ```
   opswgw.gw_list: <gw_ip_address>:<gw_port>,<gw_up_
   address>:<gw_port>
   ```

CE – Tunnel setup error

The Command Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do If A Tunnel Setup Occurs Error During a CE Test?

Contact your Opsware administrator.

CE – Gateway denied access

The Gateway is working, but refused to proxy the connection on behalf of the Agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the Agent to access the Command Engine.

What Can I Do if the Gateway is Denied Access During a CE Test?

Contact your Opsware administrator.

CE – Gateway name resolution error

The server running the Gateway in the Opsware core was unable to resolve the host name "way". It must be able to do this in order to proxy connections on behalf of managed servers in Satellite realms.

What Can I Do if a Name Resolution Error Occurs on the Gateway During a CE Test?

Log into the server where the core Gateway is located and use a command such as ping or host to confirm that the host name “way” can be resolved (for example: “host way”).

If you cannot connect, contact your Opsware administrator so that you can check the DNS configuration of the core Gateway server.

CE – Internal Gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do if an Internal Gateway Error Occurs During a CE Test?

Contact your Opsware administrator.

CE – Gateway could not connect to server

The Gateway could not establish a connection to the Command Engine. The situation might be because the Command Engine is not running, or because the Gateway is resolving the Command Engine host name (“way”) to the wrong IP address. It is also possible that a firewall might be blocking the connection.

What Can I Do if the Gateway Can’t Connect to Server During a CE Test?

Check that the name “way” resolves to the correct IP address and that the Gateway can establish a connection to port 1018 at that IP. See “Resolving Host Name” on page 771 and “Verifying that a Port is Open on a Managed Server” on page 769 in this appendix.

CE – Gateway time-out

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do if the Gateway Times Out During a CE Test?

Ensure that network connectivity is available between the Gateways in the path between the realm of the managed server and the Opsware core.

CE – No callback from Agent

The Command Engine was able to contact the Agent, but the Agent did not call back to retrieve its command. However, the Agent reports that it can connect to a Command Engine. This most likely means that the managed server's name resolution mechanism (such as DNS) is not configured to point the server to a different Opsware core facility than is currently stored for the server by Opsware SAS.

What Can I Do if There is No Callback from Agent?

It is possible that the MID file is missing. If the MID file is missing, it can be recreated easily by creating a file called 'mid' in the correct location which contains the value of the “Server ID” from the Server’s Properties page in the OCC.

If the MID file is not missing, then ensure that the server's name resolution mechanism (DNS, NIS, and so on) is properly configured so that the host names “spin” and “way” resolve to the appropriate IP addresses or Opsware core services in the same core, and that those hosts can be reached from the server on ports 1004 and 1018, respectively. If this is the case, it is likely that the server has been redirected to a different core recently, and the Agent has not yet registered, which will cause Opsware SAS to update its records. It this issue remains unresolved for more than 12 hours, contact Opsware, Inc. Customer Support.

Agent to Data Access Engine (DAE)

This test checks that the Agent can retrieve its device record from Data Access Engine. The fifteen possible results are:

- DAE – OK
- DAE – Untested
- DAE – Unexpected error
- DAE – Connection refused
- DAE – Connection time-out
- DAE – DNS does not resolve
- DAE – Old Agent version
- DAE – Realm is unreachable
- DAE – No Gateway defined
- DAE – Tunnel setup error

- DAE – Gateway denied access
- DAE – Gateway name resolution error
- DAE – Internal Gateway error
- DAE – Gateway could not connect to server
- DAE – Gateway time-out

DAE – OK

No troubleshooting necessary.

DAE – Untested

This result is returned when a functional area cannot be tested, because of a previous failure that prevents further testing. For example, if the Agent cannot reach the Data Access Engine then no other tests are possible.

What Can I Do If a Test Is Not Run During a DAE Test?

First resolve all tests that failed, and then run the Communication Test again.

DAE – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During a DAE Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware, Inc. Customer Support.

DAE – Connection refused

This result indicates that the Agent is receiving a TCP reset packet when attempting to connect to the Data Access Engine on port 1004. The likely cause is that the Agent is connecting to the wrong IP address. A firewall might also be blocking the connection.

What Can I Do If the Connection is Refused During a DAE Test?

1. Check that the name “spin” resolves to its correct IP address. See “Resolving Host Name” on page 771 in this appendix for more information.
2. Check to make sure that a firewall is not refusing the connection to this IP address.

DAE – Connection time-out

This result indicates that the Agent is not receiving any reply packets when it attempts to initiate a TCP connection to the Data Access Engine on port 1004. The likely cause is that the Agent is connecting to the wrong IP address. In other words, the Agent does not know the correct IP address of the Command Engine. A firewall might also be blocking the connection.

What Can I Do If the Connection Times Out During a DAE Test?

Follow the same steps specified in "What Can I Do If the Connection is Refused During a DAE Test?" on page 757.

DAE – DNS does not resolve

This result indicates that the Agent cannot resolve the host name "spin" to a valid IP address. In other words, the Agent does not know the correct IP address of the Data Access Engine.

What Can I Do If the Data Access Engine Name Does Not Resolve During a DAE Test?

Log into the server and use a command such as Telnet to confirm that the host name "spin" can be resolved (for example: telnet spin 1004"). If not, check the DNS configuration of the server to make sure that the host name "spin" is configured to its correct IP address. See "Resolving Host Name" on page 771 in this appendix for more information.

DAE – Old Agent version

This result indicates that the Agent was unable to contact the Data Access Engine, and the test is unable to determine the exact cause, because the Agent is out of date.

What Can I Do If the Agent is Out of Date During an DAE Test?

If this error occurs, it will likely be for one of two reasons: either the host name of the Data Access Engine ("spin") did not resolve, or the connection was refused.

- If you believe that the host name of the Data Access Engine ("way") did not resolve, then see "DAE – DNS does not resolve" on page 758.
- If you determine that the connection was refused, see "DAE – Connection refused" on page 757.

Alternatively, you can upgrade the Agent to the latest version (contact Opsware, Inc. Customer Support) and re-run the test. See “Agent Reachability Communication Tests” on page 204 in Chapter 6 for information about how to install an Agent.

DAE – Realm is unreachable

The Satellite realm where the managed server is located is unreachable. This error means that a path of tunnels between the gateways in the Opsware core and the realm of the managed server cannot be established.

What Can I Do if the Realm is Unreachable During a DAE *Test?*

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact your Opsware administrator for assistance in troubleshooting the Gateway network

DAE – No Gateway defined

The managed server is in a Satellite realm, but its Agent is not properly configured to use a Gateway. Agents located in satellites must use a Gateway to contact the core.

What Can I Do If No Gateway is Defined During a DAE *Test?*

To troubleshoot this error, try the following:

1. Create or open the opswgw.args file on the managed server. The opswgw.args file is located on the managed server at:

   - **Unix/Linux**: /var/lc/cogbot/etc
   - **Windows**: %SystemDrive%\Program Files\Common Files\Loudcloud\cogbot\etc

2. Make sure this file contains a single line as shown:

   ```
   opswgw.gw_list: <gw_ip_address>:<gw_port>,<gw_up_
   address>:<gw_port>
   ```

DAE – Tunnel setup error

The Data Access Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do if a Tunnel Setup Error Occurs During a DAE *Test?*

Contact your Opsware administrator.

DAE – Gateway denied access

The Gateway is working, but refused to proxy the connection on behalf of the Agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the Agent to access the Data Access Engine.

What Can I Do if the Gateway is Denied Access During a DAE *Test?*

Contact your Opsware administrator.

DAE – Gateway name resolution error

The server running the Gateway in the Opsware core was unable to resolve the host name "spin". It must be able to do this in order to proxy connections on behalf of managed servers in Satellite realms.

What Can I Do if There is a Name Resolution Error on the Gateway During a DAE *Test?*

Log into the server where the core Gateway is located and use a command such as ping or host to confirm that the host name "spin" can be resolved (for example: "host spin").

If you cannot connect, contact your Opsware administrator so that you can check the DNS configuration of the core Gateway server.

DAE – Internal Gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do if an Internal Gateway Error Occurs During a DAE *Test?*

Contact your Opsware administrator.

DAE – Gateway could not connect to server

The Gateway could not establish a connection to the Data Access Engine. This might be because the Data Access Engine is not running, or because the Gateway is resolving the Data Access Engine host name ("spin") to the wrong IP address. It is also possible that a firewall might be blocking the connection.

What Can I Do if the Gateway Can't Connect to Server During a DAE *Test?*

Check that the name "spin" resolves to the correct IP address and that the Gateway can establish a connection to port 1018 at that IP. See "Resolving Host Name" on page 771 in this appendix for more information and See "Verifying that a Port is Open on a Managed Server" on page 769 in this appendix for more information.

DAE – Gateway time-out

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do if the Gateway Times Out During a DAE Test?

Ensure that network connectivity is available between the Gateways in the path between the managed server's realm and the Opsware core.

Agent to Software Repository (SWR)

This test checks that the Agent can establish an SSL connection to the Software Repository.

There 16 possible results are:

- SWR – OK
- SWR – Untested
- SWR – Unexpected error
- SWR – Connection refused
- SWR – Connection time-out
- SWR – DNS does not resolve
- SWR – Old Agent version
- SWR - Server identification error
- SWR – Realm is unreachable
- SWR – No Gateway defined
- SWR – Tunnel setup error
- SWR – Gateway denied access

- SWR – Gateway name resolution error
- SWR – Internal Gateway error
- SWR – Gateway Could not connect to server
- SWR – Gateway time-out

SWR – OK

No troubleshooting necessary.

SWR – Untested

This result is returned when a functional area cannot be tested because of a previous failure that prevents further testing. For example, if the Agent cannot reach the Software Repository, then no other tests are possible.

What Can I Do If a Test Is Not Run During a SWR Test?

First resolve all tests that failed, and then run the Communication Test again.

SWR – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During a SWR Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware, Inc. Customer Customer Support.

SWR – Connection refused

This result indicates that the Agent is receiving a TCP reset packet when attempting to connect to the Software Repository on port 1003. The likely cause is that the Agent is trying to connect to the wrong IP address. A firewall might also be blocking the connection.

What Can I Do If the Connection is Refused During an SWR Test?

1. Check that the name "theword" resolves to the correct IP address. For this information, see "Resolving Host Name" on page 771.
2. Check to make sure that a firewall isn't refusing the connection to this IP address.

SWR – Connection time-out

This result indicates that the Agent is receiving a TCP reset packet when attempting to connect to the Software Repository on port 1003. The likely cause is that the Agent is connecting to the wrong IP address. In other words, the Agent does not know the correct IP address of the Software Repository. A firewall might also be blocking the connection.

What Can I Do If the Connection Times Out During an SWR Test?

Follow the same steps specified in "What Can I Do If the Connection is Refused During an SWR Test?" on page 762.

SWR – DNS does not resolve

This result indicates that the Agent cannot resolve the host name "theword" to a valid IP address. In other words, the Agent does not know the correct IP address of the Software Repository.

What Can I Do If the Software Repository Name ("theword") Does Not Resolve During an SWR Test?

Log into the server and use a command such as Telnet to confirm that the host name "theword" can be resolved (for example: "telnet theword 1003"). If not, contact your Opsware administrator so that you can check the DNS configuration of the server.

SWR – Old Agent version

This result indicates that the Agent was unable to contact the Software Repository, and the test is unable to determine the exact cause because the Agent is out of date.

What Can I Do If the Agent is Out of Date During an SWR Test?

If this error occurs, it will likely be for one of two reasons: either the host name of the Software Repository ("theword") did not resolve, or the connection was refused.

- If you think that the host name of the Software Repository ("theword") did not resolve, then see "SWR – DNS does not resolve" on page 763.
- If you determine that the connection was refused, see "SWR – Connection refused" on page 762.

Alternatively, you can upgrade the Agent to the latest version (contact Opsware, Inc. Customer Support) and re-run the test. For information on how to install an Opsware Agent, refer to Chapter 2 of the Opsware® SAS 5 User's Guide, in the Server Assimilation section.

SWR - Server identification error

Whenever an Agent makes a request of the Software Repository, the identity of the server is validated to confirm that the server should be allowed access to the information requested. This error indicates that the Software Repository was unable to identify the server being tested, or incorrectly identified that server.

What Can I Do If I Get a Server Identification Error?

The Software Repository identifies servers based on the incoming IP address of the request. To troubleshoot this error, try the following:

1. Check the Device Properties tab for the server in the Opsware Command Center to see if Network Address Translation (NAT) is in use. If it is, make sure that NAT is statically configured, and that only one server is using the NAT address. If multiple servers are using the same IP address, you will need to reconfigure the NAT device. For more information, please refer to Chapter 2 of the Opsware® SAS 5 User's Guide.

2. If the Agent is installed on a cluster, check that each node in the cluster has a unique IP address at which it can be reached. You might have to add static routes to the server to ensure that connections made from that server to the Opsware core use the unique IP. If NAT is not in use, you can alternately mark the correct interface as the "primary" interface through the Network Configuration tab for the server in the Opsware Command Center. For more information, please refer to Chapter 2 of the Opsware® SAS 5 User's Guide.

3. The server's IP address might have changed recently. If this is the case, stop and restart the Agent. For instructions on how to stop and start an Agent, see "Restarting an Opsware Agent" on page 770.

SWR – Realm is unreachable

The Satellite realm where the managed server is located is unreachable. This error means that a path of tunnels between the gateways in the Opsware core and the realm of the managed server cannot be established.

What Can I Do if the Realm is Unreachable During a SWR Test?

This error could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration. Contact your Opsware administrator for assistance in troubleshooting the Gateway network.

SWR – No Gateway defined

The managed server is in a Satellite realm, but its Agent is not properly configured to use a Gateway. Agents located in satellites must use a Gateway to contact the core.

What Can I Do If No Gateway is Defined During a SWR Test?

