



# Opsware<sup>®</sup> SAS 7.0 User's Guide: Application Automation

**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](http://www.opsware.com)

Opware SAS Version 7.0

Copyright © 2000-2008 Opware Inc. All Rights Reserved.

Opware Inc. Unpublished Confidential Information. NOT for Redistribution. All Rights Reserved.

Opware is protected by U.S. Patent Nos. 6,658,426, 6,751,702, 6,816,897, 6,763,361 and patents pending.

Opware, SAS Web Client, Model Repository, Data Access Engine, Web Services Data Access Engine, Software Repository, Command Engine, Opware Agent, Model Repository Multimaster Component, and Code Deployment & Rollback are trademarks and service marks of Opware Inc. All other marks mentioned in this document are the property of their respective owners.

Additional proprietary information about third party and open source materials can be found at <http://www.opware.com/support/sas700tpos.pdf>.

# Table of Contents

<b>Preface</b>	<b>25</b>
<hr/>	
<b>Contents of this Guide</b>	<b>25</b>
<b>Conventions in this Guide</b>	<b>27</b>
<b>Icons in this Guide</b>	<b>27</b>
<b>Guides in the Documentation Set and Associated Users</b>	<b>28</b>
<b>Opware, Inc. Contact Information</b>	<b>29</b>
<b>Chapter 1: Visual Application Manager</b>	<b>31</b>
<hr/>	
<b>Overview of Visual Application Manager</b>	<b>32</b>
In This Release – VAM 7.0	32
VAM and the Opware Clients	34
<b>Overview of VAM Features</b>	<b>34</b>
<b>VAM Usage Examples</b>	<b>36</b>
<b>How VAM Works</b>	<b>39</b>
Discovering, Modeling, Managing Business Applications	40
Data Collection and Display	41
VAM Business Application	42
<b>Launching VAM</b>	<b>48</b>

Launching VAM from Servers, Devices, or Device Groups . . . . .	50
Launching VAM from the SAS Client Library . . . . .	52
Launching VAM from Search Results – SAS Client or OMDB Client . . . . .	53
Launching VAM from Generated Reports – SAS Client or OMDB Client . . . . .	53
<b>VAM User Interface . . . . .</b>	<b>54</b>
VAM Toolbars . . . . .	55
Menus and Menu Options . . . . .	60
<b>Adding and Removing Devices in VAM . . . . .</b>	<b>60</b>
<b>VAM Maps . . . . .</b>	<b>62</b>
Tiers Map . . . . .	63
Server Map . . . . .	65
Network Map . . . . .	68
Symbols Used in Maps . . . . .	77
<b>VAM Properties . . . . .</b>	<b>80</b>
Exporting Properties Information to .csv . . . . .	81
Tiers Tree: Tiers, Process Family, Signature Properties . . . . .	82
Devices Tree: Server and Network Device Properties . . . . .	85
Storage and SAN Properties . . . . .	93
SAN Link Properties . . . . .	98
<b>VAM Options . . . . .</b>	<b>98</b>
Virtualization Settings . . . . .	99
Scan Time-Out Preference . . . . .	99
Discovery Settings . . . . .	100
Reset All Settings . . . . .	100
<b>Accessing Servers and Devices From VAM . . . . .</b>	<b>100</b>

---

Opening a Device Explorer . . . . .	101
Opening a Remote Terminal . . . . .	101
Opening a Global Shell . . . . .	101
<b>Creating Business Application Definitions . . . . .</b>	<b>102</b>
Business Application Templates . . . . .	103
Creating Business Application Contacts . . . . .	103
Sending Email to Business Application Contacts. . . . .	104
Business Application Tiers . . . . .	104
Cutting and Copying a Tier. . . . .	105
Pasting a Tier. . . . .	106
Application and Storage Signatures . . . . .	106
Signature Evaluation Order . . . . .	107
Creating an Application or Storage Signature . . . . .	109
Editing Signatures . . . . .	110
Deleting Signatures . . . . .	111
Cutting and Copying Signatures . . . . .	111
Pasting a Signature . . . . .	111
<b>VAM Business Application Management . . . . .</b>	<b>112</b>
Opening a Business Application . . . . .	112
Saving a Business Application . . . . .	113
Saving a Business Application as an Application Template. . . . .	113
<b>Running Scripts on Devices. . . . .</b>	<b>113</b>
<b>ACLs and Server Pool Configurations . . . . .</b>	<b>115</b>

Viewing ACLs .....	115
Comparing ACLs .....	116
Viewing Server Pool Configuration .....	117
Comparing Server Pool Configuration .....	117
<b>Comparing Snapshots .....</b>	<b>118</b>
Creating a Snapshot .....	119
Opening a Snapshot .....	119
Scheduling a Snapshot .....	119
“Source” and “Comparison” Snapshot .....	120
Comparison Types .....	120
Comparing Snapshots .....	125
<b>Filtering VAM Data .....</b>	<b>127</b>
Creating a Data Filter in VAM .....	128
Filter Criteria .....	129
<b>VAM Scan Error Messages .....</b>	<b>131</b>
<b>Scan Results Comparison Heuristics .....</b>	<b>134</b>
<b>Chapter 2: Audit and Remediation .....</b>	<b>137</b>
<b>Overview of Audit and Remediation .....</b>	<b>138</b>
Audit and Remediation Examples .....	138
Audits .....	139
Audit Policies .....	140
Audits and the Compliance View .....	140
Snapshots .....	140
<b>Terms and Concepts .....</b>	<b>140</b>
<b>Audits .....</b>	<b>143</b>

---

Audit Comparison Types.....	143
The Auditing Process.....	145
Audit Elements.....	146
<b>Creating an Audit.....</b>	<b>148</b>
<b>Saving an Audit as Audit Policy.....</b>	<b>150</b>
<b>Viewing Server Audit and Snapshot Usage.....</b>	<b>151</b>
<b>Configuring an Audit.....</b>	<b>152</b>
Audit Sources: Server, Snapshot, or Snapshot Specification.....	154
Opsware Audit and Remediation Rules.....	157
Server Objects Used in Audits and Snapshots.....	160
<b>Audit and Remediation Rules.....</b>	<b>163</b>
Configuration Rules: Expected (Target) and Remediation Values....	163
<b>Configuring Specific Audit and Snapshot Rules.....</b>	<b>167</b>

Configuring Application Configuration Rule .....	168
Configuring COM+ Rule .....	174
Configuring Custom Scripts Rule .....	175
Configuring Files Rule .....	178
Configuring Hardware Rule .....	180
Configuring IIS Metabase Rule.....	181
Configuring Internet Information Server Rule.....	182
Configuring Local Security Settings Rule .....	183
Configuring Registered Software Rule.....	185
Configuring Runtime State Rule .....	186
Configuring Software Rule .....	188
Configuring Storage Rule .....	188
Configuring Windows .NET Framework Configurations Rule.....	189
Configuring Windows Registry Rule .....	191
Configuring Windows Services Rule.....	193
Configuring Windows/UNIX Users and Groups Rule.....	194
Configuring The Opware Network (TON) Rule .....	195
<b>File Inclusion and Exclusion Rules .....</b>	<b>198</b>
Inclusion and Exclusion Rule Types .....	199
Example: Including all .txt Files in a Snapshot or Audit.....	200
Example: Including Only File a in a Snapshot or Audit .....	201
Example: Including last temp.txt file and exclude all else .....	202
File Rule Overlap .....	202
Parameterizing Filenames for Opware/Custom Attributes.....	204
<b>Audit Rule Exceptions .....</b>	<b>206</b>



---

Rules That Cannot Have Exceptions .....	206
Considerations When Applying Exceptions to Device Groups .....	206
Adding a Rule Exception to an Audit .....	207
Editing or Deleting a Rule Exception .....	208
<b>Audit Policies .....</b>	<b>208</b>
Creating an Audit Policy .....	209
Locating an Audit Policy in the Folder Library .....	210
Linking and Importing Audit Policies .....	211
<b>Running an Audit .....</b>	<b>213</b>
Running an Audit from the Library .....	213
Running an Audit on a Server from All Managed Servers .....	214
Re-running an Audit from Audit Results .....	216
<b>Scheduling an Audit .....</b>	<b>217</b>
Scheduling a Recurring Audit .....	217
Editing an Audit Schedule .....	218
Viewing a Completed Audit Job .....	219
<b>Viewing and Remediating Audit Results .....</b>	<b>220</b>
Accessing Audit Results .....	221
Remediation Methods: Rule, Server, or All .....	224
Viewing and Remediating Audit Results Differences .....	228
Viewing Audit Results with Exceptions .....	232
Searching for Audits .....	233
Deleting Audits .....	233
Deleting Audit Results .....	234
Archiving Audit Results .....	235
<b>Snapshots .....</b>	<b>235</b>

Snapshot Specification and Snapshot . . . . .	236
Snapshot Used in an Audit . . . . .	236
Audit Policies and Snapshot Specification . . . . .	237
Snapshot Specification Elements . . . . .	237
The Snapshot Process. . . . .	240
<b>Creating a Snapshot Specification . . . . .</b>	<b>241</b>
Creating a Snapshot Specification from a Server . . . . .	241
Creating a Snapshot Specification from the Library . . . . .	241
<b>Configuring a Snapshot Specification . . . . .</b>	<b>242</b>
Configuring a Snapshot Specification . . . . .	242
Configuring Snapshot Specification Rules . . . . .	244
Saving a Snapshot Specification as an Audit Policy . . . . .	245
<b>Running a Snapshot Specification . . . . .</b>	<b>245</b>
<b>Scheduling Snapshot Jobs . . . . .</b>	<b>246</b>
Scheduling a Recurring Snapshot Job . . . . .	247
Viewing and Editing a Snapshot Job Schedule . . . . .	248
Deleting a Snapshot Job Schedule. . . . .	249
<b>Locating Snapshots . . . . .</b>	<b>250</b>
Searching for Snapshots. . . . .	250
Archiving Snapshots . . . . .	253
Deleting a Snapshot Specification . . . . .	254
Deleting a Snapshot. . . . .	254
<b>Copying Objects from a Snapshot to a Server . . . . .</b>	<b>255</b>
Copying Objects to a Server from a Snapshot. . . . .	256
<b>Chapter 3: Server Compliance . . . . .</b>	<b>257</b>
<b>    Overview of SAS Server Compliance . . . . .</b>	<b>257</b>

---

Compliance Usage: Proactive and Reactive . . . . .	259
<b>Compliance Terms and Concepts . . . . .</b>	<b>260</b>
Server Compliance Categories . . . . .	261
Compliance Statuses . . . . .	262
Compliance Status Thresholds – Policy, Server, and Group . . . . .	266
Changing Device Group Compliance Settings . . . . .	267
<b>Viewing Compliance in the SAS Client . . . . .</b>	<b>268</b>
Viewing Individual Server Compliance . . . . .	268
Viewing Compliance for Multiple Servers . . . . .	273
Viewing Group Compliance in the Device Group Explorer . . . . .	277
<b>Adding and Removing Compliance View Columns . . . . .</b>	<b>280</b>
<b>Filtering By Compliance Status . . . . .</b>	<b>282</b>
<b>Refreshing For Latest Compliance Information . . . . .</b>	<b>283</b>
<b>Scanning for Compliance . . . . .</b>	<b>283</b>
<b>Exporting Compliance View Information . . . . .</b>	<b>284</b>
<b>Compliance Remediation . . . . .</b>	<b>285</b>
Group Compliance Remediation . . . . .	286
Compliance Remediation for Servers . . . . .	287
<b>Audit Compliance . . . . .</b>	<b>289</b>
Audit Compliance Status . . . . .	289
Audit Compliance Remediation . . . . .	290
<b>Software Compliance . . . . .</b>	<b>293</b>
Software Compliance Status . . . . .	293
Software Compliance Remediation . . . . .	294
<b>Patch Compliance . . . . .</b>	<b>298</b>
Patch Compliance Status . . . . .	298
<b>Application Configuration Compliance . . . . .</b>	<b>302</b>