To troubleshoot this error, try the following:

1. Create or open the opswgw.args file on the managed server. The opswgw.args file is located on the managed server at:

   - **Unix/Linux**: /var/lc/cogbot/etc

   - **Windows**: %SystemDrive%\Program Files\Common Files\Loudcloud\cogbot\etc

2. Make sure that this file contains a single line as shown:

   ```
   opswgw.gw_list: <gw_ip_address>:<gw_port>,<gw_up_
   address>:<gw_port>
   ```

SWR – Tunnel setup error

The Data Access Engine could not establish a connection through any of its defined Gateways. This could be due to a network problem, a malfunctioning or failed Gateway, or a Gateway misconfiguration.

What Can I Do If a Tunnel Setup Error Occurs During a SWR Test?

Contact your Opsware administrator.

SWR – Gateway denied access

The Gateway is working but refused to proxy the connection on behalf of the Agent. This error most likely means that the Gateway is misconfigured such that the Gateway will not allow the Agent to access the Software Repository.

What Can I Do if the Gateway is Denied Access During a SWR Test?

Contact your Opsware administrator.

SWR – Gateway name resolution error

The server running the Gateway in the Opsware core was unable to resolve the host name “theword”. It must be able to do this in order to proxy connections on behalf of managed servers in Satellite realms.

What Can I Do if a Name Resolution Error Occurs on the Gateway During a SWR *Test?*

Log into the server where the core Gateway is located and use a command such as ping or host to confirm that the host name “theword” can be resolved (for example: “host theword”).

If you cannot connect, contact your Opsware administrator so that you can check the DNS configuration of the core Gateway server.

SWR – Internal Gateway error

Due to an internal error, the Gateway was unable to proxy the connection. This typically occurs when the Gateway is overloaded.

What Can I Do if an Internal Gateway Error Occurs During a SWR *Test?*

Contact your Opsware administrator.

SWR – Gateway Could not connect to server

The Gateway couldn't establish a connection to the Software Repository. This error might be because the Software Repository is not running, or because the Gateway is resolving the Software Repository host name (“theword”) to the wrong IP address. It is also possible that a firewall might be blocking the connection.

What Can I Do if the Gateway Can’t Connect to Server During a SWR *Test?*

Check that the name “theword” resolves to the correct IP address and that the Gateway can establish a connection to port 1018 at that IP address. For more information, see “Resolving Host Name” on page 771 and “Verifying that a Port is Open on a Managed Server” on page 769.

SWR – Gateway time-out

The Gateways on the two ends of a tunnel could not communicate with each other, most likely due to a network connectivity problem.

What Can I Do if the Gateway Times Out During a SWR Test?

Ensure that network connectivity is available between the Gateways in the path between the realm of the managed server and the Opsware core.

Machine ID Match (MID)

This test checks whether the MID that the Agent reported matches that recorded in the Model Repository (Opsware data repository).

You can receive four possible errors from the Machine ID (MID) Communication Test:

- MID – OK
- MID – Untested
- MID – Unexpected error
- MID – MID mismatch

MID – OK

No troubleshooting necessary.

MID – Untested

This result is returned when a functional area cannot be tested, because of a previous failure that prevents further testing. For example, if the Agent cannot reach the Model Repository, then no other tests are possible.

What Can I Do If a Test Is Not Run During an MID Test?

First resolve all tests that failed, and then run the Communication Test again.

MID – Unexpected error

This result indicates that the test encountered an unexpected condition.

What Can I Do If I Get an Unexpected Error During an MID Test?

First resolve all tests that failed, and then run the Communication Test again. If the unexpected error recurs, check to see if any additional details in the error message indicate the problem. If the error cannot be resolved, contact Opsware, Inc. Customer Support.

MID – MID mismatch

This result indicates that the MID that the Agent reported does not match the recorded MID in the Model Repository for that Agent. The likely cause is that the Command Engine is running the test against the wrong Agent.

What Can I Do If the MID is Mismatched During an MID Test?

To troubleshoot this error, try the following:

1. Check the Device Properties tab for the server in the Opsware Command Center to see if NAT is in use for this server. If it is, make sure that static, 1-to-1 NAT is being used. Opsware SAS requires that all managed servers be reachable on a distinct, consistent IP address, so configurations that assign addresses dynamically or use port-based translation are not supported.
2. If the Agent is installed on a cluster, check that each node in the cluster has a unique IP address at which it can be reached. You might have to add static routes to the server to ensure that connections made from that server to the Opsware core use the unique IP. If NAT is not in use, you can alternately mark the correct interface as the "primary" interface via the Network Configuration tab for the server in the Opsware Command Center.
3. The IP address might have changed recently. If this is the case, stop and restart the Agent. For these instructions, see "Restarting an Opsware Agent" on page 770.

Common Troubleshooting Tasks

The following list of troubleshooting tasks are common to more than one Communication Test error:

- Verifying that an Agent is Running
- Verifying that a Port is Open on a Managed Server
- Restarting an Opsware Agent
- Checking Management IP of a Managed Server
- Checking Network Gateway Configuration
- Resolving Host Name

Verifying that an Agent is Running

To verify that an Agent is running on a server, perform the following steps:

1. On Solaris, HP-UX, or AIX, enter this argument at the command line:

```
/usr/ucb/ps auxwww | grep cog
```

You should get this result if the Agent is running:

```
/opt/OPSW/bin/python /opt/OPSW/blackshadow/shadowbot/
daemonbot.pyc --conf /var/lc/cogbot/etc/cogbot.args
```

2. On Linux, enter this argument at the command line:

   ```
   ps auxwww | grep cog
   ```

   You should get this result if the Agent is running:

   ```
   /opt/OPSW/bin/python /opt/OPSW/blackshadow/shadowbot/
   daemonbot.pyc --conf /var/lc/cogbot/etc/cogbot.args
   ```

3. On Windows, from the Administrative Tools | Services, check to make sure that the 'shadowbot' service is running.

Verifying that a Port is Open on a Managed Server

For some errors, you will need to verify that the port is open on the server where the Agent is installed. To do this, perform the following steps:

1. Check if the port is open.

2. On Solaris, HP-UX, AIX, or Linux enter:

   ```
   'netstat -an | grep 1002 | grep LISTEN'
   ```

   If the port is open on the box, you should get back the following:

   ```
   *.1002    *.*    0      0 24576      0 LISTEN
   ```

3. On Windows, at the command prompt enter:

   ```
   'netstat -an | find "1002" | find "LISTEN"':\
   ```

   If the port is open on the box, you should get back the following result:

   ```
   TCP0.0.0.0:10020.0.0.0:0LISTENING
   ```

4. Confirm that the port is actually open. To do this, from the computer where the Agent is installed, Telnet to port 1002 by using both localhost and the external IP address of the server. Performing the Telnet will help you confirm that a connection refused message is being caused by the lack of an open port on the managed server rather than a problem with networking hardware between the core and the managed server.

Restarting an Opsware Agent

For Solaris, Linux, or AIX, perform the following steps:

1. To stop an Opsware Agent on Solaris, Linux, or AIX, execute the following command:

   ```
   /etc/init.d/cogbot stop
   ```

2. To restart the Opsware Agent on Solaris, Linux, or AIX, execute the following command:

   ```
   /etc/init.d/cogbot start
   ```

For HP-UX, perform the following steps:

1. To stop an Opsware Agent on HP-UX, execute the following command:

   ```
   /sbin/init.d/cogbot stop
   ```

2. To restart an Opsware Agent on HP-UX, execute the following command:

   ```
   /sbin/init.d/cogbot start
   ```

For Windows, execute the following command to stop and start an Opsware Agent:

```
net stop shadowbot
```

```
net start shadowbot
```

Checking Management IP of a Managed Server

To check the Management IP of a managed server, perform the following steps:

1. To view the management IP of the managed server, log into the Opsware Command Center.
2. From the Navigation panel, click Servers ➤ Manage Servers.
3. From the Manage Servers list, click the display name of the server for which you want to check the Management IP.
4. Select the Network tab of the server’s properties.
5. Check to make sure that the Management IP address matches the IP address of the managed server.

Checking Network Gateway Configuration

To check the network Gateway configuration, perform the following steps:

1. On Solaris, enter this command to check routing table:

```
netstat -rn
```

Your results should look like this:

```
default              192.168.8.1         UG        1   5904
```

where `192.168.8.1` is the IP of the Gateway.

2. On Linux, enter this command to check routing table:

   ```
   route -n
   ```

   Your results should look like this:

   ```
   0.0.0.0       192.168.8.1    0.0.0.0       UG   0     0
   0 eth0
   ```

   where `192.168.8.1` is the IP of the Gateway.

3. On Windows, enter this command to check routing table:

   ```
   route print
   ```

   Your results should look something like this:

   ```
   0.0.0.00.0.0.0192.168.8.1192.168.8.12020
   ```

   where `192.168.8.1` is the IP of the Gateway.

4. In each case, you should also ping 192.168.8.1 (IP) to confirm that you can actually reach the Gateway.

Resolving Host Name

All managed servers (those with agents) must be able to resolve unqualified Opsware SAS service names for the following components:

- spin (Data Access Engine)
- way (Command Engine)
- theword (Software Repository)

If you need to ensure that one of these host names resolves correctly, contact your Opsware administrator to find out what qualified host name or IP address these service names should resolve to.

1. Try to ping the host. For example, execute the following command if you wanted to resolve the host name, way:

```
ping way
```

2. If the host name cannot resolve, you will get the following errors:

   Linux/Solaris/AIX/HP-UX:

   ```
   ping: unknown host way
   ```

   Windows:

   ```
   Ping request could not find host way. Please check the name
   and try again.
   ```

3. If the host name can resolve, you might get back various permutations of these types of messages (OS independent):

   ```
   way is alive
   ```

   Or

   ```
   pinging way (ip) with 32 bytes of data
   ```

Appendix D: Content Pack

IN THIS APPENDIX

This section discusses the following topics:

- Overview of the Content Pack
- Application Configurations
- Selection Criteria for Snapshots
- Global Shell Scripts

Overview of the Content Pack

This section describes content for Opsware SAS application configurations, selection criteria for snapshots, and Global Shell scripts.

Application Configurations

The following list details the types of Opsware Application Configuration Management (ACM) configurations provided in this release:

- **System Application Configurations**: Application Configurations that manage operating system configuration files.
- **Service Application Configurations**: Application Configuration that manage service configuration files.
- **“Application” application configurations**: Application Configurations that manage application configuration files.

Application configurations are designed to help you manage typical configuration files associated with applications or operating systems. Using ACM inside the OCC Client enables you to edit values in a configuration file and push those changes to the application on managed servers.

Each Application Configuration takes the name of the original configuration file it manages and uses the TPL file extension (though you can use any file extension you wish). Some application configurations consist of a single configuration template, while other application configurations contain several Application Configuration in addition to pre-install, post-install scripts, and post error scripts.

See "Application Configuration Management" on page 599 in Chapter 16 for more information on how to create an Application Configuration and use the Application Configurations described here.

For more information on how to install and access these Application Configurations, please refer to the Opsware SAS Content Pack product DVD.

Application Configuration Description Key

The following key describes the categories used in to explain the Application Configurations:

- **Application Configuration name**: Gives the Application Configuration name.
- **Platforms supported**: Lists all supported operation systems/platforms that the Application Configuration can run on.
- **Functionality**: Explains what the Application Configuration does when it is pushed to a server.
- **Pre-install/Post-install/post error scripts**: Explains what (if any) pre-install, post-install, and/or errors scripts are included with Application Configuration and what they are used for.
- **Files created**: Lists files (if any) created by the pre-install/post-install/post error scripts associated with the Application Configuration when the application configuration is pushed to the managed server.
- **Files modified**: Lists what files are modified on the managed server when the application configuration is pushed to the managed server.
- **Notes**: Describes any special considerations or instructions regarding the application configuration.
- **Limitations**: Lists any known issues or limitations to application configuration.

At this time, Application Configurations will not start services that are not already running. In the event you wish to configure a Unix or Linux service that is not already running on a system, please start the service before using Application Configuration or you may get an error from the AppConfig post-script execution. This error can be ignored, as the configuration has in fact been pushed to the server, but the service has not been started.

System Application Configurations

This section describes the following system Application Configurations:

- cron_allow.tpl
- crontab.tpl
- fstab.tpl
- group.tpl
- hosts.tpl
- hosts_allow.tpl
- hosts_equiv.tpl
- lmhosts.tpl
- passwd.tpl
- resolve.tpl
- shadow.tpl
- vfstab.tpl

cron_allow.tpl

The cron_allow.tpl application configuration allows you to manage the cron.allow file. Table D-1 explains the usage of the application configuration.

Table D-1: cron_allow.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Unix and Linux platforms.
Functionality	Allows you to manage the /etc/cron.allow file, enabling you to add or subtract users.

Table D-1: cron_allow.tpl application configuration (continued)

CATEGORY	DESCRIPTION
Pre-install/Post-install/Post error script	post-cron_allow.sh Runs to send a SIGHUP to cron[d] after a change is made so changes are put into effect.
Files created	Pre-install/Post-install/Post error script
Files modified	/etc/cron.allow
Notes	None

crontab.tpl

The crontab.tpl allows you to manage the crontab file. Table D-2 explains the usage of the application configuration.

Table D-2: crontab.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Unix and Linux platforms.
Functionality	Allows you to edit user and/or system crontab files. See the man page(s) for crontab on the specific system(s) of interest, to locate the system and/or user crontab files. On Solaris for example, these are located in the /var/spool/cron/crontabs directory with each crontab file bearing the username owning the cron, in other words, /var/spool/cron/crontabs/root. Under Linux, the main crontab file is in /etc/crontab. Additionally, other files such as /etc/cron.d/* on Linux may also be parsed by this template. The fields, which are documented by the man pages for cron and crontab, are as follows: Minute Hour Day Month Weekday Command Additionally, the MAILTO, HOME, SHELL, and PATH environment variables may also be set in this file. See the man page for further details.

Table D-2: crontab.tpl application configuration (continued)

CATEGORY	DESCRIPTION
Pre-install/Post-install/Post error script	post-crontab.sh Sends a SIGHUP signal to the cron daemon to force reloading of the configuration files/crontab entries.
Files created	None
Files modified	/etc/cron.d/crontab /var/spool/cron/crontabs/* (Depends on platform. Refer to the specific operating system's man page for crontab.)
Notes	None
Limitations	A limitation exists that disallows the embedded DOS Carriage Return character in the command to be executed. This should not pose a problem for any command with the possible exception of cronned jobs that are doing DOS-to-Unix-like translations via cron. This can be achieved by using the `dos2unix` or similar command as an alternative, if this must be done via cron on a system using an application configuration.

fstab.tpl

The fstab.tpl application configuration allows you to manage the fstab file. Table D-3 explains the usage of the application configuration.

Table D-3: fstab.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Linux and HP-UX platforms
Functionality	Allows configuration of at boot (and noauto) file system mounts.
Pre-install/Post-install/Post error scripts	post-fstab.sh should be executed, which will mount any new file systems that have been added to /etc/fstab that do not have the 'noauto' flag set.
Files created	None

Table D-3: fstab.tpl application configuration (continued)

CATEGORY	DESCRIPTION
Files modified	fstab
Notes	In the event an administrator removes an existing mount point or file system from the fstab file, the post script will not be able to unmount these file systems at this time, so these file systems should be manually unmounted as needed.

group.tpl

The group.tpl application configuration allows you to manage the group file. Table D-4 explains the usage of the application configuration.

Table D-4: group.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Unix and Linux platforms.
Functionality	Allows you to manage user groups, such as adding or removing user groups.
Pre-install/Post-install/Post error scripts	None
Files created	None
Files modified	/etc/group
Notes	The encrypted password feature is not honored by this template for /etc/group.

hosts.tpl

The application configuration allows you to manage the hosts file. Table D-5 explains the usage of the application configuration.

Table D-5: hosts.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Unix and Linux platforms.

Table D-5: hosts.tpl application configuration (continued)

CATEGORY	DESCRIPTION
Functionality	Allows you to manage the static table for host names, enabling you to add or remove entries from the /etc/hosts file.
Pre-install/Post-install/Post error scripts	None
Files created	None
Files modified	/etc/hosts
Notes	None

Limitations	The current application configuration does not handle IPv6 addresses, such as those seen by default on some versions of SuSE or SLES, similar to the following:
	`# special IPv6 addresses`
	`::1            localhost ipv6-localhost`
	`ipv6-loopback`
	`fe00::0         ipv6-localnet`
	These style of entries (with IPv6 addresses) are not currently able to be parsed, or thus manipulated via the Application Configuration.
	Some systems may declare some default IPv6 addresses. If IPv6 is not being used, these may be commented out, after which ACM will properly handle the resulting file.

hosts_allow.tpl

The hosts_allow.tpl application configuration allows you to manage the host.allow file. Table D-6 explains the usage of the application configuration.

Table D-6: hosts_allow.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Linux and Sun Solaris platforms

Table D-6: hosts_allow.tpl application configuration (continued)

CATEGORY	DESCRIPTION
Functionality	Allows you to manage the static table for host names, enabling you to add or remove entries from the LMHOSTS file.
Pre-install/Post-install/Post error scripts	None
Files created	None but needs to have a SIGHUP sent after changes.
Files modified	/etc/hosts.allow
Notes	None

hosts_equiv.tpl

The hosts_equiv.tpl application configuration allows you to manage the hosts.equiv file. Table D-7 explains the usage of the application configuration.

Table D-7: hosts_equiv.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Linux and Sun Solaris platforms
Functionality	Allows you to manage the hosts.equiv file.
Pre-install/Post-install/Post error script	None
Files created	None
Files modified	/etc/hosts.equiv
Notes	If running via inetd/xinetd, no changes needed. If running standalone, a HUP will need to be sent to the process.

lmhosts.tpl

The lmhosts.tpl application configuration allows you to manage the LMHOSTS file. Table D-8 explains the usage of the application configuration.

Table D-8: lmhosts.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Windows platforms
Functionality	Allows you to manage the LMHOSTS file
Pre-install/Post-install/Post error scripts	None
Files created	None
Files modified	LMHOSTS
Notes	None

passwd.tpl

The passwd.tpl application configuration allows you to manage the passwd file. Table D-9 explains the usage of the application configuration.

Table D-9: passwd.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Unix and Linux platforms.
Functionality	Allows you to manage user authentication information including userid, username, home directory, login shell, gid, password, status.
Pre-install/Post-install/Post error scripts	None
Files created	None
Files modified	/etc/passwd
Notes	None

Table D-9: passwd.tpl application configuration (continued)

CATEGORY	DESCRIPTION
Limitations	passwd.tpl allows a `'#'` character to be entered in the password field. If the value set containing `'#'` in the password is pushed onto disk, the subsequent import of the file will fail. Note that this should not cause problems in most cases, as most versions of Unix use a shadow password file. The current password template has an issue with systems that are not using a shadow password file, and will generate a validation error if an encrypted password is entered. It may still be used to import values, and to update other fields as long as an encrypted password is not entered or updated via the OCC Client.

resolve.tpl

The resolve.tpl application configuration allows you to manage the resolv.conf file. Table D-10 explains the usage of the application configuration.

Table D-10: resolve.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Unix and Linux platforms.
Functionality	Allows you to manage DNS Resolver client information, enabling you to add, remove entries in /etc/resolv.conf
Pre-install/Post-install/Post error scripts	None
Files created	None
Files modified	/etc/resolv.conf
Notes	None

shadow.tpl

The shadow.tpl application configuration allows you to manage the shadow file. Table D-11 explains the usage of the application configuration.

Table D-11: shadow.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Unix and Linux platforms.
Functionality	Allows you to manage encrypted passwords, enabling you to add or remove entries in the /etc/shadow file.
Pre-install/Post-install/Post error scripts	None
Files created	None
Files modified	/etc/shadow and /etc/passwd
Notes	None

vfstab.tpl

The vfstab.tpl application configuration allows you to manage the vfstab file. Table D-12 explains the usage of the application configuration.

Table D-12: vfstab.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Sun Solaris platforms
Functionality	Allows configuration of at boot (and noauto) file system mounts.
Pre-install/Post-install/Post error scripts	post-vfstab.sh should be executed, which will mount any new file systems that have been added to /etc/vfstab that do not have the noauto flag set.
Files created	None
Files modified	/etc/vfstab

Table D-12: vfstab.tpl application configuration (continued)

CATEGORY	DESCRIPTION
Notes	In the event an admin removes an existing mount point/filesystem from the vfstab file, the post script will not be able to unmount these file systems at this time, so these file systems should be manually unmounted as needed. Only rudimentary error checking is being done at this point, primarily checking if: • A permission denied error was encountered (NFS or user privileges). • An invalid mount point or device was specified. Either condition will return a generic error.

Service Application Configurations

This section describes the following service configuration Application Configurations:

- apache2_httpd_conf.tpl
- exports.tpl
- inetd_conf.tpl
- xinetd_conf.tpl

apache2_httpd_conf.tpl

The apache2_httpd_conf. application configuration allows you to manage the http.conf file. Table D-13 explains the usage of the application configuration.

Table D-13: apache2_httpd_conf.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Linux and Sun Solaris platforms.
Functionality	Allows you to manage Apache 2 Web Server configurations. See http://www.apache.org/docs-2.0 for information on the configuration file itself.

Table D-13: apache2_httpd_conf.tpl application configuration (continued)

CATEGORY	DESCRIPTION
Pre-install/Post-install/Post error scripts	post-apache2_httpd_conf.sh Sends a SIGHUP signal to the Apache httpd via the init script's reload option.
Files created	None
Files modified	/etc/httpd/conf/httpd.conf

Notes	Due to a current parser implementation, a limitation exists that limits the number of levels some directive blocks may be nested, such as:
	`<IfModule>`
	`<Directory>`
	IfModule blocks, including a nested Directory block, may only be nested 2 levels deep at this time.

exports.tpl

The exports.tpl application configuration allows you to manage the exports file. Table D-14 explains the usage of the application configuration.

Table D-14: exports.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Linux platforms
Functionality	Allows you to manage Linux NFS exports via the exports file.
Pre-install/Post-install/Post error scripts	The post-exports.sh post-configuration script forces a reload of the NFS exports file, and exports any new file systems changed by the configuration being applied.
Files created	None
Files modified	/etc/exports
Notes	None

Table D-14: exports.tpl application configuration (continued)

CATEGORY	DESCRIPTION

Limitations	You must specify a value for the netgroup or host in the exports file, due to a current limitation. Most sites supply this by default, but export entries of the format:
	`/pub          (ro, insecure, all_squash)`
	should be changed to:
	`/pub         *(ro, insecure, all_squash)`
	in order for ACM to be able to properly recognize these files; otherwise an error may occur during importing values.

inetd_conf.tpl

The inetd_conf.tpl application configuration allows you to manage the inetd.conf file. Table D-15 explains the usage of the application configuration.

Table D-15: inet_conf.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	Sun Solaris, Red Hat Linux 6.2
Functionality	Allows the configuration and addition of services to be handled by the Internet services daemon, enabling you to manage the /etc/inetd.conf file and configure allowable services for a specified system.
Pre-install/Post-install/Post error scripts	The post-script sends a SIGHUP to the inetd daemon to initiate a daemon configuration reload, allowing service changes to be available immediately.
Files created	None
Files modified	/etc/inetd.conf
Notes	None

xinetd_conf.tpl

The xinetd_conf.tpl application configuration allows you to manage the xinetd.conf file. Table D-16 explains the usage of the application configuration.

Table D-16: xinetd_conf.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	All supported Linux platforms except Red Hat Linux 7.1 and previous versions
Functionality	Allows the configuration and addition of services to be handled by the Extended Internet services daemon.
Pre-install/Post-install/Post error scripts	The post-configuration script will send a SIGHUP to the xinetd daemon to initiate a daemon configuration reload, allowing service changes to be available immediately.
Files created	None

Table D-16: xinetd_conf.tpl application configuration (continued)

CATEGORY	DESCRIPTION
Files modified	/etc/xinetd.d/* /etc/xinetd.conf
Notes	None

"Application" application configurations

This section describes the following application configurations:

- iislockd_ini.tpl
- urlscan_tpl

iislockd_ini.tpl

The iislockd_ini.tpl application configuration allows you to manage the iislocked.ini file. Table D-17 explains the usage of the application configuration.

Table D-17: iislocked.tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	Windows 2000
Functionality	Allows you to manage the iislockd.ini file by allowing you to add, modify or delete applications and their settings.
Pre-install/Post-install/Post error scripts	None
Files created	None
Files modified	iislockd.ini
Notes	None

urlscan_tpl

The urlscan_tpl application configuration allows you to manage the urlscan.ini file. Table D-18 explains the usage of the application configuration.

Table D-18: urlscan_tpl application configuration

CATEGORY	DESCRIPTION
Platforms Supported	Windows 2000
Functionality	Allows you to manage the urlscan.ini file by enabling you to change security settings within the file.
Pre-install/Post-install/Post error scripts	None
Files created	None
Files modified	urlscan.ini
Notes	None

Selection Criteria for Snapshots

In Opsware Server Compliance, snapshot selection criteria specify the information you want to capture about the state and configuration of a managed server or server group, at a particular point in time. A snapshot template identifies this selection criteria. See “Selection Criteria” on page 300 in Chapter 8 for more information.

Since selection criteria are specific to an operating system, such as Unix, Linux, or Windows, you cannot have one (universal) snapshot template for different operating systems of managed servers. Opsware SAS provides the following snapshot selection criteria for various operating systems:

- Linux-FilesystemRuleset
- Solaris-FilesystemRuleset
- Unix/Linux-Apache2Ruleset
- Unix/Linux-SunOne6.1RuleSet
- Windows 2000 SecurityRuleSet(WindowsSecureDC)
- Windows 2003 SecurityRuleSet(WindowsSecureDC)

See "Loading Snapshot Selection Criteria" on page 317 in Chapter 8 for information about how to load selection criteria into a snapshot template.

Linux-FilesystemRuleset

The following snapshot selection criteria records information about a Linux managed server:

```
/etc/init.d/cogbot
/etc/sysctl.conf
/etc/hosts.deny
/etc/hosts.allow
/etc/hosts
/etc/mtab
/etc/exports
/etc/fstab
/etc/nsswitch.conf
/etc/resolv.conf
/etc/xinetd.conf
```

By default, the `/etc/hosts.deny` and `/etc/hosts.allow` files are not automatically created on all versions of Red Hat Linux and SuSE® Enterprise operating systems during their original installations. These files are typically created by your system administrator to restrict access to a managed server from another managed server.

The `/etc/sysctl.conf` and `/etc/xinetd.conf` files are typically not found in a SuSE® Enterprise operating system. Therefore, if you want to use this selection criteria to record information about a managed server that has a SuSE® Enterprise operating system, remove these two files from the selection criteria before you create the snapshot. See "Editing a Snapshot Template" on page 318 in Chapter 8 for more information.

Solaris-FilesystemRuleset

The following snapshot selection criteria records information about a Solaris managed server:

```
/etc/hosts.allow
/etc/inet/hosts
/etc/init.d/cogbot
/etc/system
/etc/mnttab
/etc/vfstab
/etc/nsswitch.conf
/etc/dfs/dfstab
/etc/inet/inetd.conf
```

By default, the `/etc/hosts.allow` file is not included when a Sun Solaris operating system is installed. This file is typically created by your system administrator to restrict access to a managed server.

During a standard Sun Solaris operating system installation (whether it is performed with Opsware SAS or not), the `/etc/inet/hosts` and `/etc/inet/inetd.conf` files are created as symlinks.

Unix/Linux-Apache2Ruleset

The following snapshot selection criteria records information about a Unix/Linux-Apache2 managed server:

```
/etc/init.d/httpd
/opt/apache-2.0.49/conf/mime.types
/opt/apache-2.0.49/conf/ssl.conf
/opt/apache-2.0.49/conf/httpd.conf
```

This selection criteria requires the Apache 2 ISMtool installed and is provided as an Opsware DCML Exchange Tool (DET) export.