Application Configuration (App Config) Compliance Status . . . . .	303
<b>Chapter 4: SAS Client Reports</b>	<b>307</b>
<b>Overview of SAS Client Reports . . . . .</b>	<b>307</b>
<b>Reports Features . . . . .</b>	<b>308</b>
<b>Opsware SAS Client Reports . . . . .</b>	<b>308</b>
<b>User Permissions . . . . .</b>	<b>313</b>
<b>Launching the Reports Feature . . . . .</b>	<b>313</b>
<b>Reports Display . . . . .</b>	<b>314</b>
<b>Running a Report . . . . .</b>	<b>317</b>
<b>Report Results . . . . .</b>	<b>318</b>
Graphical Report . . . . .	318
List Report . . . . .	320
Exporting a Report . . . . .	321
Printing a Report . . . . .	321
<b>Chapter 5: Patch Management for Windows</b>	<b>323</b>
<b>Overview of Patch Management for Windows . . . . .</b>	<b>323</b>

---

Patch Management for Windows Features . . . . .	324
Library . . . . .	326
Patch Management for Windows Prerequisites . . . . .	328
Microsoft Patch Database . . . . .	329
Opware SAS Integration . . . . .	330
Support for Windows Patch Testing and Installation Standardization .	330
Supported Windows Patch Types . . . . .	331
Supporting Technologies for Patch Management . . . . .	332
Windows Hotfixes . . . . .	332
Searching for Patches and Policies . . . . .	333
Roles for Windows Patch Management . . . . .	334
<b>Patch Management Process . . . . .</b>	<b>335</b>
<b>Patch Properties . . . . .</b>	<b>338</b>
Patch Dependencies and Supersedence . . . . .	340
Viewing Windows Patches . . . . .	341
Editing Windows Patch Properties . . . . .	341
Importing Custom Documentation for a Patch . . . . .	342
Deleting Custom Documentation for a Patch . . . . .	342
Finding Vendor-Recommended Windows Patches . . . . .	343
Finding Servers That Have a Windows Patch Installed . . . . .	343
Finding Servers That Do Not Have a Windows Patch Installed . . . . .	343
Importing a Patch . . . . .	344
Automatically Importing Windows Patches . . . . .	344
Exporting a Windows Patch . . . . .	347
Exporting Windows Patch Information . . . . .	348
Deleting a Patch . . . . .	349
<b>Policy Management . . . . .</b>	<b>350</b>

Patch Policy . . . . .	350
Patch Policy Exception . . . . .	352
Precedence Rules for Applying Policies . . . . .	353
Remediation Process. . . . .	354
Remediating Patch Policies . . . . .	355
Setting Remediate Options . . . . .	357
Setting Reboot Options for Remediation . . . . .	357
Specifying Pre and Post Install Scripts for Remediation. . . . .	358
Scheduling a Patch Installation for Remediation . . . . .	359
Setting Up Email Notifications for Remediation. . . . .	360
Previewing a Remediation . . . . .	361
Verifying Patch Policy Compliance. . . . .	362
Creating a Patch Policy. . . . .	363
Deleting a Patch Policy . . . . .	363
Adding a Patch to a Patch Policy . . . . .	364
Removing a Patch from a Patch Policy . . . . .	364
Attaching a Patch Policy to a Server . . . . .	364
Detaching a Patch Policy from a Server . . . . .	365
Setting a Patch Policy Exception. . . . .	366
Finding an Existing Patch Policy Exception. . . . .	366
Copying a Patch Policy Exception . . . . .	367
Removing a Patch Policy Exception. . . . .	367
<b>Patch Compliance . . . . .</b>	<b>368</b>

---

Patch Compliance Scans . . . . .	368
Ways to Start a Patch Compliance Scan. . . . .	368
Starting a Patch Compliance Scan Immediately . . . . .	369
Refreshing the Compliance Status of Selected Servers. . . . .	369
Viewing Scan Failure Details . . . . .	369
Patch Compliance Icons . . . . .	370
Patch Compliance Levels . . . . .	370
Patch Compliance Rules. . . . .	371
Patch Compliance Reports. . . . .	371
<b>Patch Administration for Windows . . . . .</b>	<b>372</b>
Setting the Patch Availability . . . . .	373
Importing the Microsoft Patch Database. . . . .	373
Selecting Windows Products to Track for Patching . . . . .	374
Scheduling a Patch Compliance Scan . . . . .	375
Setting the Patch Policy Compliance Level. . . . .	376
Importing Windows Patch Utilities . . . . .	376
Exporting Windows Utility Files . . . . .	376
Editing the Customized Patch Policy Compliance Level . . . . .	377
<b>Locales for Windows Patching . . . . .</b>	<b>377</b>
Supported Locales. . . . .	378
Overview of Locale Configuration Tasks . . . . .	378
Configuring the Opsware Core for Non-English Locales . . . . .	378
Selecting the Locales of Patches to Import . . . . .	379
End User Requirements for Non-English Locales . . . . .	380
<b>Patch Installation . . . . .</b>	<b>380</b>

Installation Flags . . . . .	381
Application Patches . . . . .	382
Service Packs, Update Rollups, and Hotfixes . . . . .	383
Installing a Windows Patch . . . . .	383
Setting Windows Install Options . . . . .	384
Setting Reboot Options for a Windows Patch Installation . . . . .	385
Specifying Install Scripts for a Windows Patch Installation . . . . .	386
Scheduling a Windows Patch Installation . . . . .	388
Setting Up Email Notifications for a Windows Patch Installation . . . . .	388
Previewing a Windows Patch Installation . . . . .	389
Viewing Job Progress of a Windows Patch Installation . . . . .	390
<b>Patch Uninstallation . . . . .</b>	<b>391</b>
Uninstallation Flags . . . . .	392
Uninstalling a Windows Patch . . . . .	393
Setting Uninstall Options . . . . .	394
Setting Reboot Options for a Windows Patch Uninstallation . . . . .	394
Specifying Install Scripts for a Windows Patch Uninstallation . . . . .	396
Scheduling a Windows Patch Uninstallation . . . . .	397
Setting Up Email Notifications for a Windows Patch Uninstallation . . . . .	397
Previewing a Windows Patch Uninstallation . . . . .	398
Viewing Job Progress of a Patch Uninstallation . . . . .	398
<b>Chapter 6: Patch Management for Unix . . . . .</b>	<b>401</b>
<b>Overview of Patch Management for Unix . . . . .</b>	<b>401</b>



---

Patch Management for Unix Features . . . . .	402
Opsware SAS Integration . . . . .	404
Support for Unix Patch Testing and Installation Standardization . . . . .	404
Library . . . . .	406
Search Feature . . . . .	407
<b>Patch Management Roles for Unix . . . . .</b>	<b>407</b>
<b>Patch Management for Specific Unix Operating Systems . . . . .</b>	<b>408</b>
Supported Unix Versions and Patch Types . . . . .	408
Underlying Technologies for Patch Management on Unix . . . . .	410
AIX Patches . . . . .	411
Solaris Patches . . . . .	412
HP-UX Patches . . . . .	412
Patch Uploads for Unix . . . . .	413
Patch Uploads for Specific Unix Versions . . . . .	413
<b>Patch Properties . . . . .</b>	<b>414</b>
Viewing Unix Patches . . . . .	415
Editing Unix Patch Properties . . . . .	415
Finding Servers That Have a Unix Patch Installed . . . . .	416
Finding Servers That Do Not Have a Unix Patch Installed . . . . .	416
Exporting a Patch . . . . .	416
Deleting a Patch . . . . .	417
<b>Software Policies . . . . .</b>	<b>417</b>
Patch Compliance Reports . . . . .	418
<b>Patch Administration for Unix . . . . .</b>	<b>418</b>
Setting the Default Patch Availability . . . . .	418
<b>Patch Installation . . . . .</b>	<b>419</b>

Installation Flags . . . . .	420
Application Patches . . . . .	421
Installing a Unix Patch . . . . .	421
Setting Unix Install Options . . . . .	423
Setting Reboot Options for a Unix Patch Installation . . . . .	423
Specifying Install Scripts for a Unix Patch Installation . . . . .	424
Scheduling a Unix Patch Installation. . . . .	426
Setting Up Email Notifications for a Unix Patch Installation. . . . .	426
Previewing a Unix Patch Installation. . . . .	427
Viewing Job Progress of a Unix Patch Installation. . . . .	428
<b>Patch Uninstallation . . . . .</b>	<b>428</b>
Uninstallation Flags . . . . .	429
Uninstalling a Unix Patch. . . . .	430
Setting Uninstall Options . . . . .	431
Setting Reboot Options for a Unix Patch Uninstallation. . . . .	431
Specifying Pre and Post Install Scripts for a Unix Patch Uninstallation	432
Scheduling a Unix Patch Uninstallation . . . . .	433
Setting Up Email Notifications for a Unix Patch Uninstallation . . . . .	434
Previewing a Unix Patch Uninstallation . . . . .	434
Viewing Job Progress of a Patch Uninstallation. . . . .	435
<b>Chapter 7: Software Management</b>	<b>437</b>
<b>Overview of Software Installation. . . . .</b>	<b>437</b>
<b>Software Installation Process . . . . .</b>	<b>438</b>
<b>Ways to Install Software in Opware SAS . . . . .</b>	<b>440</b>
<b>Installing or Uninstalling Software on a Server . . . . .</b>	<b>441</b>

---

Installing or Uninstalling Software.....	443
<b>Installing Software Using a Software Policy .....</b>	<b>449</b>
Attaching a Software Policy to a Server .....	449
Attaching a Server to a Software Policy .....	451
Overview of Software Policies Remediation .....	453
Ways to Open the Remediate Window .....	453
Remediating Software Policies .....	454
<b>Uninstalling Software Using a Software Policy .....</b>	<b>460</b>
Detaching a Software Policy from a Server .....	460
<b>Overview of Software Template .....</b>	<b>461</b>
<b>Overview of Running ISM Controls .....</b>	<b>462</b>
Ways to Open the Run ISM Control Window .....	463
Running ISM Controls .....	464
<b>Software Policy Compliance .....</b>	<b>467</b>
Checking Software Compliance Scan .....	468
<b>Software Policy Reports .....</b>	<b>468</b>
<b>Chapter 8: Script Execution .....</b>	<b>471</b>
<b>Overview of Script Execution .....</b>	<b>471</b>
Script Execution Features.....	471
<b>Script Execution Process.....</b>	<b>472</b>
Types of Scripts .....	472
<b>Managing Scripts .....</b>	<b>473</b>

Creating a Script . . . . .	474
Opening a Script in the SAS Client . . . . .	478
Editing Script Properties . . . . .	479
Viewing All the Software Policies Associated with a Script . . . . .	480
Viewing Script Version History . . . . .	480
Locating Scripts in Folders . . . . .	481
Exporting a Script . . . . .	481
Renaming a Script . . . . .	482
Deleting a Script . . . . .	482
<b>Executing Scripts . . . . .</b>	<b>482</b>
Ways to Open the Run Script Window . . . . .	483
Running a Server Script (Saved Script or Ad-Hoc Script) . . . . .	484
Running an OGFS Script . . . . .	490
<b>Chapter 9: Application Configuration Management . . . . .</b>	<b>495</b>
<b>Overview of Application Configuration Management (ACM) . . . . .</b>	<b>495</b>
<b>Application Configuration Creation and Use . . . . .</b>	<b>496</b>
<b>ACM Components . . . . .</b>	<b>498</b>
Configuration Template . . . . .	499
Application Configuration . . . . .	499
Value Set Editor . . . . .	499
Configuration Markup Language (CML) . . . . .	503
<b>Application Configuration Inheritance . . . . .</b>	<b>503</b>
Application Configuration Default Values . . . . .	504
Application Instance Values . . . . .	505
<b>Sequence Merging and Inheritance . . . . .</b>	<b>508</b>

---

<b>Sequence Replace</b> .....	<b>509</b>
Sequence Append .....	509
Sequence Prepend .....	511
<b>Using ACM</b> .....	<b>512</b>
Creating an Application Configuration .....	513
Creating a Configuration Template .....	515
Searching for Application Configurations .....	517
Viewing Application Configuration Template Sources .....	517
Adding or Removing Configuration Templates .....	518
Deleting Application Configurations .....	518
Loading a Template File .....	519
Setting a Configuration Template to Run as a Script .....	521
Specifying Template Order .....	522
Editing Default Values for an Application Configuration .....	523
Attaching an Application Configuration to a Server or Device Group .....	526
Setting Application Configuration Values on a Server or Device Group ...	528
Loading Existing Values into a Configuration Template .....	530
Pushing Changes to a Server or Group .....	532
Scheduling an Application Configuration Push .....	533
Comparing Two Configuration Templates .....	535
Comparing a Template Against an Actual Configuration File .....	535
Scanning Configuration Compliance .....	537
Scheduling a Configuration Compliance Scan .....	538
Restoring to a Previous State .....	540
<b>Chapter 10: Operating System Provisioning</b> .....	<b>541</b>
<b>Supported Operating Systems for OS Provisioning</b> .....	<b>542</b>

<b>OS Provisioning Overview</b> .....	<b>543</b>
Server Lifecycle for OS Provisioning .....	544
<b>OS Provisioning</b> .....	<b>546</b>
Overview of OS Provisioning .....	546
Solaris OS Provisioning .....	547
Windows OS Provisioning .....	547
Network Setup for OS Provisioning .....	548
<b>Hardware Preparation</b> .....	<b>548</b>
<b>New Server Booting</b> .....	<b>550</b>
Booting New Servers with Different Operating Systems .....	550
OS Build Agent .....	551
Booting a Windows (DOS), Linux, or VMware ESX Server with PXE ..	551
Booting a Windows Server with PXE Using WinPE .....	554
Booting a Solaris Server Over the Network .....	556
Ways that the OS Build Agent Locates the Opware Build Manager .	557
Installing OS Build Agents .....	557
Verifying Installation of an OS Build Agent .....	557
Recovering when an OS Build Agent Fails to Install .....	558
<b>OS Installation with the SAS Client</b> .....	<b>559</b>

---

Create an OS Installation Profile .....	559
Create an OS Sequence .....	560
Select Servers in the Unprovisioned Servers List .....	565
Before Running an OS Sequence .....	566
Run an OS Sequence .....	570
Reprovisioning a Managed Server .....	571
<b>Appendix A: VAM Platform Support</b> .....	<b>573</b>
<b>Supported Platforms in VAM</b> .....	<b>574</b>
<b>Appendix B: Glossary</b> .....	<b>579</b>
<b>Index</b> .....	<b>597</b>

---





# 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: Visual Application Manager** : Describes how to use the Visual Application Manager tool to draw detailed layout views of the operational architecture and behavior of distributed business applications in your IT environment. Provides instructions about how to create, edit, and export physical and logical drawings that can help you diagnose and resolve problems.

**Chapter 2: Audit and Remediation**: Describes how to define server configuration policies and make sure that servers in your facilities meet those policy standards. When servers are found to be 'out of compliance' (not configured the way you want them to be), you can remediate the differing server configurations.

**Chapter 3: Server Compliance**: Describes how the Compliance View allows you to view at a glance the overall compliance levels for all the devices in you facility and helps you to remediate compliance problems. The Compliance View displays compliance tests for software policies, application configurations, audits, patches, and duplex status. Each of

these compliance tests is based upon an Opware Server Automation System (SAS) "policy" (user or system defined) which define a unique set up server or device configuration settings or values that help ensure your IT environment is configured the way you want it to be.

**Chapter 4: SAS Client Reports:** Provides information about how to create reports in the SAS Client and how you can perform actions on objects within the reports. These reports include: Server Reports, Compliance Reports, Sarbanes-Oxley (SOX) Reports, Network Reports, User and Security Reports, and Custom Reports.

**Chapter 5: 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 by using patch policies and patch policy exceptions.

**Chapter 6: Patch Management for Unix:** Provides information about managing patches for Unix operating systems by using software policies. 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 7: Software Management:** Provides information about installing and uninstalling software using software policies, installing software using software policy template, running ISM Controls, and performing software compliance scans.

**Chapter 8: Script Execution:** Provides information about creating, and executing Server scripts and OGFS scripts using the SAS Client.

**Chapter 8: Application Configuration Management:** Provides information about managing application configurations through the SAS Client, and includes such topics as creating Application Configurations, Application Configuration inheritance, editing value sets, and applying Application Configurations to a server.

**Chapter 9: 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 SAS Client to install operating systems using OS sequences.

---

**Appendix A: VAM Platform Support:** Provides information about the operating system platforms and architecture that VAM supports scanning and displaying application (process families), server, and device information.

**Appendix B: Glossary:** Defines terminology and acronyms that are unique to Opware SAS.

## Conventions in this Guide




This guide uses the following typographical and formatting conventions.

NOTATION	DESCRIPTION
<b>Bold</b>	Identifies field menu names, menu items, button names, and inline terms that begin with a bullet.
<code>Courier</code>	Identifies text that is entered or displayed at the command-line prompt, such as Unix commands, Opware 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.
<i>Italics</i>	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 icons.

ICON	DESCRIPTION
	This icon represents a note. It identifies especially important concepts that warrant added emphasis.

ICON	DESCRIPTION
	<p>This icon represents a requirement. It identifies a task that must be performed before an action under discussion can be performed.</p>
	<p>This icon represents a tip. It identifies information that can help simplify or clarify tasks.</p>
	<p>This icon represents a warning. It is used to identify significant information that must be read before proceeding.</p>

## Guides in the Documentation Set and Associated Users

- The *Opware® SAS User's Guide: Server Automation* is intended for system administrators responsible for all aspects of managing servers in an operational environment. It describes how to use Opware SAS, introducing the system and the user interface. It provides information about managing servers, remediating servers, script execution, configuration tracking, deploying and rolling back code, and agent deployment. It also explains how to use the Opware Global Shell and open a Remote Terminal on managed servers.
- *Opware® SAS User's Guide: Application Automation* is intended for system administrators responsible for performing the day-to-day functions of managing servers. It reviews auditing and compliance, software packaging, visual application management, application configuration, and software and operating system installation on managed servers.
- The *Opware® SAS Administration Guide* is intended for administrators responsible for monitoring and diagnosing the health of the Opware SAS core components. It also documents how to set up Opware user groups and permissions.
- The *Opware® SAS Planning and Installation Guide* is intended for advanced system administrators responsible for planning all facets of an Opware SAS installation. It documents all the main features of Opware SAS, scopes out the planning tasks

---

necessary to successfully install Opsware SAS, explains how to run the Opsware Installer, and details how to configure each of the components. It also includes information on system sizing and checklists for installation.

- The *Opsware® SAS Policy Setter's Guide* is intended for system administrators responsible for setting up OS provisioning, configuration tracking, code deployment, and software management.
- The *Opsware® SAS Content Utilities Guide* is intended for advanced system administrators responsible for importing content such as software packages into Opsware SAS. It documents the following command-line utilities: OCLI 1.0, IDK, and DET (CBT).
- The *Opsware® Platform Developer's Guide* is intended for software developers responsible for customizing, extending, and integrating Opsware SAS. It documents how to create Web Services, Java RMI, Python, and CLI clients that invoke methods on the Opsware API.

## **Opsware, Inc. Contact Information**

For more information, see the Opsware, Inc. main web site and phone number:

- <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, see the Opsware Knowledge Base:

- <https://download.opsware.com/kb/kbindex.jspa>

To contact Opsware Customer Support, see the following email address and phone number:

- [support@opsware.com](mailto:support@opsware.com)
- +1 (877) 677-9273



# Chapter 1: Visual Application Manager

## IN THIS CHAPTER

This section discusses the following topics:

- Overview of Visual Application Manager
- Discovering, Modeling, Managing Business Applications
- Launching VAM
- VAM User Interface
- Adding and Removing Devices in VAM
- How VAM Works
- Accessing Servers and Devices From VAM
- Creating Business Application Definitions
- VAM Business Application Management
- Running Scripts on Devices
- ACLs and Server Pool Configurations
- Comparing Snapshots
- Filtering VAM Data
- VAM Scan Error Messages
- Scan Results Comparison Heuristics

## Overview of Visual Application Manager

The Visual Application Manager (VAM) 7.0 helps you manage the operational architecture and behavior of distributed business applications in your IT environment by displaying detailed application information in physical and logical drawings.

VAM enables you to scan selected servers or devices in your data center so you can visualize all aspects of your business applications and how it interacts with other components on your network. VAM enables you to create models of your business application's tiers, hierarchies, and signatures (applications and storage), which map to actual application processes and other devices. This provides a detailed and comprehensive picture of how all a business application's components interact.

VAM's detailed picture of your business application includes all related physical and virtual servers, network and storage devices, and remote connections. When you better understand a business application's processes and interrelationships, you can understand how the business applications are distributed and you are likely to be more effective in troubleshooting errors when they occur.

VAM enables you to take snapshots of a business application (on a one time or recurring basis) and compare the results, so you can view and compare differences in your business application at a specific point in time. You can compare two snapshot results to see changes that have occurred and remediate any differences in the results.

You can also view compliance information in VAM, so you can monitor server and device compliance levels and troubleshoot those that are out of compliance.

VAM is tightly integrated with other features in the Opware SAS Client, including the Opware Network Automation System (NAS) Client and the Opware Application Storage Automation System (ASAS). This enables you to perform change management tasks, such as reconfiguring, patching, auditing, and remediating software and patch policies. (The kinds of data you can visualize in VAM, however, depends upon the Opware products you are licensed to run.)

### In This Release – VAM 7.0

VAM 7.0 is released as part of the Opware Server Automation System (SAS) 7.0 and provides the following new features:

- Support for visualizing logical storage dependencies and physical storage connections (SAN fabric) in your data center and displaying how they relate to your business applications – and servers map to the storage devices they are connected to.



VAM displays file systems and their relationship to local and remote storage, FC Adapters, FC switches (physical and virtual) connections and ports, disk arrays, NAS filers, LUN mappings, RAID configurations, SAN Fabrics, and more. Storage information is visible in the VAM maps, panes, tables, and tiers.

- Display of Weblogic and Microsoft IIS properties information and in the VAM maps
- Support for network devices your NAS installation supports
- ACL and Server Pool Configuration visibility:
  - For network devices (such as firewalls, load balancers, routers, switches), you can view load balancer and ACL configuration information
  - For Load Balancers, you can view Server Pool Configuration (in addition to ACLs)
- Visualization of network device and storage compliance
- Ability to create email contacts and associate them with business applications – and send emails to the contacts on the list
- Integration of virtualization relationships into the maps, with the ability to show or hide virtual relationships, where applicable
- Ability to start stop, pause (VMs only), and restart VMs and Zones
- New Infrastructure pane provides rich detailed information about devices, file systems, compliance policies, network interfaces, and more
- Ability to label all IPC lines with their associated protocol
- Ability to add new devices to an existing business application
- Display of routing table information on a server's Properties pane
- Increased OS platform support
- Improved filtering (searching)
- Snapshot scheduling on demand or on a recurring basis
- Improved snapshot comparison, and enhanced ease of use
- Display of DNS servers in the Properties pane and Server and on the Network Maps (LAN and SAN)
- Easier business application signature Properties editing, including extended application signature discovery using environment variables
- Ability to run scripts on devices or in the Opsware Global File System (OGFS).

- Support to run scripts on network devices
- Sharing of VAM objects and VAM through URLs using drag and drop into emails or chat windows, or through embedding VAM launch URLs in web pages.
- New VAM discovery settings to determine remote server dependencies
- Ability to export tables to CSV

### **VAM and the Opsware Clients**

Opsware's Visual Application Manager (VAM) and is a separately licensed product that requires the Opsware Server Automation System (SAS) in order to run.

In order to visualize networking information with Opsware's Network Automation System (NAS) inside of VAM, you must have both a licensed version of NAS integrated with your Opsware SAS core, plus an additional license to run VAM showing NAS data.

Additionally, in order to view storage devices and SAN information from the Opsware Storage Automation System (ASAS) inside of VAM, you must have both a licensed version of ASAS integrated with your Opsware SAS core, plus an additional license to run VAM showing ASAS data.

You can also visualize servers and devices in search and report results from inside the Opsware Operational Management Database (OMDB), which is also a separately licensed product.

If you have not purchased VAM, NAS, ASAS or OMDB, but would like to, contact your Opsware sales representative.

### **Overview of VAM Features**

VAM enables you to perform the following tasks:

- Discover, map, and visualize the process families, connections, dependencies, and storage of multi-tiered business applications
- Visualize business applications that run on virtual servers, showing virtual servers in relationship to their hypervisors, as well as virtual switches and port groups (VMware ESX only)
- Visualize business application information in multiple physical and logical layouts, such as an application view, a server view, a network view (including virtual network devices), a storage and SAN view that displays logical and physical storage

connections, and an infrastructure view that provides inventory and infrastructure information related to objects scanned.