Unix/Linux-SunOne6.1RuleSet

The following snapshot selection criteria records information about a Unix/Linux-SunOne6.1 managed server:

```
/var/netscape/server4/https-admserv/config/admpw
/var/netscape/server4/https-web/config/mime.types
/var/netscape/server4/https-web/config/magnus.conf
/var/netscape/server4/https-web/config/obj.conf
```

Windows 2000 SecurityRuleSet(WindowsSecureDC)

The following snapshot selection criteria records security information about a Windows 2000 managed server:

```
machine\software\microsoft\driver signing
machine\software\microsoft\windows
nt\currentversion\setup\recoveryconsole
machine\software\microsoft\windows nt\currentversion\winlogon
machine\software\microsoft\windows\currentversion\policies\syst
em
machine\system\currentcontrolset\control\lsa
```

```
machine\system\currentcontrolset\control\print\providers\lanman
print services\servers
machine\system\currentcontrolset\services\lanmanserver\paramete
rs
machine\system\currentcontrolset\control\session manager
machine\system\currentcontrolset\services\netlogon\parameters
machine\system\currentcontrolset\Session Manager\SubSystems
machine\system\currentcontrolset\SecurePipeServers\winreg
```

Use HKEY_LOCAL_MACHINE to search for these registry keys.

This snapshot selection criteria is intended to capture Windows Registry entries; however, it may also record additional server object information, such as the GUID (Globally Unique Identifier) of the object. When you use this snapshot in an audit, you may see different values for this information. This is best explained by the following example:

Suppose you have created a snapshot of a Windows 2000 server where the securedc.inf security template has already been applied. Consider this to be a golden snapshot of a known working server–which is a managed server that is compliant (performs as expected) and is also referred to as a reference server or a baseline server.

Subsequently, you create another Windows 2000 server (using exactly the same process as the reference server), and then apply the securedc.inf security template (exactly the same way you applied it to the previous reference server).

In the audit process, when the golden snapshot (of the Windows 2000 reference server) is compared to the other (newly-minted) Windows 2000 server, you may expect to see no server differences in the audit results. However, you will indeed see differences that are primarily due to fixed strings which are server specific, such as the host name of the server domain in which it is registered, and so on.

Windows 2003 SecurityRuleSet(WindowsSecureDC)

The following snapshot selection criteria records security information about a Windows 2003 managed server:

```
machine\software\microsoft\driver signing
machine\software\microsoft\windows
nt\currentversion\setup\recoveryconsole
machine\software\microsoft\windows nt\currentversion\winlogon
```

```
machine\software\microsoft\windows\currentversion\policies\syst
em
machine\system\currentcontrolset\control\lsa
machine\system\currentcontrolset\services\lanmanserver\paramete
rs
machine\system\currentcontrolset\services\ldap
machine\system\currentcontrolset\control\session manager
machine\system\currentcontrolset\control\print\providers\lanman
print services\servers
machine\system\currentcontrolset\services\netlogon\parameters
machine\system\currentcontrolset\Session Manager\SubSystems
machine\system\currentcontrolset\SecurePipeServers\winreg
machine\software\policies\microsoft\windows\safer\codeidentifie
rs
```

Use HKEY_LOCAL_MACHINE to search for these registry keys.

This snapshot selection criteria is intended to capture Windows Registry entries; however, it may also record additional server object information, such as the GUID of the object. When you use this snapshot in an audit, you may see different values for this information. This is best explained by the following example:

> Suppose you have created a snapshot of a Windows 2003 server where the securedc.inf security template has already been applied. Consider this to be a golden snapshot of a known working server—which is a managed server that is compliant (performs as expected) and is also referred to as a reference server or a baseline server.
>
> Subsequently, you create another Windows 2003 server (using exactly the same process as the reference server), and then apply the securedc.inf security template (exactly the same way you applied it to the previous reference server).
>
> In the audit process, when the golden snapshot (of the Windows 2003 reference server) is compared to the other (newly-minted) Windows 2003 server, you may expect to see no server differences in the audit results. However, you will indeed see differences that are primarily due to fixed strings which are server specific, such as the host name of the server domain in which it is registered, and so on.

Global Shell Scripts

In Opsware Global Shell , you can execute shell scripts on managed servers that are running the following operating systems: AIX, HP-UX, Linux, Solaris, and Windows. These scripts are executed on servers that are in the Opsware Global File System (OGFS). See "Opsware Global File System (OGFS)" on page 389 in Chapter 10 for more information.

Opsware, Inc. recommends that you use these shell scripts to find the following information:

- Processes that are running in ACTIVE state on managed servers
- Processes that exceed the CPU threshold on managed servers
- Disk space that is available across all managed servers
- Disk space that is available across all managed servers and which is also less than a value that you specify
- Processes in the CPU on any managed server

Opsware SAS provides the following shell scripts:

```
FindByProcessName.sh
FindProcessMoreThanXCpu.sh
FindProcessMoreThanXCpu.sh
FindSpaceAvailable.sh
FindSpaceLessThanX.sh
FindTopNCpu.sh
```

If you are invoking a script from the top level of the `Server` directory (`/opsw/Server`), processing of a script that must recurse and run against multiple servers may require extra time.

FindByProcessName.sh

Arguments: At least one argument, which must be a string for a process name.

Returns: Results of a string match for a process that is in ACTIVE state on any managed server, depending on the location of the script invocation.

Usage: `FindByProcessName.sh java`

FindProcessMoreThanXCpu.sh

Arguments: At least one argument, which must be a positive integer.

Returns: All processes that exceed the CPU value provided on any managed server, depending on the location of the script invocation.

Usage: `FindProcessMoreThanXCpu.sh 10`

When this script is run on an HP-UX 11.11 server, it displays all processes instead of limiting the processes to ones that are based on CPU usage. This occurs because, in the HP-UX 11.11 operating system, the `ps -e pcpu,comm` command that is used does not return the CPU values as defined in `pcpu`. An operating system patch for `ps` may be required.

When this script is run on AIX servers, it displays the process name but not the CPU utilization percentage (%).

FindSpaceAvailable.sh

Arguments: No arguments.

Returns: All disk space that is available across all managed servers, depending on the location of the script invocation.

Usage: `FindSpaceAvailable.sh`

The results of this script may vary by ½ MB when it is run on the Red Hat Linux AS2.1 operating system.

The process of running this script may become suspended if the NFS server is not running or is not responding. This script can be updated to call `rosh` with the `-w` option that specifies the time-out value. See “Remote Opsware Shell (rosh)” on page 401 in Chapter 10 for more information.

FindSpaceLessThanX.sh

Arguments: At least one argument, which must be a positive integer. This value is treated as a percentage (%).

Returns: All disk space that is available across all managed servers and which is also less than X, depending on the location of the script invocation.

Usage: `FindSpaceLessThanX.sh 10`

This script does not support network drives on Windows operating systems.

FindTopNCpu.sh

Arguments: At least one argument, which must be a positive integer.

Returns: All N processes in the CPU on any managed server, depending on the location of the script invocation. If the number of processes found is less than N, all processes are returned.

Usage: `FindTopNCpu.sh 5`

This script returns incorrect results for servers that are running Windows operating systems. This script uses Microsoft® Windows Management Instrumentation (WMI), as suggested in the Microsoft Windows Resource Toolkit. (Please refer to pstop.vbs in the Microsoft Windows Resource Toolkit for the method calls.) Essentially, CPU is measured by iteratively going through each process in a process list, by looking at the Perfcounter of CPU, waiting for 1500ms, and then by looking at the Perfcounter again. The difference is then reported as a percentage of CPU usage. It is possible that you will find two processes that are using 90% of the CPU. This does not mean that the processes used that CPU at that one instance; it means that, during that period of time, the processes used CPU that was close to 99% of the total CPU available.

Appendix E: Glossary

IN THIS APPENDIX

This appendix describes terminology and acronyms used in Opsware SAS.

ACM *See* Application Configuration Management.

Ad-Hoc scripts A script that is created (or uploaded) and then immediately executed by a user. The script is intended for one-time use and is not stored in Opsware SAS.

administrator *See* Opsware administrator.

Agent *See* Opsware Agent.

Agent Installer An application that installs the Opsware Agent on a server.

Agent Uninstaller An application that uninstalls the Opsware Agent on a server.

application configuration Contains application configuration templates associated with an application.

Application Configuration Management (ACM) An Opsware feature that enables you to manage and modify configuration files for applications on managed servers.

application Provisioning See Software Provisioning.

audit A process that compares Opsware managed servers to determine how objects may differ. When an audit reveals a difference between servers, you can install software and server objects to remediate the discrepancy.

audit job The process that performed the audit.

audit template A definition of source, one or more targets, and selection criteria that will be examined during the audit process to compare servers, server groups, and existing snapshots.

Automated Configuration Tracking An Opsware feature that allows users to monitor critical configuration files and configuration databases. When Opsware SAS detects a change in a tracked configuration file or configuration database, the system can perform a number of actions, including backing up the configuration file or sending an email to a designated individual or group.

available patch A patch that the patch administrator has tested and marked as available. Only patches that have been marked as available can be installed by anyone other than a patch administrator. (The patch administrator can install an unavailable patch in order to test it.)

available server A reserve of new, unconfigured Opsware-enabled servers ready for quick deployment. The provisioned server can be moved into the Live environment to replace existing servers, add capacity, or support new applications. While optional, this provides faster recovery options in cases of hardware failure.

backup A feature in Automated Configuration Tracking that performs a backup of a file or database when it detects a change to a tracked configuration file or database. This action is performed only if the backup action is selected in the configuration tracking policy for the file or database.

backup (CDR) Process of saving the entire contents of the current Live directory for a specific service to the Backup directory. Code Deployment & Rollback (CDR) saves the backup copy to the local disk for the host on which the Backup operation was run. Only one backup copy is maintained at any time for a service.

backup event An event that causes configuration files or configuration databases to be backed up. Types of backup events include manual, full, and triggered.

blocked attachment An attachment that is not installed when that template is applied. The attachment also does not appear in child templates or folders.

Boot Server A part of the OS Provisioning feature that supports network booting of Sun and x86 systems with inetboot and PXE respectively. The processes used to provide this support include the Internet Software Consortium DHCP server, and Sun Solaris TFTP and NFS.

Build Manager A part of the OS Provisioning feature that facilitates communication between the OS Boot Agent and the Command Engine for OS provisioning.

CDR *See* Code Deployment & Rollback (CDR).

change log An audit trail of changes made to a node (read-only). Tracks changes made to a node. Identifies who has recently modified the node to add or remove software packages, add or remove operating systems, add or move servers, and create or remove subordinate nodes.

Code Deployment & Rollback (CDR) An Opsware feature used to push updated code and content to staging host servers.

Code Deployment Role A specific role that authorizes access to capabilities and functions with the Opsware Code Deployment & Rollback feature.

Command Engine The Opsware SAS component that enables distributed programs to run across many servers. The Command Engine handles the entry of scripts into the Opsware Model Repository (the script storage location in Opsware SAS) and the versioning of stored scripts. Command Engine scripts are written in Python and run on the Command Engine server.

Communication Test A feature that helps in identifying managed servers with unreachable Opsware Agents. A Communication Test lists all servers with unreachable agents, returns specific errors associated with each unreachable agent, and provides troubleshooting information to resolve the error. The Communication Test runs various tests like Command Engine to Agent Communication, Crypto Match, Agent to Command Engine Communication, Agent to Data Access Engine, Agent to Software Repository Communication, and Machine ID mismatch to determine if an Opsware Agent is reachable.

configuration template A set of values that represent the configuration file of an application.

configuration tracking policy The configuration tracking policy defines the set of files or configuration databases to be monitored, and the actions to be taken when change is detected to a tracked file.

configuration tracking reconcile Process by which new configuration tracking polices or changes to existing configuration tracking polices are deployed on servers.

core *See* Opsware core.

custom attributes Attributes such as miscellaneous parameters and named data values that users can set for servers in the Opsware Command Center. Used when performing a variety of Opsware functions, including network and server configuration, notifications, and CRON script configurations.

custom extension Custom Command Engine scripts that extend Opsware SAS functionality to customers to cover their specific needs.

customer An account within Opsware SAS that has access to designated resources, such as servers and software.

cutover A feature in CDR, that causes the Update directory and current Live directory to be identical. Performed automatically by determining the differences between the Update directory and current the Live directory. The files that are different are synchronized from the Update directory to the current Live directory.

Data Access Engine The XML-RPC interface to the Model Repository that simplifies interaction with various clients such as the Opsware Command Center, system data collection, and monitoring agents on servers.

data center Legacy term. *See* facility.

Data Center Intelligence Reporting An interface for mining the data that is contained in the Model Repository about all managed servers.

deactivated server Server removed from Opsware management even though its history still exists.

deployment Within CDR, automatically pushes code and content from a staging server to a live network server.

deprecated A possible state of a package or patch in Opsware SAS. A deprecated package or patch can no longer be installed on a managed server, but might still be installed on a server before the patch or package was deprecated.

device Legacy term. *See* server.

Distributed Scripts An Opsware feature that allows you to manage scripts in your managed environment.

Dormant Opsware Agent An Opsware Agent that runs in the dormant mode after its installation when Opsware SAS core is not available on the network. The dormant agent periodically attempts to contact the core and when the core is available, it performs the initialization tasks to complete its installation.

dynamic group A server group that contains servers added to or removed from the group based on a set of user-defined rules.

email notification list In the Automated Configuration Tracking feature, an email can be sent to the email addresses in the email notification list whenever a change to a tracked file or configuration database is detected.

Environment Tree The Environment Tree manages characteristics about a customer's unique data center environment, including hardware, location of servers, network infrastructure, application names, business units, and service levels assigned to servers and applications. The information contained in the Environment Tree, combined with the information contained in the Software Tree, is utilized by the Opsware Automation Platform to model and simulate operational changes before they are executed in the production environment.

facility The collection of servers that a single Opsware core manages. A facility can be all or part of a data center, server room, or computer lab.

full backup During a full backup, all tracked configuration files that were selected to be backed up are backed up (and not just the files that have changed). Full backup is performed if you select backup as the action for a tracked configuration file.

full reconcile A reconcile process that reconciles a server with all of the nodes that it has been assigned to.

gateway See Opsware Gateway.

Global Shell A terminal window for the Opsware Global File System (OGFS) in your Opsware SAS.

group See server group.