- Organize recognized application and storage signatures into multi-tier applications to create a logical view that can be analyzed to verify correct operation
- Map actual business application process families to application and storage signatures and highlight them with custom color schemes
- Create, schedule and compare VAM snapshots
- Filter business application and snapshots to find exactly the data you are looking for
- Create and share business application templates that represent an ideal application definition
- Run scripts on devices or the Opware Global File System (OGFS) on a one-time or scheduled basis
- Troubleshoot and resolve problems by launching the SAS Client Device Explorer, Network Device Explorer, Global Shell, Remote Terminal, and NAS interface to perform in-depth analysis or to perform actions on the systems under investigation
- Export maps to .gif, .jpg, and .svg files
- Export tables (Properties and Infrastructure tabs) to .csv

### **VAM Prerequisites**

VAM requires an Opware Agent version 7.0 on managed servers, which enables VAM to scan them.

### **Backwards Compatibility with Previous Versions**

VAM 7.0 supports visualizing devices and applications from SAS 7.0, network devices from NAS 7.0, and storage devices and connections from ASAS 1.0.

### **Supported Operating Systems**

VAM collects and displays data about managed servers that are running AIX, Linux, HP-UX, Solaris, VMware ESX, and Windows operating systems. If you are running non-standard kernels on a Linux operating system, VAM might depend on the kernel version, in addition to the operating system version.

For more detailed information on VAM platform support, see Appendix A, “VAM Platform Support”.

## **Business Application Management**

VAM manages all business application information that is collected, displayed, and captured in a Business Application, which can be opened, closed, saved (as a file or template), and edited.

Each Business Application contains a definition (which can be empty) and zero or more snapshots (see “Comparing Snapshots” on page 118), which contain the visual maps (Tiers, Server, Virtual, and Devices). The business application definition specifies a business application's logical construction in terms of tiers, subtiers, signatures, devices, and process families contained in those tiers.

Each VAM snapshot represents the state of a set of managed servers and network and storage devices, the process families running on those servers, the connections among those process families, and any external clients and dependencies. You can refresh snapshots to view the latest business application data, and these results are saved in the Business Application.

A VAM business application that does not have any snapshots attached to it is useful for creating a business application definition that can be saved as a template. See “VAM Business Application Management” on page 112 for instructions on how to save, open, and edit a Business Application. See “Business Application Templates” on page 103 for information on creating a VAM template.

## **VAM Usage Examples**

Understanding how VAM functions within the context of a real data center is best illustrated with some general usage examples:

- Launch VAM
- Discover and Map Business Applications on Servers
- View Related Networking and Storage Information
- Define Business Application Definition
- Troubleshoot Problems and Take Action

### **Launch VAM**


An application administrator starts a new job at a company and one of his first tasks is to add a new feature to a business application, which was maintained by a prior employee, but the former employee left very little documentation. The administrator was provided with the application's source code but does not understand how all pieces of the application work together from an operational standpoint.

To gain a better picture of the application, he opens the Opware SAS Client, selects a group of servers (some of which use remote storage) that the application runs on, and launches VAM.

### **Discover and Map Business Applications on Servers**

VAM scans the selected servers and discovers all applications, signatures, processes and process families, files systems, local and remote storage, and any connections related to all the applications running on the selected servers. VAM displays detailed “maps” of the application (processes and process families) and servers and connections associated with the servers selected (virtual and physical), as well as any related network relationships and remote storage and SAN connections.

The application administrator examines this information and sees a short list of items that contains two servers, one of which is his server, and two network devices. Looking at the Server Map, he selects the box that represents his server, and a properties pane opens to display more detailed information about the server. He notices that the server has virtual machine-related information, so he concludes that his business application may be running on a virtual machine instance.

He then clicks **Show Virtual/Physical Containment Relationships**  on the VAM toolbar, and now the map showing that his server is a virtual machine running on a hypervisor. He double-clicks the hypervisor server and once it expands, he sees his server within it. He now understands that his business application runs on a VMware virtual machine (VM), and has visibility into the physical host (hypervisor) on which the VM runs. He also notices that the server is connected to a SAN disk array, and he can see the connection between the file system on the server and the storage device.

### **View Related Networking and Storage Information**

The administrator then selects the Network Map and sees his server again, but notices that it has a green line connecting it to another box. By clicking on that box and examining the properties, he determines that the other box is a VMware vSwitch, which in turn is connected to a Cisco switch. He can see precisely which VLAN, port group, switch port,

and network interfaces are involved when his business application communicates over the network. He now understands how his business application fits into the network, both physical and virtual.

He takes a closer look at the lines emanating from his server, and notices a prominent, thick black line pointing at some IP address, so he clicks on it. He sees that the line represents 64 connections to some another host on port 1433. It looks like the database that he knows his business application uses. Inside of the VAM, he selects the Server Map, right-clicks, and selects **Add Devices**. He locates the server, which has that same IP address, selects it and his original server, and adds them to the currently loaded Business

Application. He clicks **Refresh Snapshot** , which updates the current view to include the newly added server information.

Now he sees that the thick black line is pointing at the new server, and after drilling down, he discovers it pointing specifically at an SQL Server process. He continues this until he finds his business application running across and depending on 10 separate servers. He also sees that the Oracle process family runs on two different file systems on the server, which are being stored on a SAN disk array. The connection to the disk array is black, so he knows this is not a problem with the remote storage device.

### **Define Business Application Definition**

The application administrator naturally does not want to have to perform all of this manual mapping and discovery each time he wants to view and manage his business application. He knows that the vendor's documentation contains a logical architectural diagram of the business application, so to make his job easier, he uses VAM to create an business application template.

His first step is to create the logical tiers of the business application. He selects the Tree tree and creates four main tiers for the business application: Web, Application, Database, and Storage. He creates sub-tiers for authentication services and integration services. He then defines application signatures to add to each tier, specifying which tier a recognized signature should fall in. For the Storage tier, he creates storage signatures to capture any related storage devices, such as Fibre Channel Adapters on servers, disk arrays, LUN mappings, NAS filers, SAN switches, and so on.


In order to create reusable application and storage signatures for each tier, he specifies the criteria used to recognize it, including process names, open files, listener ports, command line, IDs, and so on.

He continues to do this for each tier in the business application, and then color codes the signatures in each tier. When the business application is visualized in the Tiers or Server or Storage maps, he will be able to see each tier of the business application in different colors. Next time he launches VAM, the business application will map and display according to his definition.

Finally, he saves his business application definition as a template, so it can be reused by others who want to work with the same or similar business application.

### **Troubleshoot Problems and Take Action**

To help keep track of the state of a business application at any given time, the application



administrator continually clicks **Refresh Snapshot**  on the VAM toolbar in order to create new snapshots. Each snapshot is saved inside of a Business Application that contains his business application definition, which can be used later to compare previous snapshots of the business application with a current state to find any important differences and troubleshoot errors.

For example, if at some point his business application malfunctions and stops working, the administrator can open his saved the business application, select the Compare feature, and visualize the differences between snapshots that compare the current state of the business application with the last known good state. Comparing snapshots will show if specific devices are not communicating with other devices. For example, he can drill into the network map and see that an interface is missing from his VMware ESX hypervisor from the same diagram and select Open Remote Terminal to remedy to problem.

## **How VAM Works**


VAM's main function is to visualize business applications in great detail, and to display the relationships among all their parts and processes and the servers and devices they depend on to function.

VAM scans a server (or multiple servers) and network and storage devices to gather this information and displays it in the maps and tiers, visualizing all processes and process families, connections, and devices related to the business application. Each VAM session, which can be saved as a Business Application, allows you to create, visualize, analyze, define, share, and troubleshoot your business applications.

Clicking **Refresh Snapshot**  allows you to scan the current state of the VAM business application and save it. These scan results (called a “Snapshot”) can be compared on a one to one basis using the compare feature (activated by the Compare  toolbar button).

### Discovering, Modeling, Managing Business Applications

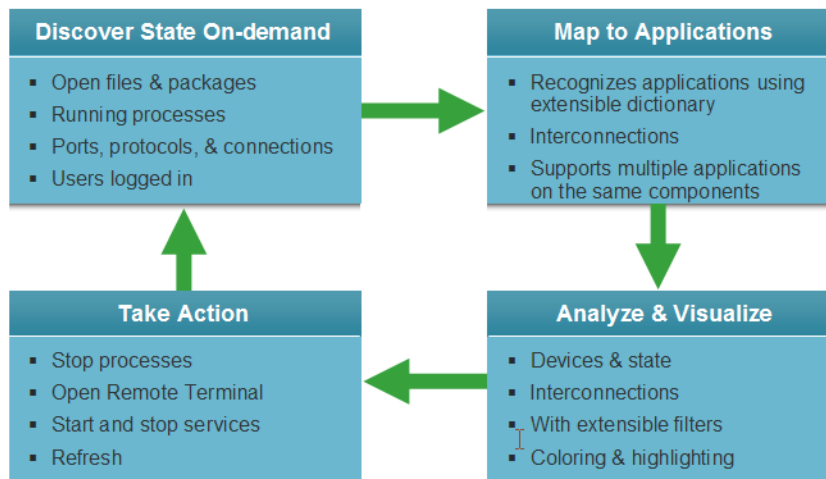
When you launch VAM, it scans, collects, models, maps, and displays information about the operational architecture and behavior of business applications in your IT environment. VAM manages this information in a Business Application that contains application and storage definitions and VAM snapshots.

When you click **Refresh Snapshot**  on the VAM toolbar, VAM creates a new snapshot that captures all the information gathered when you scan a business application (and the servers and devices it runs on) as well as your business application definitions.

VAM uses information gathered from the Opsware SAS, Opsware NAS, and Opsware ASAS data models, leveraging the architecture to collect more data on-demand (such as processes that are running, open ports, and the number of users logged in). It also maps business application data to visualize and analyze your operational environment.

Figure 1-1 shows the process you follow using VAM to display application data in your IT environment.

Figure 1-1: VAM Processes





## Data Collection and Display

VAM scans devices (servers and network and storage devices) and draws layout maps based on data that is collected in real-time results of a VAM snapshot. Device data is captured directly from servers and then recorded in snapshots. Network device data is scanned and then recorded in scan results by NAS – where it is retrieved by the VAM from the Opsware NAS data model. Storage data that relates to the selected device is scanned from the ASAS data model.

When you launch VAM, a set of programs runs on the selected managed devices and captures data. This scanning process collects data about processes running on those devices and the connections between them. It also collects detailed configuration information and current run-time state information about connections and processes. VAM then merges the server data, network, and storage device data to show how servers, interfaces, switches and switch ports, file systems and local and remote storage are connected together.

VAM collects and displays the following information about managed servers, network, and storage devices:

- Processes and process families (grouped into application signatures) that are running on managed servers
- TCP and UDP connections between these processes
- Detailed configuration information
- Current runtime information about servers, connections, and processes
- File systems on servers and how they relate and are connected to process families and local and remote storage devices
- Servers, interfaces, adapters, switches and vSwitches, and switch port connections
- Local and remote storage devices and how they connect to servers, SAN switches, and other storage devices

See “Processes, Process Families, and Extended Process Families” on page 45 for an explanation of how VAM interprets this data. See “Filtering VAM Data” on page 127 for instructions on how to search the data that was collected by object type, such as by process family, network interface, and so on.

## **VAM Business Application**

A business application is a complex collection of services that typically run across multiple servers, networking (LAN and SAN), and storage devices. A business application in VAM consists of business application definitions (tiers, application and storage signatures, and properties definitions) visible in the Tiers tree, and a collection of maps that visualizes relationships between a business application's signatures, processes (and process families), file systems, storage devices, and external clients and dependencies.

A VAM business application maps to actual instances of business applications that are running on servers that VAM has scanned and displayed. A business application, as seen in the Tiers Map, is a collection of processes running on a managed server that maps to a VAM Application definition, as specified in the Tiers tree. A business application can also include storage devices and how they relate to and connect with process families running on servers.

The VAM business application is further explained in the following sections:

- Tiers Tree
- Creating Tiers to Model Business Applications
- Application Signatures
- Processes, Process Families, and Extended Process Families
- Storage Signatures

For information on how to create a VAM application, see "Running Scripts on Devices" on page 113.

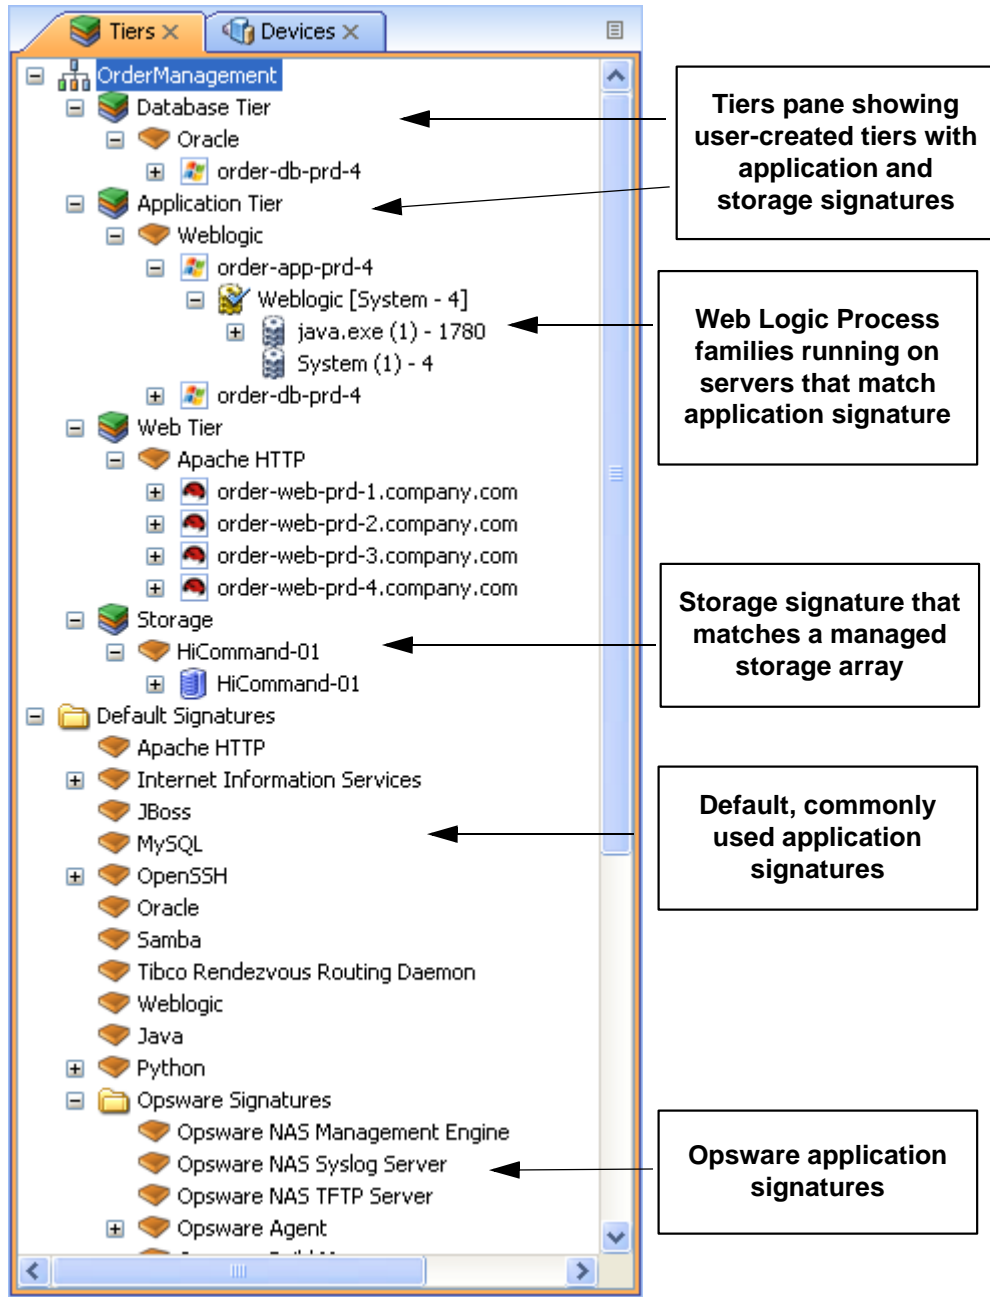
### **Tiers Tree**

The Tiers tree is a logical view of a business application that provides a hierarchical representation of a business application's infrastructure. The Tiers tree provides a different way of looking at what is visually displayed in the Tiers map. The Tiers tree contains tiers and subtiers, which in turn can contain application and storage signatures.

Inside of each application signature is the server or device that the process families run on, and inside of the server or device is the process family itself. A storage signature is similar to an application signature except that its signature matches storage volumes rather than process families (as in an application signature). This allows you to quickly identify and categorize storage dependencies.

See Figure 1-2 for a picture of the Tiers tree.


Figure 1-2: Tiers Tree





Signatures at the bottom of the Tiers tree (such as ones that are associated with the root tier of a business application) do not appear in the Tiers map – instead, they are highlighted in the Network Map, Server Map, and Storage and SAN Maps.

---

If there are no matching process families for a signature, a warning icon  appears next to it, and the tier that contains it, in the Tiers tree.

### **Creating Tiers to Model Business Applications**

Creating tiers enables you to model the logical structure of a business application, representing all of its processes and process families as a diagram of elements that run across multiple servers, displaying the connections among them, clients connecting to them, and dependencies to which they connect. Tier definitions can contain a server filter, which restricts the servers whose process families will match the tier's application signatures; or, a tier can contain a storage filter to define storage volumes that model the relationships between application processes, the file systems they run on, and the types of storage (both local and remote) they use to store data.


Each application consists of a set of tiers and sub-tiers, such as a Web tier running Apache on Linux, an application tier running WebLogic on Windows, a database tier running Oracle on Solaris, and a storage tier to represent disk arrays and NAS filers.

A tier is represented in the Tiers tree by the  icon, which can contain either an application tier object, that consists of application signatures and optional sub-tiers, or storage objects that define local or remote storage devices.

For information on how to create business application tiers, see “Business Application Tiers” on page 104.

### **Application Signatures**

An application signature is an object that represents a process or process family that comprise an application, such as Apache, Oracle, BEA WebLogic, Microsoft® SQL Server, and so on.

An application signature is represented in the Tiers tree by the  icon. An application signature object consists of a signature and preferences.

A signature is a set of rules that you provide and that VAM uses to identify a process family. This set of rules uses data such as process name, open files, command line, connected to port, modules, executable path, and listener port. If VAM discovers the process or process family during a scan according to the signature rule definition, then the process or process family will be added to the signature and highlighted in the maps.


Preferences specify the alias of the application component. These are displayed in the differing background and foreground text color of the different maps.

For more information on creating application signatures, see See “Creating an Application or Storage Signature” on page 109.

### **Processes, Process Families, and Extended Process Families**


In VAM, a process is a running instance of a program in a Unix or Windows environment. A process is discovered and aggregated into process families and extended process families.

A *process family* is a collection of processes that are part of the same Unix session (same name and GID) or a collection of processes that are part of the same Windows session (same name and login session ID).

A process family is represented in the Application (or Device) Tree by this  icon. (Single processes are always grouped visually into process families, and so are also represented by the process family icon.) If the process family is connected to something else (another process family, for example), it is represented in the Application (or Device)

Tree by the  icon.

An *extended process family* is a set of processes that the VAM has heuristically computed to be related, but are not necessarily members of the same process hierarchy.


An extended process family is represented by the  icon.


### **Storage Signatures**

If your core has been enabled to view storage and you are licensed for ASAS, then you can also model logical storage hierarchies to help visualize and understand how your storage devices and SAN relate to the processes used by your business application.

A storage signature is similar to an application signature except that its signature matches storage volumes rather than process families (as in an application signature). This allows you to quickly identify and categorize storage dependencies.

Storage signatures are created in the Tiers tree. The Tiers Map displays them according to modifiable settings of color and name in the signature's properties.

A storage signature is represented in the Tiers tree by the  icon. A storage signature object consists of the name of the storage volume, LUN name and ID, exported path, and any relevant manufacturer information, such as manufacturer name and model number of the related device.

Storage signatures that do not match any actual volumes in the current scan are flagged with a  warning icon.

For more information on creating storage signatures, see See “Creating an Application or Storage Signature” on page 109 and “Examples of Regular Expressions” on page 130.

### **Devices Tree**

The Devices tree is a logical, text-based view of top-level information about managed servers, process families, and network and storage devices. This tree hierarchically displays the same top level information that is shown in the Network Map, Server Map, Storage Map and SAN Map.

The Devices tree contains servers, network and storage devices (physical and virtual) as its top nodes. Below the servers are process families and extended process families. Network devices contain VLANs, ports, and port groups.

Storage devices in the Devices Tree display the following elements:

- NAS filers and their exported file systems and mapped LUNs used by servers
- SAN arrays and their volumes LUN mapped to servers:
- SAN switches and their fibre channel ports.

The Devices tree will also show virtual devices that were scanned when you launched, including the following:

- All hypervisors for all scanned virtual servers are displayed, regardless of whether the hypervisor was scanned.
- Virtual servers are grouped as children of their hypervisors, whether or not the server is managed.
- VMware virtual switches (vSwitches) are grouped under each hypervisor. vSwitches can be expanded to view their port groups.

- Solaris Global zones can be expanded to list all running processes, but this list of processes includes processes on non-global zones not included in the current scan.

The Devices Tree also shows attributes in the Properties pane for the following objects:

WebLogic

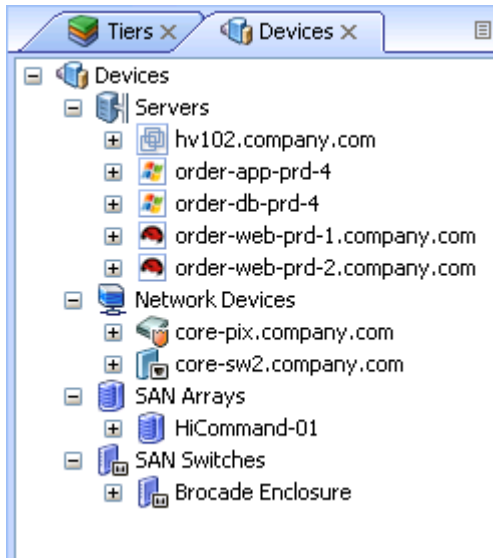
- Applications
- Web Applications
- EJBs
- JDBC Connection Pools

Microsoft IIS

- Web Sites
- FTP Sites
- Binding

Figure 1-3 illustrates the Devices tree, showing servers, network devices, storage and SAN devices.

Figure 1-3: Devices Tree




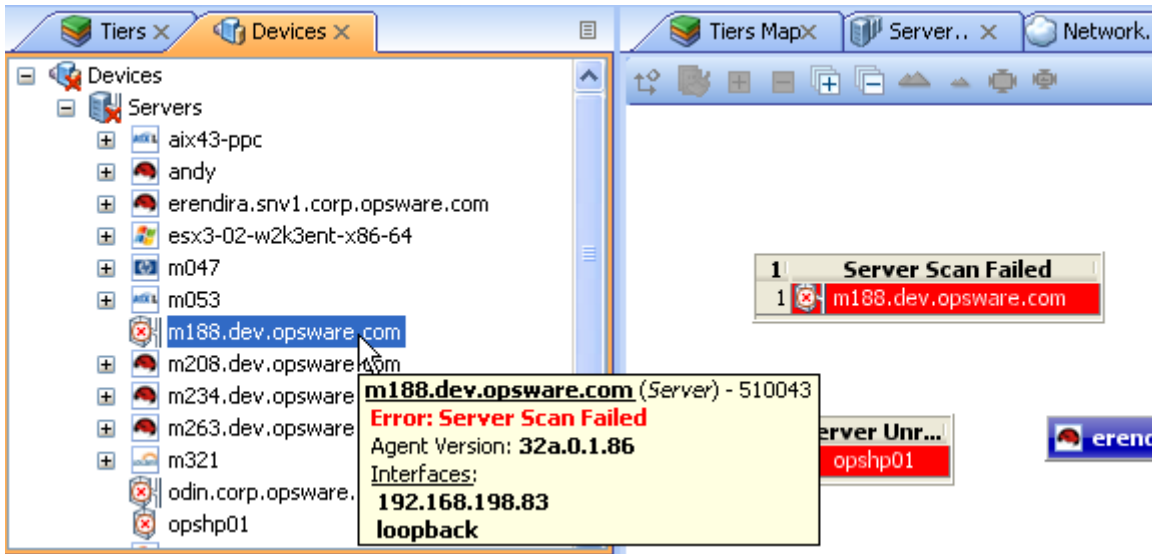
If a server device has an error associated with it, then it will appear in the Devices tree with an error icon on it , for example if the server is unreachable by Opware SAS. When you move your mouse pointer over the device node in the tree, a tooltip message indicates the nature of the error, as shown in Figure 1-4.