IDK Intelligent Software Module (ISM) Development Kit. The tools from Opsware Inc. used to build and upload ISMs.

Import Media tool A utility script included with Opsware SAS that is used to import OS media from the Media Server to Opsware SAS.

inclusions/exclusions criteria Specifies how to include and exclude directories and files during the snapshot or audit process.

incremental backup During an incremental backup, only targets that have changed since the last backup (and that have been selected to be backed up) are backed up. Incremental backup is performed if you select backup as the action for a tracked configuration file.

inherited attachment An attachment that is inherited from an ancestor folder or a template.

initialization Legacy term. *See* OS Provisioning.

IP Range Groups A designated set of servers assigned to a customer account, grouped by either a physical or a logical list.

IP Ranges A designated grouping of servers.

ISM Intelligent Software Module. A set of file and directories that include application bits, installation scripts, and control scripts. When an ISM is uploaded into an Opsware core, a node is created for the application and installable packages are attached to the node.

ISM control A script within an ISM package that can be run on a managed server.

job Any major process run by the Opsware Command Center or the Opsware Command Center Client such as Communication Test or Install Software.

Live directory In CDR, the directory that stores the actual code and content required to run a live site.

local attachment An attachment that is attached directly to a folder or a template.

MAC *See* Media Access Control Address (MAC).

Machine ID (MID) A unique identifier that Opsware SAS uses to identify the server. Opsware SAS assigns a unique number to the server when it first registers and stores the Machine ID and uses it to identify each server.

managed server A Server that has an Opsware Agent installed on it and is under the control of a particular Opsware core.

management IP The IP address that Opsware SAS uses to communicate with the Opsware Agent on the server.

manifest Within CDR, a list of files that indicate the results or preview of an update to be performed. Each entry in the list specifies the file size, last-modified date and timestamp, and the full directory path to the listed file.

Media Access Control Address (MAC) The network interface card's unique hardware number. The MAC is used as the server's physical address on the network.

Media Resource Locator (MRL) A network path in URL format that is registered with Opsware SAS. The path defines the installation media for an OS.

Media Server Contains the vendor-supplied OS media used during OS provisioning over the network. The OS media on the Media Server is accessed over the network by using NFS for Linux and Solaris OS provisioning, and by using SMB for Windows OS provisioning.

MID *See* Machine ID.

Model Repository The Opsware database that stores information about managed server configurations within Opsware SAS. It contains all information necessary to build, operate, and maintain an Opsware-managed site, including a list of all servers under management, the hardware associated with these servers, including memory, CPUs, storage capacity, and the configuration of these servers, including IP addresses, DNS configuration, and so on.

Model Repository Multimaster Component The application that propagates and synchronizes changes from each Model Repository database to all other Model Repository databases.

Modeling and Change Simulation Engine Opsware SAS enables users to first model and simulate proposed operational changes to their environment before propagating these changes to production servers and applications. Utilizing the information contained in the Software and Environment Trees, the Modeling and Change Simulation Engine maintains a model of the various hardware and software configurations and other customer characteristics associated with each of the production servers under Opsware SAS control.

MRL *See* Media Resource Locator (MRL).

multimaster core An Opsware core that belongs to a multimaster mesh.

multimaster infrastructure component See Model Repository Multimaster Component.

multimaster mesh A set of two or more Opsware cores that are linked by synchronizing the data in the Model Repositories at each of the cores. The Model Repositories in each of the cores are continually updated so that they are exact duplicates of each other. All the Opsware cores in a multimaster mesh can be managed through a single Opsware Command Center.

My Jobs A page in the Opsware Command Center that displays a list of jobs from the Model Repository such as software installation or server provisioning.

My Scripts Private scripts that can only be executed by the user who created the script. My Scripts are created for personal use.

name-value pairs Legacy term. See custom attributes.

node A hierarchical set of categories or types that classify hardware, software, configuration, or other components of a site's infrastructure. Simplifies server management (for example, servers within Opsware SAS) and the software applications and configurations associated with those servers.

node-based configuration tracking policy A configuration tracking policy defined for a particular software node for a particular application.

OCLI *See* Opsware command Line Interface (OCLI).

OGFS See Opsware Global File System.

Opsware administrator Responsible for overall administration, policy, and practices for individuals accessing Opsware SAS. Can add users and define access to specific Opsware SAS features that allow users to view site information and deploy new code and content to their site.

Opsware Agent Intelligent software on Opsware-managed servers that is used to make changes to the servers. Depending on the request, the agent might use Global Opsware services. Some functions supported include software installation and removal, software and hardware configuration, server status reporting, and auditing.

Opsware Automation features Opsware SAS is made up of a set of Opsware Automation features. Opsware Automation features are the components that automate particular IT processes. The Opsware Automation features include the following functions: Software Provisioning, Patch Automation, Configuration Tracking, Code Deployment and Rollback, Script Execution, and Data Center Intelligence Reporting.

Opsware Discovery and Agent Deployment A feature that helps deploy Opsware Agents to a large number of servers through the Opsware Command Center Client.

Opsware Command Center A web-based user interface for managing the Opsware environment.

Opsware Command Line Interface (OCLI) An alternative interface to the Opsware Command Center. The OCLI allows you to perform some actions not possible though the browser-based interface of the Opsware Command Center, such as uploading multiple packages, patches, AIX filesets, and so forth, in a batch operation.

Opsware core The server side of Opsware SAS server-agent architecture. A core consists of the Opsware components (such as the Model Repository, the Software Repository, the Data Access Engine, and the Command Engine) for a particular installation.

Opsware Gateway Provides connectivity with an Opsware core either directly or through a network of gateways. All traffic between the servers in the Satellite and the core that manages them is routed through Opsware Gateways.

Opsware Global File System (OGFS) The Opsware Global File System is a single, unified file system view of all file systems for all managed servers in Opsware SAS.

Opsware installation Either a standalone core, multimaster core, or Opsware Satellite.

Opsware model space The Opsware Global File System (OGFS) file system structure that is derived from the Model Repository.

Opsware Satellite Installed in a remote facility, an Opsware satellite provides network connection and bandwidth management for a core that manages remote servers. A Satellite must be linked to at least one core, which may be either standalone or multimaster.

Opsware SAS The server management application to preserve the knowledge of system administrators, network engineers, and database administrators in a centralized knowledgebase. Automates previously manual tasks associated with the deployment, support, and growth of a data center infrastructure.

OS Build Agent A part of the OS Provisioning feature that is responsible for registering bare metal servers in Opsware SAS and guiding the installation process.

OS media Installation software for an OS from the software vendor that is distributed on a CD-ROM, or DVD, or can be obtained by downloading the software from the vendor’s FTP site.

OS Provisioning Process of installing a basic set of software components, including an operating system and an Opsware Agent to add a server into the Opsware managed environment. After provisioning is complete, the server is ready to be managed by Opsware SAS.

Package Repository Legacy term. *See* Software Repository.

package A collection of executables, configuration, or script files that are associated with an Opsware-installable application or program. In Opsware SAS a package contains software package files registered in the Software Repository. Contains software for operating systems, applications (for example, BEA WebLogic, IBM WebSphere), databases, customer code, and software configuration information.

packaging server A managed server that has the IDK installed on it. Visual Packager requires a packaging server for each type of operating system for the packages you plan to create.

partial reconcile A reconcile process that only reconciles servers based on the nodes that the user has currently selected.

patch management administrator Administrator responsible for testing patches and defining patch options, such as installation and uninstallation scripts. A patch cannot be installed by other personnel until the patch administrator has marked the patch available through the Opsware Command Center.

Patch Management An Opsware feature that allows you to upload, test, and deploy patches in a safer and uniform way.

permission A setting within a User Group that allows or disallows access to Opsware SAS features and resources. A resource is usually a set of managed servers or software nodes. The set of managed servers corresponds to a facility, customer, or server group.

platform The name and version of an operating system.

post-install script A shell script invoked on a managed server immediately after a software package is installed on a managed server.

post-uninstall script A shell script invoked on a managed server immediately after a software package is removed from the managed server.

pre-install script A shell script invoked on a managed server immediately before a software package is installed on a managed server.

pre-uninstall script A shell script invoked on a managed server immediately before a software package is removed from the managed server.

preview reconcile Before Opsware SAS installs software on a server, it performs a preview reconcile, and determines what will happen when the actual reconcile is performed (for example, what packages will be installed or removed, what server reboots are required, and so forth).

primary IP A locally-configured IP address of the management interface.

private group A type of server group that can be edited, or deleted by the Opsware user who created the server group.

privileges *See* Permissions and User Group.

public group A type of server group that can be created, edited, or deleted by any Opsware user who has Manage Public Server Groups permissions.

realm One or more Opsware Gateways that service the managed servers contained within an Opsware realm. In Opsware SAS, a realm is a routable IP address space, which is serviced by one or more gateways. All managed servers that connect to an Opsware core via a gateway are identified as being in that gateway's realm.

reconcile The process of updating the actual software configuration of a server based on the specified configuration stored in the Model Repository.

reconcile output After a reconcile operation completes, Opsware SAS displays the reconcile output for each server that was reconciled. The reconcile output aggregates output from the various installation, uninstallation, or post-installation scripts, messages from Opsware SAS, and messages from the system utilities that reconcile uses to perform the installation and uninstallation of packages, operating systems, and patches.

Reconcile Software Wizard A Wizard that can enable a user to directly invoke the reconcile process on a selected server or a group of servers.

reference server A managed server that is compliant (performs as expected) and is also referred to as a known working server or a baseline server.

remote terminal A terminal window for a Unix server or an RDP client window for a Windows server.

restore A function of the Automated Configuration feature that allows the user to return the configuration file or database to a previous state, when the backup action for a tracked file or database is selected.

restore Within CDR, the process of restoring the previous Live directory from the Backup directory to the Live directory.

restore queue Queue in which configuration files are placed before they are restored to a server.

Role Legacy term. *See* node or user group.

rollback Within CDR, returns a site to the state prior to the last cutover. During rollback, restores the set of modified and deleted files to the Live directory.

rosh The remote Opsware shell is a command that makes a client connection enabling you to remotely run programs on managed servers.

Satellite See Opsware Satellite

Script Execution See Distributed Scripts.

selection criteria Rules that instruct Opsware SAS what server objects you want to collect information about, how to collect the server objects, and (optionally) file comparison and inclusions/exclusions criteria. Selection criteria is required for the snapshot and audit processes.

sequence Process within CDR that simplifies deployment by grouping a collection of service operations and synchronizations that can be performed as a single task.

Sequence Editor In CDR, a predefined User Group to create, modify, or delete a sequence definition.

Sequence Performer (Production) In CDR, a predefined User Group to directly perform or request performance of a sequence action on production hosts.

Sequence Performer (Staging) In CDR, a predefined User Group to directly perform or request performance of a sequence action on staging hosts.

Sequence Requester (Production) In CDR, a predefined User Group to request performance of a sequence action on production hosts.

Sequence Requester (Staging) In CDR, a predefined User Group to request performance of a sequence action on staging hosts.

servers Any specific hardware. Specific nodes are attached to servers that determine the specific software, configuration, and other server attributes.

server assimilation Opsware SAS assimilates servers that are already functioning in the operational environment, which allows users to deploy and manage new applications installed on those servers. Assimilating servers installs Opsware Agents on the servers and registers them with the Model Repository.

server baselines Process of defining and provisioning servers with standard configurations. Opsware templates can be used to automate the building of complete server baselines.

Server Explorer A feature of the Opsware Command Center Client that allows you to browse and manage servers and server groups in your facility.

server group A feature used to organize servers into groups in order to perform the same action on all of the servers. Server groups can be comprised of individual servers as well as other server groups.

Server ID The primary key in the Opsware Model Repository that represents a given server. The Server ID is used internally in Opsware SAS.

server lifecycle The various server states assigned to a server by Opsware SAS. Server states include Unprovisioned, Available, Installing OS, and Managed.

server management Process by which users can manage and track servers in an Opsware-managed environment. Opsware SAS forces changes to the operating environment by first changing the centralized configuration information in the Model Repository and then changing the actual configuration of physical servers.

Server Pool Servers that have registered their presence with Opsware SAS, but do not have a full operating system installed.

server provisioning The process of installing a basic set of software components that include the operating system, an Opsware Agent, and other system utilities and debugging tools to manage the server. Configuration is defined in the Model Repository.

server reconcile A process that compares a designated server image from the Model Repository with a specific server, checking for configuration, content, versions, and so forth, to determine if the live server is current and up-to-date. Includes OS, applications, upgrades, and patches.

Server Search A feature that allows you to search for servers based on a variety of criteria, including OS version, installed package, customer, and installed patch.

Server Status A feature that defines server availability. The three major status conditions are USE, STAGE, and STATE.

server-based configuration tracking policy A configuration tracking policy that is defined for a particular server or group of servers, rather than for a particular software node (application).

service A host application (for example, BEA WebLogic, Allaire ColdFusion, Microsoft IIS, Apache Web Server, or iPlanet Application Server).

Service Editor In CDR, a predefined User Group to define and modify or delete service definitions.

Service Performer (Production) In CDR, a predefined User Group to directly perform or request performance of service operations on production hosts (servers).

Service Performer (Staging) In CDR, a predefined User Group to directly perform or request performance of service operations on staging hosts.

Service Requester (Production) In CDR, a predefined User Group to request performance of service operations on production hosts.

Service Requester (Staging) In CDR, a predefined User Group to request performance of service operations on staging hosts.

service-instance Multiple independent instances of a service running on a host (for example, BEA WebLogic, which can run single or multiple instances).

Service Levels User-defined categories that are used to group servers in an arbitrary way. For example, a user can group servers by functionality, tier, application, or ontogeny.

Shared Scripts Public scripts that every Opsware SAS user can access.

Site Backup directory In CDR, the directory that stores a complete backup of the Live directory when the user issues a Backup service operation.

Site Previous directory In CDR, the directory that stores the files that have changed between the current Live directory and its previous state prior to the last cutover. It holds all the changes necessary to revert the Live directory back to the state that it was in before the last cutover.

snapshot A record of how an Opsware managed server is configured at a particular point in time. Snapshots allow administrators to audit the configuration of servers and deploy files and software to correct discrepancies. A snapshot can be based on specified server objects. Server Compliance records one snapshot per server.

snapshot job The process that created a snapshot of a server or server group.

snapshot template A definition of a target and selection criteria that will be examined during the snapshot process to capture and record information about a managed server.

Software Provisioning An Opsware feature that allows system administrators to install, configure, and remove packaged software in a systematic way across servers that are distributed over many different facilities. Software provisioning can also involve the automatic execution of installation and post-installation scripts. Software can be provisioned by using the Install Software Wizard, the Install Template Wizard, or by attaching a server to a node and then reconciling the server.

Software Repository The central repository for all software managed by Opsware SAS. It contains software packages for operating systems, applications, databases, customer code, and software configuration information.

Software Repository Cache An Opsware Satellite component that contains local copies of files. The Software Repository Cache stores files from the Software Repository of an Opsware core or from another Software Repository Cache, and supplies the cached files to Opsware Agents on managed servers.

Software Repository Replicator A component providing backup functionality for Software Repositories running in a multimaster mesh.

Software Tree The Software Tree records a variety of information for software applications and operating systems, including data about how changes to a given software application might impact other existing applications.

source In the snapshot process, this is the managed server that information is recorded about. In the audit process, this is an existing snapshot or server you are comparing selection criteria *from*.

standalone core An Opsware core that manages servers in a single facility. Unlike a multimaster core, a standalone core does not communicate with other cores.

static group A server group in which the servers are added to and removed from the group manually.

synchronization Process within CDR to move modified files from a directory on a source host to a directory on a destination host.

Synchronization Editor In CDR, a predefined User Group to create, modify, or delete a synchronization definition.

Synchronization Performer In CDR, a predefined User Group to directly perform or request performance of a synchronization action.

Synchronization Requester In (CDR), a predefined User Group to request performance of a synchronization action.

target In the snapshot process, this is the managed server or server group you are recording information about. In the audit process, this is an existing snapshot, server, or server group you are comparing selection criteria *to*.

template Used to install a set of (usually related) applications through a single invocation of a wizard.

template inheritance Process by which templates and folders inherit all attachments of the folder they reside in. Inheritance is propagated from parent (folder) to child (template or folder) and to all children of children.

tunnel A TCP connection between two Gateways that carries multiplexed TCP or UDP connections.

Update directory The directory that CDR writes to when synchronizing modified files in source and destination hosts. After synchronization, the Update directory is different from the current Live directories. After cutover, the Update directory and current Live directory are identical.

user An individual with access to the Opsware SAS. An Opsware user belongs to one or more User Groups, which control the access of its members.

User Group Represents a role played an organization’s Opsware users. The permissions specified for a user group determine what the group’s members can do with Opsware SAS.

Value Set Editor Enables you to change the values in a configuration file by editing that file's value set. Each entry configuration file is represented inside the value set editor as a “value set” (a key name and a value).

Web Service API A web services interface that facilitates the integration of operations and business support systems with Opsware SAS. The Opsware Web Services APIs allow other IT systems, such as customers’ existing monitoring, trouble ticketing, billing, and virtualization technology, to exchange information with Opsware SAS.

Web Services Data Access Engine A web services interface to the Model Repository that provides increased performance to other Opsware SAS components.

Wizard A graphical user interface that groups a series of data collection operations, actions, and jobs into a logical, easy-to-understand workflow presentation.

Index

A

- accessing
  - CDR 676
  - OCC Client 99
  - OCC client options 106
  - Opsware Command Center 82
  - Server Explorer 158
  - server groups 175
- ACM. *See* Application Configuration Management.
- adding
  - application configuration, server or group to 629
  - configuration template 621
  - custom attributes for servers 284
  - customized tracking policies 651
  - package, new content to 377
  - servers and server groups 249
  - servers to My Servers 122
  - service levels, Opsware Command Center to 287
- Ad-Hoc Scripts
  - creating 431
  - definition of 797
  - executing 431
- adopting software, how to 704
- Agent to Command Engine (CE)
  - connection time-out 752
  - DNS does not resolve 753
  - gateway
    - gateway could not connect to server 755
    - gateway denied access 754
    - gateway time-out 755
    - internal gateway error 755
    - name resolution error on gateway 754
    - no gateway defined 754
  - no callback from agent 756
  - OK 751
  - old agent version 753
  - realm is unreachable 753
  - tunnel setup error 754
- Agent to Data Access Engine (DAE)
  - connection refused 757
  - connection time-out 758
  - DNS does not resolve 758
  - gateway
    - gateway could not connect to server 760
    - gateway denied access 760
    - gateway time-out 761
    - internal gateway error 760
    - name resolution error on gateway 760
    - no gateway defined 759
    - tunnel setup error 759
  - OK 757
  - old agent version 758
  - realm is unreachable 759
  - unexpected error 757
  - untested 757
- Agent to Software Repository (SWR)
  - connection refused 762
  - connection time-out 763
  - DNS does not resolve 763
  - gateway
    - gateway could not connect to server 766
    - gateway denied access 765
    - gateway time-out 766
    - internal gateway error 766
    - name resolution error on gateway 765
    - no gateway defined 765
  - internal setup error 765
  - OK 762
  - realm is unreachable 764
  - server identification error 764
  - unexpected error 762
  - untested 762
- Agent. *See* Opsware Agent.
- agent-server architecture, Opsware SAS 67
- AIX
  - APARs
    - about 525, 526
    - uploading 526
  - LPPs, about 525, 707
  - reconcile on 707
- APARs. *See* AIX APARs.
- Application Configuration Management
  - adding, configuration template 621
  - adding, server or group to 629

  - application configuration types 773
  - application configuration, definition of 797
  - application configuration, description key 774
  - application configuration, overview 603
  - applying and inheriting values 607
  - CML 607
  - comparing, configuration templates 637
  - configuration out of synch with server 639
  - configuration template, overview 603
  - creating
    - application configuration 617
    - configuration template 618
  - default values, inheritance 607
  - definition of 797
  - deleting, configuration template 621
  - editing, default values 626
  - instance values, inheritance 609
  - loading
    - template 622
    - values 633
  - main components 602
  - overview 66, 96
  - preserve values 606
  - process 600
  - pushing changes, servers or groups to 634
  - roll back 641
  - scheduling
    - audit 638
    - push 635
  - sequence merging
    - append mode 613
    - overview 612
    - preprend mode 615
    - primary key option 614
    - replace mode 613
  - service application configurations 784
  - setting values, on 630
  - setting, templates to run as scripts 624
  - show inherited values 606
  - specifying, template order 625
  - system application configurations 775
  - value set editor
    - definition of 811
    - overview 603
- applications, patches
  - installing 542
- asset tracking, overview 113
- assigning, servers to nodes 710
- audit
  - audit job, definition of 797
  - audit results
    - copying, objects 360
    - deleting 361
    - exporting 361
    - overview 352
    - viewing 353, 356, 357, 359
  - audit template, overview 337
  - copying, audit template 346
  - creating, packages 372
  - definition of 298, 334
  - deleting, audit template 347
  - editing, audit template 346
  - loading, audit selection criteria 344
  - overview 334
  - performing 342
    - audit results, from 337
    - audit template, from 337
    - managed server, from 336
    - snapshot, from 336
  - schedule for application configuration 640
  - scheduling 348
  - selection criteria
    - clear all overrides 302
    - comparison criteria 301
    - content 301
    - inclusions/exclusions 301, 302
    - managing 345
    - override default options 301
    - overview 300
    - saving 344
    - setting up 340
  - specifying, source - target, audit template in 339
- audit schedule 347
- audit template
  - copying 346
  - creating 338
  - definition of 797
  - deleting 347
  - editing 346
  - specifying, source - target 339
- available
  - patches, definition of 798
  - servers, definition of 798

B

- backups
  - backup event, definition of 798
  - configuration tracking
    - deleting 665
    - full backup 646
    - history search options 661

incremental backup 646
manual backup 657
restoring 669
full, definition of 801
incremental backup, definition of 801
baseline server 792, 793
baseline server, definition .298
blocked attachment, definition of 798
booting
Linux servers, with PXE .577
Solaris servers, over network 580
Windows servers, with PXE 577
browsers
configuring 80
supported 79
browsing
job logs in the OCC Client 179
properties, directory and file, of 329
snapshot contents 327
Build Manager
definition of 798
OS Build Agents, locating 580

C

CDR. *See* code deployment.
change logs, definition of .798
changing
custom field values 281
use and stage values, servers for 265
character encoding
See encoding 404
checking
management IP of managed server 770
network gateway configuration 770
cloning, servers 268
code deployment
accessing 676
CDR operations previous status, viewing 696
code and content, uploading to staging 673
definition of 798
deployment requests from users, processing .695
deployment, definition of 800
features 675
live directory, definition of 802
log, accessing 696
overview 59
restore, definition of 806
restoring for CDR 687
rolling back code and content 687
sequences
definition of 807
performing 693
sequence editor, definition of 807
sequence performer, definition of 807
sequence requester, definition of 807
services
accessing 688
overview 685
performing by hostname 691
performing by service name 689
site backup directory, definition of 809
site previous directory, definition of 809
static NAT, use with 224
synchronizations, directories, overview 679
user roles, definition of .799
Code Deployment and Rollback (CDR). *See* code deployment.
COM+ objects, browsing, Server Explorer 172
Command Engine to Agent (AGT)
connection
refused 745
time out 746
gateway
gateway could not connect to server 748
gateway denied access 748
gateway time-out 749
internal gateway error 748
OK 745
realm is unreachable 747
request time-out 746
server not registered with Command Engine . . 747
tunnel setup error 748
unexpected error 745
untested 745
Command Engine, definition of 799
command line interface, definition of 804
Common Troubleshooting Tasks
checking, management IP of managed server 770
checking, network gateway configuration 770
resolving, hostname 771
restarting, an Opsware agent 770
verify that agent is running 768
verify that port is open on managed sever 769
Communication Test
about 205
creating, DCI report of 218
definition of 799
errors 207
exporting, unreachable server status list 218
running, multiple servers on 213
searching, unreachable servers for 216

- types of 206
- unexpected error 212

comparing, configuration templates 637

components
- Application Configuration Management, of 602
- Server Explorer of 159

Configuration Markup Language 607

configuration template
- adding 621
- adding, scripts 624
- comparing 637
- creating 618
- definition of 799
- deleting 621
- loading
  - CML file 622
  - values into 633
- overview 603
- specifying, order 625

configuration tracking
- backed up files, restoring 666
- backups, definition of 798
- definition of 797
- email and logging actions, overview 646
- email notification lists
  - creating 647
  - definition of 800
- file info and file versions 664
- manual backups 657
- node-based policies, definition of 803
- overview 60, 643
- policies
  - definition of 799
  - overview 60
  - reconciling customized 655
  - viewing for servers 655
- policy limits 644
- reconcile, definition of 799
- restore queue, definition of 807
- restore, definition of 806
- server-based policies, definition of 808
- servers, customizing policies for 650

configuring, browsers 80

Connection Refused
- Agent to Command Engine (CE) 752
- Agent to Data Access Engine (DAE) 757
- Agent to Software Repository (SWR) 762
- Command Engine to Agent (AGT) 745

Connection Time-out
- Agent to Command Engine (CE) 752
- Agent to Data Access Engine (DAE) 758
- Agent to Software Repository (SWR) 763
- Command Engine to Agent (AGT) 746

controllable application
- opening, Control window 449
- parameters 449
- running, control scripts 450
- viewing results 452

conventions used in the guide 44

converting, snapshot 330

copying
- audit template 346
- files, Server Explorer's File System of 164
- objects 360
- objects, snapshot from 331
- snapshot template 319

creating
- Ad-Hoc Scripts 431
- application configuration 617
- audit template 338
- communication test 212
- configuration template 166, 618
- configuration tracking email notification list 647
- package 166
  - audit results, from 372
  - managed server, from 371
  - native formats, in 373
  - snapshot, from 372
- patch installation order dependencies 533
- reports, agent installation status on 202
- scripts 417
- server groups 235
  - dynamic groups 238
  - static groups 235, 237
- snapshot 315
- snapshot template 312

Crypto Match (CRP)
- certificate mismatch 750
- OK 749
- SSL negotiation failure 750
- unexpected error 749
- untested 749

custom attributes
- adding, for servers 284
- definition of 799
- deleting, for servers 286
- editing, for servers 285
- managing, overview 283
- overview 90

custom extensions
- definition of 799
- running 442

custom fields
  changing value 281
customers
  association with servers 148
  Customer Independent, definition of 147
  definition of 799
  Not Assigned customer, definition of 147
  overview 146
cutting over
  code and content for CDR 683
  definition of 799

D

Data Access Engine, definition of 800
Data Center Intelligence reporting
  definition of 800
  features and benefits 62
  overview 62, 151
data center. *See* facility.
deactivating
  servers, definition of 800
deleting
  audit results 361
  audit template 347
  backups for configuration tracking 665
  configuration template 621
  custom attributes, for servers 286
  files, Server Explorer’s File System of 165
  scripts 422, 423
  servers 268
  snapshot 333, 807
  snapshot template 319
dependencies
  patch installation order, creating 533
depots
  converting 706
  patch management 525
  reconciling 708
deprecated packages, definition of 800
DHCP
  addresses for servers, usage 586, 595
  configuring for servers 295
  Linux servers, requirements for using 579
  re-provisioning servers, requirements for 595
  servers, booting with 572
  Solaris servers, usage of 575, 576
directories
  live, definition of 802
  site backup directory, definition of 809
  synchronization for CDR, overview 679
  Update directory, definition of 811
directories in OGFS
  Application 397
  Customer 398
  Facility 392
  Group 393
  opsw 391
  OS 398
  Realm 399
  root 391
  Script 396
  Server 394
  ServiceLevel 399
  types 390
disabling, custom configuration tracking policies . 654
DNS Does Not Resolve
  Agent to Command Engine (CE) 753
  Agent to Data Access Engine (DAE) 758
  Agent to Software Repository (SWR) 763
DNS, configuring for servers 294
domains, changing for Windows servers 296
duplicating, server groups 252
dynamic, server groups 230