Figure 1-4: Devices Tree Server Node with Tooltip Indicating Device Scan Error



For more information on possible device errors, see “VAM Scan Error Messages” on page 131.

## Launching VAM

You can launch VAM in several different ways:

- From a server or group of servers (using drag and drop or menus)
- From the SAS Client Library
- From search or report results
- From a storage or network device (or group of devices)
- From a report or search result inside of the Opware Operational Management Database (OMDB) Client.



When you launch VAM, it performs an extensive scan of the servers or devices you selected – including all virtual servers and their hypervisors.

For more information about how VAM scans a server or device, see “Data Collection and Display” on page 41.

For more information on launching VAM .from the OMDB Client, see the *OMDB User’s Guide*.



The Allow Analyze permission is required to use VAM. You also need read access to each managed server that you plan to scan. Write access to each managed server is not required to run the VAM; however, write access is required to perform any actions on the servers, such as opening a remote terminal or running a script.

To visualize virtualization dates inside of VAM, the View Virtual Server permission must be set to Yes for the user group that your user belongs to. (Without this permission, virtual servers will be displayed just like regular physical servers.) To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide*.

---

You can launch VAM from inside the SAS Client from the following different locations:

- Launching VAM from Servers, Devices, or Device Groups
- Launching VAM from the SAS Client Library
- Launching VAM from Search Results – SAS Client or OMDB Client
- Launching VAM from Generated Reports – SAS Client or OMDB Client

### Launching VAM from Servers, Devices, or Device Groups

To launch VAM servers (virtual servers, or hypervisors), devices (servers, storage devices, or network devices), or groups of devices, perform the following steps:

- 1** Launch the SAS Client from one of the following locations:
  - Click the SAS Client link in the Power Tools section of the SAS Web Client home page.
  - Double-click the SAS Client icon on your desktop (if you installed it on your desktop when you installed the SAS Client).
  - Select **Start** menu ► **All Programs** ► **Opsware SAS Client**.
- 2** From the Navigation pane, select the Devices tree.
- 3** From the Device Groups, Servers list, or Storage list (for an ASAS-enabled core), select a device and perform one of the following actions:
  - From the **Actions** menu, select **Open with** ► **Visual Application Manager**.Or
  - Right-click, and from the menu, **Open with** ► **Visual Application Manager**.

Or

- From the **Tools** menu, select **Visual Application Manager > Open Selection**.

Or

- Select the servers and drag them into an open VAM window

After scanning is completed, the VAM application window appears containing the selected device or devices in the Devices tree, Tiers tree, Properties Panes, Server Map, Network Map, Storage Map, Tiers Map, and the Infrastructure pane.

If a VAM scanning process is taking too long, you can cancel. For more information on how to set the scan timeout value, see “Scan Time-Out Preference” on page 99.



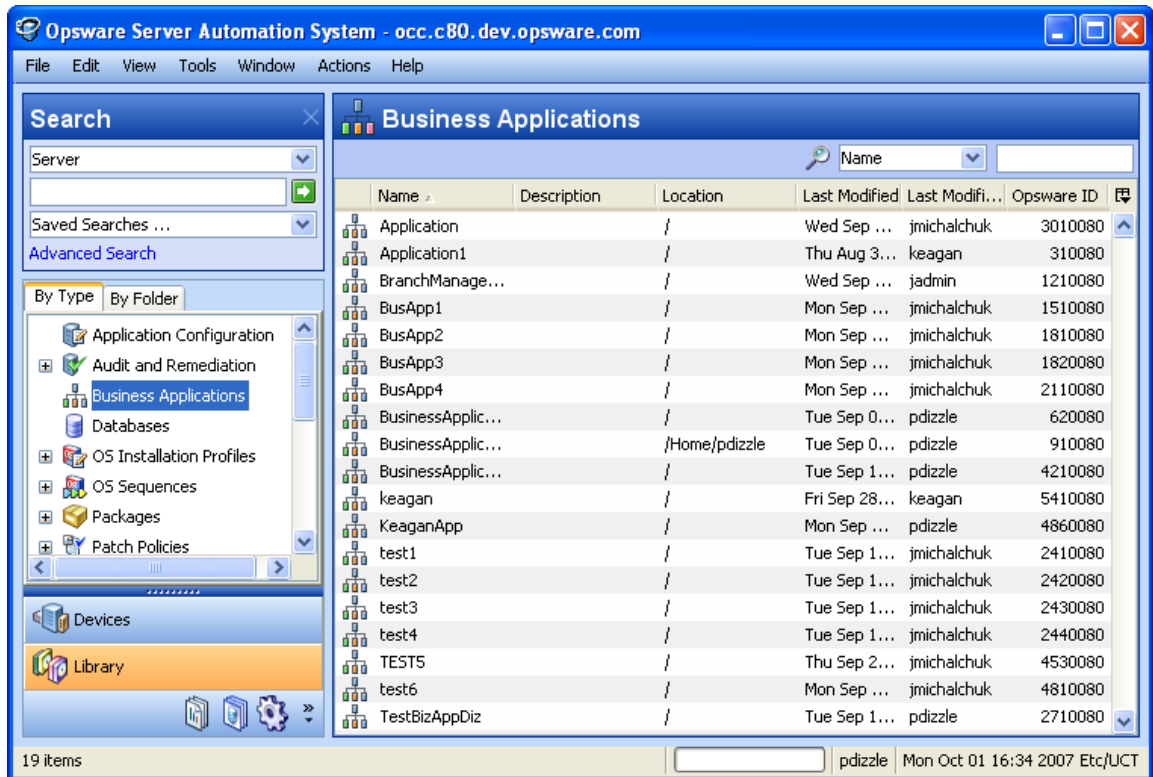
If you have selected virtual servers or a virtual server’s hypervisor to open with VAM, you will initially be asked if you want to scan virtualization relationships – in other words, whether or not to scan any virtual and host servers related to the servers that you selected. This could increase the time it takes to complete the scan, depending on how many virtual servers or hypervisors are related to your selected servers. To control virtual server scan settings, see “Virtualization Settings” on page 99.

---

## Launching VAM from the SAS Client Library

You can save your business applications in the SAS Client Library (given the proper feature and folder permissions), and you can launch Business Applications from the Library as well.

Figure 1-5: Business Applications in the SAS Client Library



For more information about creating and saving VAM business applications, see “Creating Business Application Definitions” on page 102 and “Saving a Business Application” on page 113.

To launch VAM from the SAS Client Library, perform the following steps:


- 1** From the Navigation pane, select Library ► By Type.
- 2** Select the Business Applications object. The Contents pane on the right side shows all VAM business applications you have created and have permissions to see.
- 3** To open a business application and launch VAM, select a business application, right-click, and select **Open**.

## Launching VAM from Search Results – SAS Client or OMDB Client

To launch VAM from search results in either the SAS Client or OMDB Client, perform the following steps:

- 1 Launch the SAS Client from one of the following locations:
  - Click the SAS Client link in the Power Tools section of the SAS Web Client home page.
  - Double-click the SAS Client icon on your desktop (if you installed it on your desktop when you installed the SAS Client).
  - Select **Start** menu ► **All Programs** ► **Opware SAS Client**.

Or, to launch the OMDB Client:

- Select **Start** menu ► **All Programs** ► **Opware OMDB Client**.
- 2 From the Search panel, perform a search for servers. For example, from the top drop-down list, select Servers, and click the green search button .
  - 3 In the search results, select one or more servers and then perform one of the following actions:
    - From the **Actions** menu, select **Visual Application Manager**.

Or

- From the **Tools** menu, select **Visual Application Manager** ► **Open Selection**.

After scanning is completed, the VAM application window appears containing the selected device or devices in the Devices tree, Tiers tree, Properties pane, Server Map, Network Map, Storage Map, Tiers Map, and the Infrastructure pane.

## Launching VAM from Generated Reports – SAS Client or OMDB Client

To launch VAM from report results from either the SAS Client or the OMDB Client, perform the following steps:

- 1 Launch the SAS Client from one of the following locations:
  - Click the SAS Client link in the Power Tools section of the SAS Web Client home page.
  - Double-click the SAS Client icon on your desktop (if you installed it on your desktop when you installed the SAS Client).

- Select **Start** menu ► **All Programs** ► **Opsware SAS Client**.

Or, to launch the OMDB Client:

- Select **Start** menu ► **All Programs** ► **Opsware OMDB Client**.

- 2** From the Navigation pane, select Reports.
- 3** Expand the Reports, and select a report that will display servers in its results.
- 4** From the report results, drill down and select an individual server or multiple servers, right-click, and select **Visual Application Manager**

After scanning is completed, the VAM application window appears containing the selected device or devices in the Devices tree, Tiers tree, Properties panes, Server Map, Network Map, Storage Map, Tiers Map, and the Infrastructure pane.



---

When launching VAM on device groups or when refreshing a previous scan, the servers involved in that scan will consist of the members of those device groups at the time of the scan. Membership may change over time, so two scans of the same selection may produce a different set of scanned servers.

---

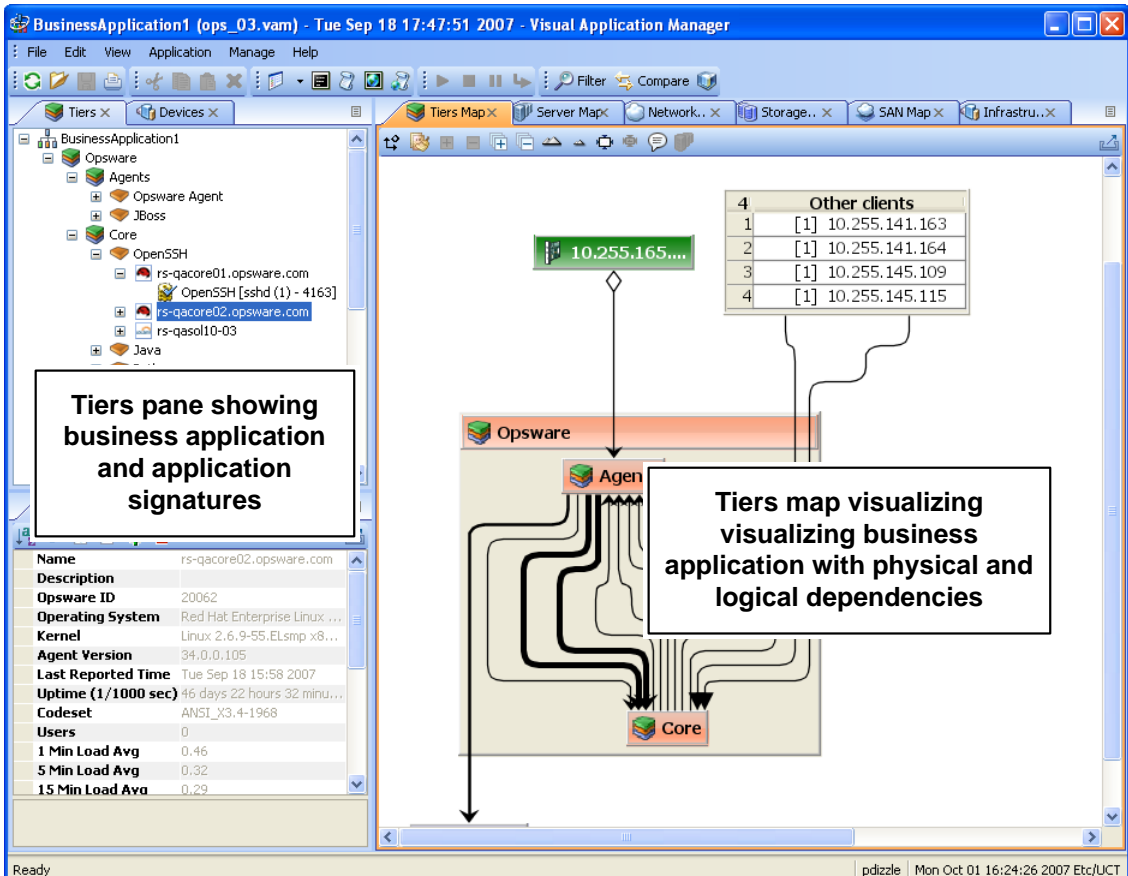
## VAM User Interface

The VAM user interface shows a business application and all its related processes, connections, and devices. It does so by providing the following user interface elements:

- Maps that display the physical layouts of applications – see “VAM Maps” on page 62.
- Trees that display the logical layouts of applications – see “Tiers Tree” on page 42 and “Devices Tree” on page 46.
- Properties panes that provide granular information about a selected object, signature, process, or connection – see “VAM Properties” on page 80.
- Dynamic tool bars and detailed tooltips to provide more information about tree and map objects.

Figure 1-6 shows the types of information that the VAM displays.

Figure 1-6: Opsware Visual Application Manager User Interface



### VAM Toolbars

The VAM toolbars allow you to open, close, resize, and organize different layout views and trees, as well as execute scripts, launch the Global Shell, create a VAM Snapshot, compare Snapshots, and more.

Depending on the tree and view selected, certain toolbar icons will be unavailable. See Table 1-1 for a description of 'VAM's toolbar icons.

Table 1-1: Toolbar Icons in VAM

TOOLBAR ICON	DESCRIPTION
<b>Main Toolbar</b>	

Table 1-1: Toolbar Icons in VAM (continued)








TOOLBAR ICON	DESCRIPTION
	<p>Refreshes the scan results by collecting and displaying new information. Each time you click this button, VAM creates a new VAM snapshot which gets saved as part of the Business Application, and which can be used in a snapshot comparison.</p>
	<p>Opens a previously saved .vam or .vat file.</p>
	<p>Saves the current business application (including maps) as a .vam or .vat file in your local file system or in the Opware Global File System (OGFS). If the business application has not been previously saved, the Save As window displays.</p>
	<p>Prints the selected map. Displays the Print window where you specify page setup (including printing across multiple pages), a title for the printed map, and so on.</p>
	<p>Cuts a selected business application component or a selected tier in the Tiers tree and saves it to the clipboard.</p>
	<p>Copies a business application component in the Tiers tree and saves it to the clipboard.</p>
	<p>Deletes a selected application tier or signature in the Tiers tree or Tiers Map.</p>
	<p>Opens the Device Explorer for the selected device – server, storage device, network device.</p>
	<p>Opens the compliance view of the server (in SAS Client) or network device (in NAS Client).</p>
	<p>Opens Opware NAS for the selected network device (if your Opware core is NAS-enabled).</p>



Table 1-1: Toolbar Icons in VAM (continued)












TOOLBAR ICON	DESCRIPTION
	Opens the Open Remote Terminal window where you select a login ID for a Remote Terminal.
	Launches the Run Script window in the SAS Client (for servers) or the NAS Client New Task - Run Command Script page (for network devices).
	Opens a Global Shell session.
	Launches the Run OGFS Script window.
	Starts the selected virtual server (VM or Solaris local zone).
	Stops the selected virtual server (VM or Solaris local zone).
	Pauses the selected virtual server (VM only).
	Restarts the selected virtual server (VM or Solaris local zone).
	Allows you to filter the currently loaded scan results to find relevant data. For more information, see “Filtering VAM Data” on page 127.
	Allows you to compare to VAM snapshots. Clicking once will display the Compare pane at the bottom of VAM window. Click again to hide the Compare pane. For more information on comparing scan results, see “Comparing Snapshots” on page 118.
	Displays virtualization relationships between virtual server or virtual switches inside of their hosts. Also applies to the Storage Map and Devices Tree.

Table 1-1: Toolbar Icons in VAM (continued)





















TOOLBAR ICON	DESCRIPTION
<b>Properties Pane Toolbar</b>	
	Alphabetize properties of the selected object in the Properties pane.
	Displays additional information about the selected Properties attribute.
	Expands all Properties categories.
	Collapses all Properties categories.
	Adds a contact to your Business Application. (Available only when you select the top-level node of a Business Application in the Tiers tree.)
	Deletes a contact to your Business Application. (Available only when you select the top-level node of a Business Application in the Tiers tree.)
	Exports Properties pane contents to .csv.
<b>Maps Toolbar</b>	
	Rotates the selected view, toggling it between a vertical and a horizontal orientation.
	Redraws all components in the selected view. Components that have been manually revised will retain their sizing.
	Expands selected tiers in the Tiers tree or closed folders in the selected map. Tiers are expanded recursively down to the business application component that they contain. Managed servers underneath the business application components are not expanded.

Table 1-1: Toolbar Icons in VAM (continued)

TOOLBAR ICON	DESCRIPTION
	Collapses all tiers in the Tiers tree or closed components in the selected view.
	Opens all tiers in the Tiers tree or signature in the selected map.
	Closes selected tiers in the Tiers tree or folders in the selected map.
	Zooms into the selected view (enlarges display size).
	Zooms out of the selected view (reduces display size).
	Resizes the selected components in a currently active view to fit within the screen size.
	Resizes all components in the currently active map to fit within the screen size.
	Labels all IPC lines with their associated protocol, such as SSH, HTTP, and so on. Protocols will display in the Server, Network and Tiers Maps. The Show Protocols mode is applied on a per-map basis
	Toggles the maps to display the host server names in the title bar for virtual servers (Server map), virtual devices (Network Map), business application components (Tiers Map), and virtual devices in the SAN Map.
	Exports the current map to an image file.

## Menus and Menu Options

This section discusses the menus and menu options that might not be self-explanatory.

### **File Menu**

If you have made changes to the business application definition and want to set this as the default, select **Set as Default Template** from the **File** menu.

If you have made changes to the business application definition and want to restore the previously saved default business application, select **Reset Default Template** from the **File** menu.

If you would like to import a business application template that has already been saved, from the **File** menu, select **Import Template** and import the selected template.

For more information on business application templates, see "Business Application Templates" on page 103.

### **View Menu**


By default, the **Animate Layout** option is ON (preceded by a check mark). This causes the map to be animated (objects are displayed in motion) each time it is drawn, including a refresh. If the **Animate Layout** option is OFF (no check mark), the map will not be animated (objects are not displayed in motion) each time it is drawn.

## Adding and Removing Devices in VAM

After you have opened and visualized devices in VAM and created your business application definition, you can add more devices to the initial snapshot in order to see how other devices – servers, network devices, storage devices, and so on – relate to the current state of your business application.

For example, you might have created a business application but you are still not sure about all the storage devices it might be using, or, you might see that there are a few managed servers that are connected to your business application, but that were not initially scanned. You can easily add these to your VAM file.

Depending on your settings in the Add Devices window that you see when you add a device, VAM will automatically refresh the VAM snapshot and scan the selected devices. If you would rather not have VAM automatically scan any newly added devices, then you

can add the devices and scan them later by clicking **Refresh Snapshot**  on the VAM toolbar.

In some cases, you might find that some of the devices you have scanned are not necessary to your snapshot, and so you want to remove them. You can easily do so by selecting the device and selecting to remove it.

There are several potential errors that can occur when a Refresh or a device scan in VAM does not work. For a list of potential errors, see “VAM Scan Error Messages” on page 131.

### **Adding Devices to VAM**

To add devices to VAM, perform the following steps:

- 1** From anywhere inside of VAM, either right-click or from the **Application** menu, select **Add Devices**.
- 2** In the Add Devices window, in the left pane you can choose a device category, and the corresponding devices appear in the right pane. You can add servers, device groups, network devices, storage devices, and so on.

The Discovered Dependencies category shows any devices that VAM has discovered to be related to or connected to some of the devices in the existing VAM snapshot.

The Refresh Scan Results option instructs VAM to automatically refresh the VAM snapshot when you click **OK**.


- 3** When you have selected the devices to add, click **OK**. Be sure to save the results or these newly added devices will not be saved in the Business Application.
- 4** If the Refresh Snapshot Results option was not selected, click **Refresh Snapshot**




on the VAM toolbar so VAM will scan the newly added devices. Any devices that have been added to VAM without being refreshed appear as a translucent icon and will not display any properties information.

### Removing Devices From VAM

to remove devices from VAM, perform the following steps:

- 1** From inside of VAM, from one of the maps or the Devices pane, select a device, right click or from the Application menu, select Remove Devices.
- 2** You are asked to confirm that you want to remove the selected device. Click **Yes** to remove the device.
- 3** To make sure your VAM snapshot is up to date, click **Refresh Snapshot**  on the VAM toolbar so VAM can update the snapshot and scan the newly added devices, and then save the Business Application.

When a device is removed from a scan, the device and all connections to and from it and all external client IP addresses are removed in the maps, trees and tables, including links to other managed servers in the scan.

When you click **Refresh Snapshot** , these lines to the other managed servers may display as a Client IP or other dependencies.

### VAM Maps

VAM provides five visual maps and one Infrastructure pane that display physical and logical drawings of managed servers, network and storage devices, and connections in your environment: the Tiers Map, Server Map, Network Map, Storage Map, SAN Map, and the Infrastructure pane.

To enable you to see and understand how your application functions, VAM provides the following maps:

- Tiers Map
- Server Map
- Network Map
- Storage Map
- SAN Map
- Infrastructure Pane

In addition to viewing the VAM maps, you can also:

- Show any virtualization relationships in the maps (virtual servers or devices, hypervisors, switches, and more. See “Showing Virtual Server Relationships in the Server Map” on page 65.
- Export a map to a .gif, .jpg, and an .svg file. See “Printing a Map” on page 79.
- Print a map on single and multiple sheets of paper. See “Printing a Map” on page 79.
- View IPC service names in the Server or Network Maps. See “Showing IPC Service Names in Maps” on page 80.


## Tiers Map

The Tiers Map displays the logical structure of a business application, including a business application’s tiers and the connections between its application and storage signatures, external clients, and other dependencies. By default, this map is initially empty until you create the tiers and define signatures that comprise an application. See “Creating a Tier” on page 105.

In addition, the Tiers Map shows the external IP addresses (Client IPs) that are connected to the application and the external IP addresses (external dependencies) that the application connects to and depends on.

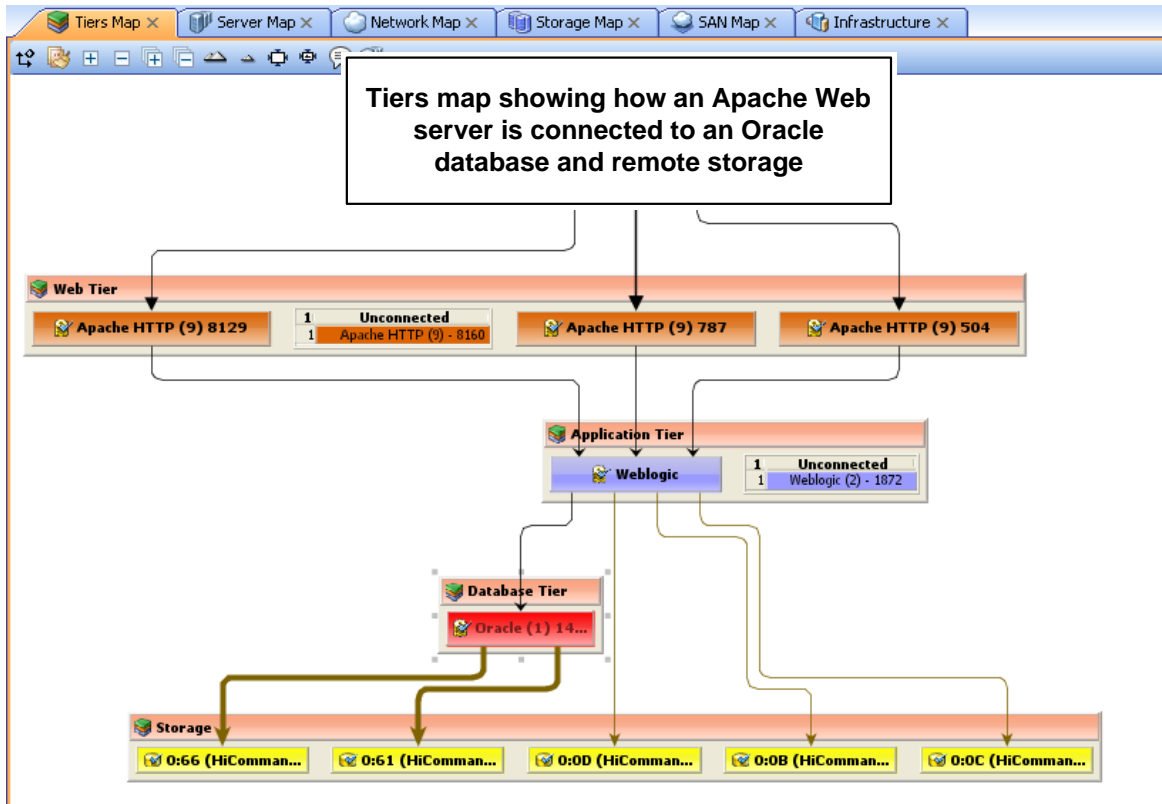
In this release, the Tiers Map also supports a new object called a storage signature, for those cores which are ASAS-enabled.