E

editing
  audit template 346
  custom attributes, for server 285
  customized configuration tracking policies 651
  default values, application configuration for . . . 626
  patch options 532
  scripts 420
  server properties 265
  service levels 288
  snapshot template 318
email
  configuration tracking and logging actions 646
  configuration tracking notification list, format . . 648
  notification list attribute, search order for 648
  notification lists for configuration tracking, definition of 800
e-mail notifications 483, 485, 507, 515
enabling, custom configuration tracking policies . 654
encoding
  audit results 355
  choose for application configuration source . . 606
  Global Shell and Remote Terminal 111
  install scripts 383
  managed server 161
  multi-byte characters 386

OGFS 404
package metadata 379
script 419, 423
switching with swenc 406
Terminal and RDP preferences 111
transcoding 408
uploading template file .622
viewing file contents 164

Environment Tree
definition of 800
overview 69

errors
communication test 207
in scripts to install software 552
non-zero return codes, for scripts 440

examples
Opsware Agent Installer, commands & options 726
searching, ways to 142
swapping servers' disks 146

executing
Ad-Hoc Scripts 431
My Scripts or Shared Scripts 426

exporting
audit results 361
unreachable server list to CSV 218

F

facilities, definition of 800

File System
copying, files 164
deleting, files 165
overview 163
renaming, files 165
viewing, contents 163

firewall 388

G

Gateway Could Not Connect to Server
Agent to Command Engine (CE) 755
Agent to Data Access Engine (DAE) 760
Agent to Software Repository (SWR) 766
Command Engine to Agent (AGT) 748

Gateway Denied Access
Agent to Command Engine (CE) 754
Agent to Data Access Engine (DAE) 760
Agent to Software Repository (SWR) 765
Command Engine to Agent (AGT) 748

Gateway Time-out
Agent to Command Engine (CE) 755
Agent to Data Access Engine (DAE) 761
Agent to Software Repository (SWR) 766
Command Engine to Agent (AGT) 749

Global Shell
definition of 801
error messages 410
opening, session 387
Opsware Global File System
overview 66, 96, 385
remote Opsware shell 401
remote terminal 386
shell scripts 794

Globally Unique Identifier .793

groups, definition of 801

GUID. *See* Globally Unique Identifier.

H

hard disks, swapping for servers 146

hardware preparation, overview 575

histories
CDR operations, viewing previous status 696
scripts, viewing versions 424
viewing, for servers 153

HKEY_CLASSES_ROOT .369

HKEY_CURRENT_CONFIG 369

HKEY_LOCAL_MACHINE 369

HKEY_USERS 369

hotfix chaining 460

HP-UX
depots
converting 706
patch management .525
reconciling 708
uploading 528
reconcile on 708

I

icons, servers icons in the OCC 258

IIS Metabase, browsing, Server Explorer 171

import media tool, definition of 801

inheritance
installation wizards 553
overview 553

inherited attachment, definition of 801

Install Operating System Wizard
custom installation 590
template installation 585

Install Patch Wizard 538

Install Patch wizard 500
Install Software Wizard 553
installation
  flags, overview 500, 510, 537
  scripts for patches, overview 536
installing
  application patches 542
  IDK on a packaging server 365
  Install Patch Wizard 538
  issues for software 552
  OCC Client 97
  operating systems
    with custom installation 590
    with templates 585
  Opsware Agents
    augmenting for servers 728
    checklist for 718
    commands and options, examples 726
    definition of 807
    preparing for 718
    using CLI 721
    using ODAD 196
    verification of 727
  OS Build Agents
    verification 581
  software
    with Install Software Wizard 553
  templates
    with Install Templates Wizard 564
Intelligent Software Module (ISM)
  definition of 801
  overview 446
Internal Gateway Error
  Agent to Command Engine (CE) 755
  Agent to Data Access Engine (DAE) 760
  Agent to Software Repository (SWR) 766
  Command Engine to Agent (AGT) 748
IP addresses
  configuring for servers 294
  management IP
    definition of 221
    viewing 222
  primary for servers, setting 224
  primary IP, definition of 222
  searching for servers by 141
  setting for servers 224
IP range groups, definition of 801
IP ranges
  definition of 801
  overview 293
ISM control, definition of .801

J

jobs
  definition of 801
  scheduling for servers .276
  snapshot, definition of .809
  time outs for 279
  viewing, snapshot jobs 333
jump 403

L

LANG environment variable 407
LC_CTYPE environment variable 407
life cycle
  definition of 808
Linux
  booting servers, with PXE 577
  reconcile on 708
  re-provisioning servers, details of 595
loading
  audit selection criteria .344
  snapshot selection criteria 317
  template 622
  values, configuration template into 633
local attachment, definition of 802
loginToServer 403

M

Machine ID (MID), definition of 802
Machine ID Match (MID)
  MID mismatch 767
  OK 767
  unexpected error 767
  untested 767
management IP
  definition of 221, 802
managing
  audit selection criteria .345
  custom attributes 283
  snapshot selection criteria 317
manifests, definition of 802
mbsacli.exe 457, 459
Media Access Control Address, definition of 802
media resource locators (MRLs), definition of . . .802
media server, definition of 802
Microsoft Baseline Security Analyzer 456
Microsoft patch management prerequisites 456
Microsoft Terminal Services 111
Model Repository 455, 458, 474, 475
Model Repository Multimaster Component, definition

of 802
Model Repository, definition of 802
model-based approach
definition of 68
servers, affects on 70
modeling and change simulation engine
definition of 802
overview 70
modeling, public server groups 230
msiexec.exe 459
multimaster
mesh, definition of 803
overview of support in Opsware SAS 77
multimaster core, definition of 803
My Customers, overview 85
My Jobs
Browsing job logs in the OCC Client 179
communication test 218
definition of 803
overview 84
My Profile, overview 90
My Scripts, executing 426
My Servers
overview 91
servers, adding to 122
servers, removing from 123

N

Name Resolution Error On Gateway
Agent to Command Engine (CE) 754
Agent to Data Access Engine (DAE) 760
Agent to Software Repository (SWR) 765
NAT tables
changing, affects of 225
See also static NAT.
navigating
Opsware Command Center 86
Opsware Command Center Client 99
navigation panel, overview 90
network
configuring for servers
after OS provisioning 595
overview 294
Solaris servers, booting over 580
no callback from Agent, Agent to Command Engine (CE) 756
No Gateway Defined
Agent to Command Engine (CE) 754
Agent to Data Access Engine (DAE) 759
Agent to Software Repository (SWR) 765
nodes
configuration policies for, definition of 803
configuration tracking policies for 644
definition of 803
entries for configuration tracking, overview . . . 649
servers, assigning and removing, overview . . . 710
notifying, jobs 277

O

OCC Client
accessing 99
general options 106
installing 97
options, accessing 106
overview 93
packaging servers preferences 109
requirements 97
terminal and shell preferences 110
unmanaged servers options 106
OCC. *See* Opsware Command Center.
ODAD. *See* Opsware Discovery and Agent Deployment.
OGFS. *See* Opsware Global File System.
OK
Agent to Command Engine (CE) 751
Agent to Data Access Engine (DAE) 757
Agent to Software Repository (SWR) 762
Command Engine to Agent (AGT) 745
Crypto Match (CRP) 749
Machine ID Match (MID) 767
Old Agent Version
Agent to Command Engine (CE) 753
Agent to Data Access Engine (DAE) 758
Agent to Software Repository (SWR) 763
opening
Control window 449
Global Shell session 387
remote terminal 166
remote terminal session, unmanaged server . . 202
remote terminal, managed server to 386
operating systems
custom installation 590
failure to install OS, recovering from 592
installing, with Opsware Command Center . . . 583
patch management, supported for 523
provisioning 567
provisioning, overview 58
re-provisioning, requirements for 595
template installation 585
uninstalling, patches for 545

ways to install on servers 584
Opsware administrator, definition of 803
Opsware Agent 458
Agent upgrade tool 732
Communication Test 204
definition of 803
dormant, definition of .800
installation and functionality, verifying 727
Installer 718
Installer options 722
installing 721
installing, using ODAD .196
limits of functionality on servers 184
overview 67, 181
registration 520
running, in a dormant mode 717
server data tracked by .185
Uninstaller options 730
uninstalling 729
Opsware Agent Installer
command line options .722
commands and options, examples 726
definition of 797
uninstaller options 730
Opsware Agent Uninstaller, definition of 797
Opsware Command Center
accessing 82
administration features .89
Code Deployment, in navigation 88
definition of 804
Environment, in navigation 88
Home page, overview 87
icon tooltips 91
icons for servers, defined 258
My Customers, overview .85
My Jobs
definition of 803
in navigation 87
overview 84
My Servers, overview 91
navigating, overview 90
navigation panel, overview 86
OS provisioning 583
patch administration in .531
Reports, in navigation 89
server lists, filtering 120
Servers, in navigation 87
Software, in navigation 87
tasks, overview 84
ways to search 125
Opsware Discovery and Agent Deployment
creating reports 202
definition of 804
installing, Opsware Agents 196
opening remote terminal sessions, unmanaged server to 202
overview 64, 94
permissions required 196
requirements 196
specifying
deployment actions .191
login settings 191
viewing
history, unmanaged server of 201
Opsware Gateway, definition of 804
Opsware Global File System
accessing with ssh 387
copying files to 388
definition of 804
directories
Application 397
Customer 398
Facility 392
Group 393
opsw 391
OS 398
overview 390
Realm 399
root 391
Script 396
Server 394
ServiceLevel 399
overview 389
server filtering 400
Opsware guides
contents of 41
conventions used 44
documentation set 46
icons in guide, explained 45
Opsware model space 400
Opsware model space, definition of 804
Opsware SAS
agent-server architecture 67
Code Deployment & Rollback, features 675
core, definition of 804
documentation set 46
features, definition of 804
installation, definition of 804
model-based approach 68
model-based approach, affecting servers 70
related documentation 46
software in, overview 550

supported operating systems 115, 116
types of users 51
Opsware Satellite, definition of 804
options
OCC Client, accessing .106
unmanaged servers, OCC Client 106
OS Build Agents
Build Manager, locating 580
definition of 804
failure to install, recovering from 582
overview 577
verifying installation 581
OS media, definition of 805
OS provisioning
booting, Windows or Linux servers 577
custom OS installation .590
definition of 805
failure to install OS, recovering from 592
hardware preparation 575
import media tool, definition of 801
Linux servers 574
managed server values .571
media resource locators (MRLs), definition of .802
media server, definition of 802
OS Build Agents
definition of 804
overview 577
using 577
overview 58
re-provisioning Linux, details of 595
Server Pool values 570
Solaris servers 573
template OS installation 585
Windows servers 574

P

package types
AIX APAR 525, 526
depot 369
HP-UX depots 524, 525, 527, 528, 708
HP-UX depots, reconciling 706
LPP 369, 525, 707
MSI 364, 369
RPM 364, 369, 523, 706, 707
Solaris 364, 369
Windows Hotfix 459, 501
packages
definition of 805
deprecated, definition of 800
overview 368
server management, overview 75
viewing
installed packages .167
packaging server
definition of 805
installing, IDK 365
setting up 364
setting, preferences 367
packaging, overview. 370
partial reconcile, definition of 805
passwords, selecting 81
patch compliance 474, 475, 477
patch compliance scan 457
patch installation, previewing 507
patch installation, scheduling 483, 506
patch management
definition of 805
installation order dependencies, creating 533
Microsoft patch releases 457
operating systems, supported 523
overview 59
patch administrators
definition of 805
patch information from Agent 520
patch testing, support for 521
roles 521
supported Unix versions 523
patch policy 473
patch policy exception 475
patch reboot options .480, 504
patch uninstallation, previewing 515
patch uninstallation, scheduling 515
patches
available for use 531
installation flags 500, 537
installing for applications 542
options, editing 532
status, setting 531
types supported 523
uninstallation flags .510, 537
uploading with wizards .529
performing
audit
audit results, from .337
audit template, from 337
managed server, from 336
overview 342
snapshot, from 336
manual backups of tracked configurations 657
sequences for CDR 693
services by hostname for CDR 691

services by service name for CDR 689
services, synchronizations, and sequences, overview 679
synchronizations for CDR 681
permissions
ODAD, required for 196
server groups, required for 227
policy setter 461
port open on managed server, verifying 769
post-install script, definition of 805
post-uninstall script, definition of 805
pre and post install scripts,specifying . 481, 505, 514
preferences
general options, OCC Client 106
packaging servers, OCC Client 109, 367
terminal and shell, OCC Client 110
pre-install script, definition of 805
preparing, installing Opsware Agents, for 718
preparing, patches for uploading 528
pre-uninstall script, definition of 806
preview reconcile
definition of 806
primary IP, definition of 222, 806
private group, definition of 806
private, server groups 229
properties
editing for servers 265
servers, overview 261
public group, definition of 806
public, server groups 229
push 403
application configuration
onto a server or group 634
scheduling 635
PuTTY 110
PXE, booting servers with 577

Q

qchain.exe 459
QNumber 168, 169, 176, 459, 484, 486

R

RDP 110
reach 403
Realm is Unreachable
Agent to Command Engine (CE) 753
Agent to Data Access Engine (DAE) 759
Agent to Software Repository (SWR) 764
Command Engine to Agent (AGT) 747
realm, definition of 806
reconcile
adopting software 704
AIX 707
definition of 806
full, definition of 801
HP-UX 708
installation order for adopted software 704
Linux 708
operating systems, supported 706
output
about 709
definition of 806
partial, definition of 805
performing 700
preview
definition of 806
process, overview 701
Reconcile Software Wizard 711
Reconcile Software Wizard, definition of 806
servers
assigning to nodes 710
reconciling directly 711
removing from nodes 710
Solaris 708
ways to perform 700
Reconcile Patches wizard 474, 478, 479
Reconcile Software Wizard 711
reconciling, customized tracking policies 655
Red Hat Linux 365
reference server 792, 793
reference server, definition of 298, 806
registry (Windows), browsing, Server Explorer . . . 171
remote Opsware shell
definition of 807
modes of operation 403
options and commands 401
overview 401
remote terminal
definition of 806
opening
managed server for 386
overview 166
unmanaged servers to 202
overview 386
removing
earlier versions of Opsware Agents from servers
UNIX 731
Windows 732
Opsware Agent, UNIX and Windows 729
servers

from My Servers 123
from nodes 710
from nodes, overview 710
from server groups 250, 251
from service levels .291
renaming files, Server Explorer's File System of . . 165
reports
overview 151
re-provisioning
Linux servers
about 596
details of 595
requirements for 595
Solaris servers
about 596
requirements for 595
Request Time-out, (AGT) .746
requirements, running OCC Client for 97
resolving
hostname 771
script time out events .439
script uploading errors 439
restarting, an Opsware Agent 770
restoring
backed up files for configuration tracking 666
code and content from old version for CDR . . .687
tracked configurations .666
rolling back
code and content to previous version for CDR 687
definition of 807
restored files for configuration tracking 669
rosh. *See* remote Opsware shell.
RPM
patching 523
reconciling RPM packages 706, 707
rsync 388
RunCommandOnServer 403
running
Communication Test
individual server 212
multiple servers 213
control scripts 450
custom extensions 442

S

saving
audit selection criteria .344
snapshot selection criteria 318
scheduling
server jobs 276
time outs affecting 279
scp 388
script management
functionality for 413
tasks, tips, procedures .416
scripts
creating 417
creation tips 417
deleting 420, 422, 423
Distributed Scripts
definition of 800
overview 60, 411
editing 420
error conditions when installing software 552
error resolution 439
execution
functionality 414
results 414
results, tasks 435
tasks, tips, procedures 424
Global Shell scripts 794
investigating
non-zero return codes 440
partial executions 440
server authentication errors 440
My Scripts, definition of 803
patch installation, overview 536
post-install script, definition of 805
post-uninstall script, definition of 805
pre-install script, definition of 805
pre-uninstall script, definition of 806
resolving, time outs 439
resolving, upload errors 439
running, control scripts .450
setting, templates to run as script 624
Shared Scripts, definition of 809
uploading 417
viewing, version histories 424
ways to initiate 415
searching
details 139
examples of 142
IP addresses 141
overview, search feature of 91
Search box 124
search feature 125, 139
server search, definition of 808
unreachable servers 216
ways to 125
security
Opsware Agents, on servers 183

server management 183
selecting, passwords 81
selection criteria, definition of 334
sequence merging, application configuration
append mode 613
prepend mode 615
primary key option 614
replace mode 613
sequences
editor for CDR, definition of 807
performing for CDR 693
server baselines, definition of 808
Server Compliance
audit results 352
audit template 337
audit, definition of 298, 334
audit, overview 334
baseline server, definition of 298
browsing
file properties 329
snapshot contents 327
converting, snapshot 330
copying
audit template 346
objects 360
objects, snapshot from 331
snapshot template 319
creating
audit template 338
snapshot 315
snapshot template 312
deleting
audit results 361
audit template 347
snapshot 333, 807
snapshot template 319
editing
audit template 346
snapshot template 318
exporting, audit results 361
loading
audit selection criteria 344
snapshot selection criteria 317
managing
audit selection criteria 345
snapshot selection criteria 317
overview 64, 94
performing audit
audit results, from 337
audit template, from 337
managed server, from 336
overview 342
snapshot, from 336
reference server, definition of 298
saving
audit selection criteria 344
snapshot selection criteria 318
scheduling
audits 348
snapshots 320
selection criteria 300
clear all overrides 302
comparison criteria 301
content 301
definition of 334, 807
inclusions/exclusions 301, 302
inclusions/exclusions, definition of 801
override default options 301
server objects 299
setting up
audit selection criteria 340
snapshot selection criteria 313
snapshot
definition of 298, 809
job, definition of 809
overview 307
template 311
template, definition of 809
working with 325
snapshot selection criteria 789
source, definition of 334, 810
specifying, source - target, audit template in . . 339
target, definition of 308, 810
viewing
audit results 353, 356, 357, 359
snapshot job 333
Server Explorer
accessing 158
creating, configuration template 166
creating, package 166
definition of 808
File System
copying, files 164
deleting, files 165
renaming, files 165
viewing, contents 163
working with 163
installed patches, overview 172
main components 159
opening, remote terminal 166
overview 64, 94, 157
server properties, overview 161

  server summary, overview 160
  services, overview 168
  viewing
    packages 167
    server history 174
  Windows IIS Metabase, overview 172
  Windows Registry, overview 171
server groups
  accessing 175
  adding, application configuration 629
  adding, servers and server groups to 249
  characteristics 228
  creating 235
    dynamic groups 238
    static groups 235, 237
  criteria for search 240
  definition of 808
  duplicating 252
  dynamic, definition of 800
  members, overview 175
  modeling, public groups 230
  overview 226
  permissions, required 227
  private, definition of 806
  public, definition of 806
  pushing changes, application configuration of 634
  removing, servers and server groups to 250, 251
  Server Group Browser, overview 174
  setting values on 630
  static, definition of 810
  types
    dynamic 230
    private 229
    public 229
    static 229
  using, server search 240
Server ID, definition of 808
server life cycle, Opsware SAS 255
server lists
  filtering 120
  overview 117
server management
  audit trail, overview 57
  definition of 808
  functions 219
  groups, definition of 801
  IP range groups, definition of 801
  IP ranges, definition of 801
  model-based approach 70
  multiple facilities, in 76
  notifying, jobs 277
  packages 75
  scheduling, tasks 276
  server asset tracking 70
  server life cycle tasks 263
  templates 75
Server Not Registered, Command Engine to Agent (AGT) 747
Server Pool
  definition of 808
  overview 117
server search, using 240
servers
  adding, application configuration 629
  asset tracking, overview 113
  assigning to service levels 290
  association with customers 148
  authentication errors in scripts, investigating 440
  booting
    over network 580
    with PXE 577
  changing value, custom field for 281
  changing, use and stage values 265
  cloning 268
  configuration policies for, definition of 808
  configuration tracking policies, customizing 650
  configuration tracking policy 644
  configuration tracking, entries for 649
  creating a package 371
  custom attributes, deleting 286
  data tracked by Opsware Agents 185
  deactivating 267
  deactivating, definition of 800
  definition of 807
  deleting from the Opsware SAS 268
  DHCP, configuring 295
  domains for Windows servers, changing 296
  filtering in server lists 120
  histories, viewing 151
  history, unmanaged server of 201
  identifying 144
  installing Opsware Agents
    augmenting 728
    verifying 727
  jobs
    scheduling 276
  life cycle
    definition of 808
  Machine ID, definition of 802
  managed, definition of 802
  management IP, definition of 802
  management of 117

media access control address, definition of . . . 802
model-based approach 70
My Servers, removing from 123
network configuration
configuring 294
overview 293
nodes, overview of assigning/removing 710
opening remote terminal sessions, unmanaged server to 202
opening, remote terminal 166
performing audit, managed server from 336
pre-assimilation checklist 718
preparing for installing Opsware Agents 718
primary IP addresses, setting 224
primary IP, definition of 806
properties
editing 265
overview 261
provisioning, definition of 808
pushing changes, application configuration of . 634
reconcile, definition of 808
reconciling directly 711
reference server, definition of 806
re-provisioning
Linux details of 595
requirements for 595
Search box, using 124
search feature, using 125
searching by IP address 141
searching for 124
See also server groups.
Server ID, definition of 808
Service Levels, definition of 809
service levels, removing from 291
setting values on 630
Solaris servers, booting 580
source, definition of 810
status, definition of 808
swapping disks, examples 146
tracking policies, viewing 655
viewing
history 174
packages 167
ways to install operating systems 584
See also server management.
service levels
adding individual 287
assigning servers to 290
editing 288
overview 287
removing, servers from 291
ways to view for servers 289
Service Levels, definition of 809
services
CDR, overview 685
code deployment, accessing 688
definition of 808
editor for CDR, definition of 809
overview 168
performer (production) for CDR, definition of . . 809
performer (staging) for CDR, definition of 809
performing by hostname for CDR 691
performing by service name for CDR 689
requester (production) for CDR, definition of . . 809
requester (staging) for CDR, definition of 809
service-instance for CDR, definition of 809
setting
application configuration values 630
configuration templates to run as scripts 624
packaging server 364
preferences, packaging servers for 367
primary IP address, server of 224
selection criteria, snapshot for 313
setup for servers
pre-assimilation checklist 718
preparation for installing Opsware Agents 718
sftp 388
Shared Scripts
definition of 809
executing 426
snapshot
browsing 327
browsing, file properties 329
converting 330
copying
objects 331
template 319
creating
overview 315
package 372
template 312
definition of 298, 809
deleting 333, 807
template 319
editing, snapshot template 318
job, definition of 809
loading, selection criteria 317
managing, selection criteria 317
overview 307
performing audit, snapshot from 336
scheduling 320
selection criteria 789

clear all overrides 302
comparison criteria .301
content 301
inclusions/exclusions 301, 302
override default options 301
overview 300
saving 318
setting 313
target, definition of 308
template, definition of 809
template, overview 311
viewing, job 333
working with 325
snapshot schedule 320
software
adopted, installation order for 704
Install Software Wizard .553
installation issues 552
overview in Opsware SAS 550
types for installation 551
Uninstall Software Wizard 558
uninstallation issues 552
ways to install 550
software provisioning
definition of 797, 809
features, summary of 551
overview 549
servers, selecting 556
Software Repository 458
Software Repository Cache, definition of 810
Software Repository Replicator, definition of 810
Software Repository, definition of 810
Software Tree
definition of 810
model-based approach, relationship to 71
nodes, explained 72
overview 69
Solaris
booting servers over network 580
OS provisioning 573
reconcile 708
servers, re-provisioning .595
source, definition of 334
specifying
application configuration template order 625
application configuration to run as script 624
deployment actions, ODAD in 191
login settings, ODAD in .191
options, new package content for 380
source - target, audit template in 339
ssh 387
stage values, overview for servers 265
standalone core, definition of 810
static group, definition of 810
static NAT
changing NAT tables, affects of 225
code deployment and rollback, use with 224
static, server groups 229
Summary Review 485, 508, 516
supported operating systems
for managed servers 115
for OCC Client 116
supported, browsers 79
swenc 406
synchronizations
definition of 810
editor for CDR, definition of 810
performer for CDR, definition of 810
performing for CDR 681
requester for CDR, definition of 810
site code and content, overview 679

T

target, definition of 308, 810
telnet 110
templates
blocked attachment, definition of 798
definition of 810
distinguished from packages and nodes 72
inheritance, definition of 811
inherited attachment, definition of 801
Install Templates Wizard 564
local attachment, definition of 802
operating systems, installing with 585
server management, use for 75
time outs
scripts, resolving in 439
server management jobs 279
transcoding 408
troubleshooting
operating systems, installation failure 592
OS Build Agents
installation failure 582
verifying installation .581
partial script executions, investigating 440
script non-zero return codes 440
script time out events 439
script uploading errors .439
unreachable servers 743
Tunnel Setup Error
Agent to Command Engine (CE) 754

Agent to Data Access Engine (DAE) 759
Agent to Software Repository (SWR) 765
Command Engine to Agent (AGT) 748
tunnel, definition of 811
types of, server groups 229

U

Unexpected Error
Agent to Command Engine (CE) 752
Agent to Data Access Engine (DAE) 757
Agent to Software Repository (SWR) 762
Command Engine to Agent (AGT) 745
Machine ID Match (MID) 767
Uninstall Patch Wizard 546
Uninstall Patch wizard 510
Uninstall Software Wizard 558
uninstallation
flags, overview 500, 510, 537
uninstalling
application patches 546
earlier versions of Opsware Agents from servers
UNIX 731
Windows 732
issues for software 552
operating system patches 545
Opsware Agent, UNIX and Windows 729
software with Uninstall Software Wizard 558
Untested
Agent to Command Engine (CE) 751
Agent to Data Access Engine (DAE) 757
Agent to Software Repository (SWR) 762
Command Engine to Agent (AGT) 745
Crypto Match (CRP) 749
Machine ID Match (MID) 767
unzip.exe 459
updating, user profile 81
uploading
code and content to staging for CDR 673
patches
with Opsware Command Center 528
with the Opsware CLI 528
with Upload Patch Wizard 529
scripts 417
use values, overview for servers 265
user profile, updating 81
user roles, definition of 811
users, definition of 811
users, of Opsware 51
using, server search 240

V

value set editor
definition of 811
value set editor, application configuration
editing default values .626
overview 603
preserve values 606
show inherited values .606
verifying
agent is running 768
installation of OS Build Agents 581
port is open on managed server 769
viewing
audit results 353, 356, 357, 359
backup histories for configuration tracking 659
contents, Server Explorer's File System of 163
history
server 153
server for 174
unmanaged server of 201
job details in the OCC .273
job details in the OCC Client 275
list of backup events for configuration tracking 660
managed servers by communication status . .215
My Jobs, Communication Test 218
package details 384
packages 167
results, controllable application of 452
script execution results
immediately 437
stored 438
script version history 424
Scripts list 420
server's configuration tracking policies 655
service levels, ways to 289
snapshot job 333
status of previous CDR operations 696
Visual Packager
adding, new content 377
creating packages 373
audit results, from 372
managed server, from 371
methods 371
snapshot, from 372
definition of 805
installing, IDK 365
overview 65, 95, 363
package, overview 368
packaging, overview 370
setting, packaging server up 364
specifying, options, new package content for . 380

viewing, details 384

W

Web services APIs, definition of 811
Web Services Data Access Engine, definition of .811
Windows Hotfix
installation flags 501
uploading 501
Windows servers
booting, with PXE 577
changing, domains 296
OS provisioning 574
Windows Update Agent 459
wizards
definition of 811
Distributed Script Execution 411
Install and Uninstall Software, overviews 553
Install Operating System
custom install 590
template install 585
Install Patch 538, 542
Install Software 553
Install Template 564
Reconcile Software 711
Reconcile Software Wizard, definition of 806
Run Custom Extension .442
Uninstall Patch 545, 546
Uninstall Software 558
Upload Patch 529
working with
Server Explorer's File System 163
snapshot 325