A storage signature is similar to an application signature except that its signature will match storage volumes rather than process families (as in an application signature). This allows you to quickly identify and categorize storage dependencies. The Tiers Map displays storage signatures according to modifiable settings of color and name in the signature’s Properties.

Storage signatures that don’t match any actual volumes in the current scan will be flagged just like an application signature, with a warning  icon.

Signatures can be attached to a tier and used to recognize process families as named elements of an application. Tiers that do not have process families are shown in the Tiers Map. You can group them within a tier and make them visually distinct by modifying their color and name. See Figure 1-7.

Figure 1-7: Tiers Map



If an application (or storage) signature does not have any process families (or storage devices) associated with it, then the tier object title bar will display a warning icon ⚠, for

example,  Tier1 .



## Server Map


The Server Map displays the physical layout of how elements of a Tiers Map to a set of servers (virtual or physical), including the process families that are running on servers and how those processes families are connected to one another.

The Server Map shows the external IP addresses (client IPs) that are connected to the application and the external IP addresses (external dependencies) that the application connects to and depends on. If one of these external connections is an Opsware SAS managed server, then VAM will display the connection as a server.

If you want to include these servers to the scan, right-click the server and select Add


Devices. After the device is added, click **Refresh Snapshot**  on the VAM toolbar.

In addition, in both the Server Map and Network Map, a DNS Servers element displays DNS servers in use by all the servers in the scan. If managed server information is known about any of these servers, then you can right-click to Add Devices. No connections will be shown to this DNS servers node.

If you want to see the type of service being used for connections between processes (and the devices they run on) in the Server Map, such as http, ssh, telnet, click the Show IPC Service Names in Maps  button.

### **Showing Virtual Server Relationships in the Server Map**

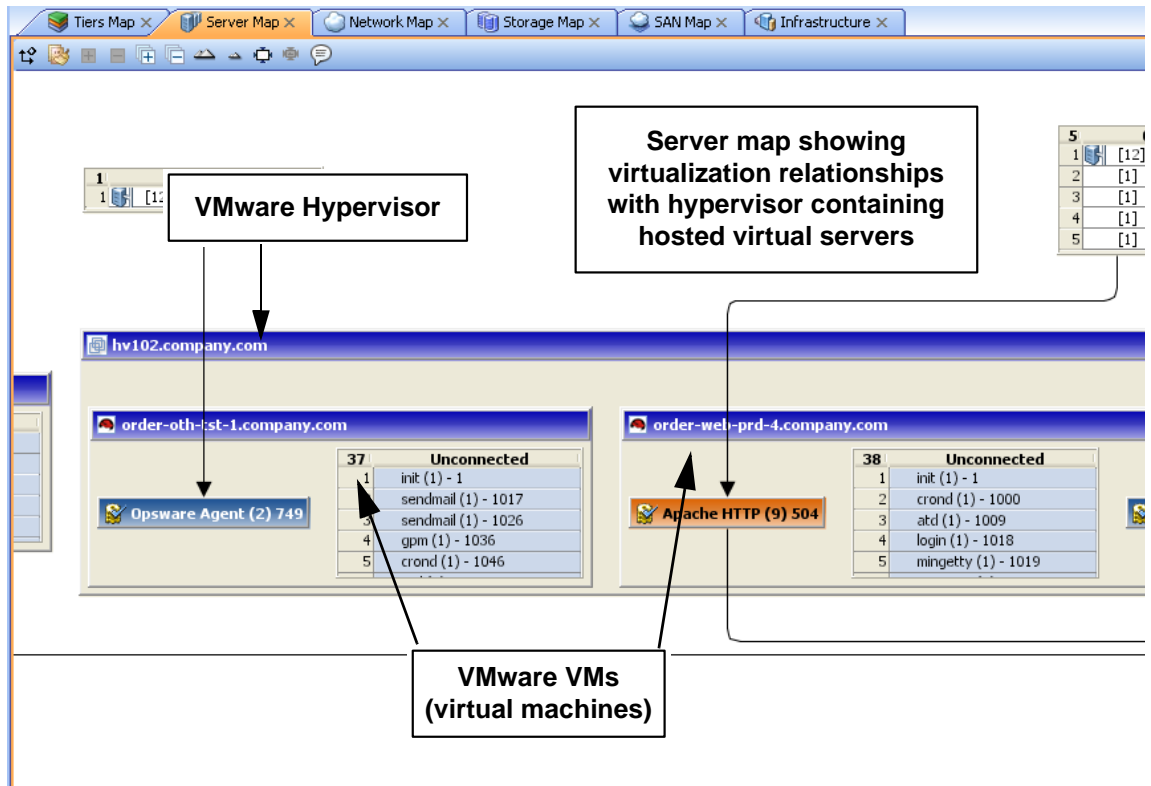
If you scanned virtual servers or devices in VAM, or any devices that contain or are connected to other virtual devices, then clicking **Show Virtual/Physical Containment**

**Relationships**  will display the relationships between virtual servers and the physical devices that host them in the Server Map (and any other map that shows virtual devices).

Clicking this button shows virtual servers inside of the hypervisor that is hosting them, so you can see which servers are virtual, and which are physical, and how they relate to one another.

Figure 1-8 shows the Server Map displaying a physical (non-virtual) server (top) and a virtual server (bottom) with the hypervisor name showing in the title bar.

Figure 1-8: Server Map



For virtual server technologies, if you choose to scan only the hypervisor but not its guests, then you will only see limited information about the virtual servers and will not be able to open them and view their contents.

For all Solaris global zones that are scanned, the Virtual Map displays a list named Other Zones. It contains all processes visible in the global zone that are actually running in a non-global zone, but that were not included in the scan.


For VMware ESX hypervisors, vSwitches are shown alongside virtual machines, in addition to the connections between the virtual machines and the vSwitches' port groups, which will always appear as a green matching-duplex ethernet connection.

The information that VAM is able to display is also dependent upon whether or not the hypervisor or virtual server has an agent installed on it.

For example:

- It is possible that a virtual server has an agent installed on it, but not its hypervisor. In this case, VAM will only display the scanned virtual server but not its hypervisor.
- It is possible to have a hypervisor that has an agent installed on it, but some or all of its guest virtual servers do not. In this case, VAM will show all guest virtual servers but only with limited virtual server information. In other words, you won't be able to open it (drill down into it).







If you open a virtual server and the arrangement of server boxes is difficult to view inside the map, click the Rotate Layout icon  on the toolbar and VAM will rotate the layout for a different unique maps of the servers.

### **Starting, Stopping, Suspending, Resetting Virtual Servers**

In addition to viewing virtual servers and their relationships, you can also stop, pause, start, and reset virtual servers inside of VAM (but not physical hypervisor servers). Currently, VAM supports VMware ESX 3.0 and Sun Solaris 10 virtual server technologies.

Using either the VAM toolbar, or by right-clicking a virtual server, you can perform the following functions on a virtual server (VMware VM or a Solaris zone):

- **Start Virtual Servers:**  If the virtual server is currently paused, it will be resumed. This is enabled if the selected virtual servers are stopped or paused.
- **Stop Virtual Servers:**  Stops selected virtual servers which are running.
- **Suspend Virtual Servers:**  Pauses selected virtual servers which are running. Only available for VMware's VMs, not Solaris local zones.
- **Restart Virtual Servers**  : Restarts if any selected objects are virtual servers which are running.

To start, stop, suspend, or reset a virtual server, perform the following steps:

- 1** From one of the VAM maps that display virtual servers (Servers, Network, Storage, SAN), select a virtual server.
- 2** From either the right-click or from the **Manage** menu select **Start, Stop, Suspend**, or **Restart Selected Virtual Server**.

- 3** If you want to see virtual servers and their relationships to their hypervisors, click

**Show Virtual/Physical Containment Relationships**  on the VAM toolbar.

### **Microsoft IIS and WebLogic in the Maps**

In the maps, you can drill into a Microsoft IIS or WebLogic and expand the following objects:

WebLogic process family:

- Web Applications
- EJBs
- JDBC Connection Pools

Microsoft IIS process family:

- Web Site
- FTP Sites

### **Network Map**


The Network Map displays a physical (and virtual) layout of how the elements of an application connect to each other within the network, including the network interfaces on a server and the devices (switches and vSwitches) to which the server is connected. VAM also displays any firewalls and load balancers in your network environment.

In this map you can see process families that are connected over network interfaces on a server, the ports and port groups, VLANs, and listeners that a server's network interfaces are connected to. All network elements are displayed in green.

The Network Map also shows external IP addresses (client IPs) that are connected to an application and the external IP addresses (external dependencies) that an application connects to and depends on. If one of these external connections is an Opsware SAS managed server, then VAM will display the connection as a server. If you want to include

these servers to the scan, click **Refresh Snapshot**  on the VAM toolbar.

Finally, the Network Map displays DNS servers in use by all the servers in the snapshot. If managed server information is known about any of these servers, then you can right-click to Add Devices.

If you want to see the type of service being used for connections between processes (and the devices they run on) in the Network Map, such as http, ssh, telnet, click the Show IPC Service Names in Maps  button.

### **Enhanced Layer 1 (L1) Network Graph**

The Network Map displays not just the devices that are directly connected to a scanned server, but also the devices that are connected to the devices that are connected to scanned servers. Specifically, the Network Map will display the following L1 network information:

- All switches that VAM can detect that are directly connected to scanned servers are shown
- All network devices along the shortest path (based on number of hops) between any 2 servers in the snapshot are shown
- All additional network devices, including:
  - Network devices that were manually added to a business application snapshot
  - Network devices that were shown in a previous snapshot of the business application, but which were not yet manually removed from the Business Application
  - All physical (L1) connections between any two devices that are plugged into each other

### **Network Speed and Duplex Matching**

The Network Map also highlights layer 1 connections that have speed or duplex mismatches between interfaces and network devices, using the following color scheme:

- Green lines and arrows indicate duplex and speed matches.
- Red lines and arrows indicates either a duplex or speed mismatch.
- Gray lines and arrows indicate that not enough information was gathered to determine the duplex or speed matches.

### **ACL and Server Pool Configurations**

For network devices (such as firewalls, load balancers, routers, switches), you can view ACL and server pool configuration (load balancers only) information.

For more information on viewing and comparing ACL and server pool configurations, see “ACLs and Server Pool Configurations” on page 115.

## Virtual Network Devices







If you want to see how virtual network devices are related to the physical devices they run on and are connected to, click **Show Virtual/Physical Containment Relationships**



on the VAM toolbar.

VMware vSwitches are shown alongside virtual machines. Connections between them will appear as a green matching-duplex ethernet connection. If VAM cannot tell from the information gathered in a snapshot the devices a group of processes are connected to, it groups them into a box named "Extraneous."

Network interfaces and devices use the following symbols:

-  **Network Device:** (As shown in the Devices tree) A switch or a vSwitch. In the Network Map, the network device will display a box with the device manufacturer's logo in the title bar. For example, if it is a Cisco switch, the icon in the Network Map will show the Cisco logo in the title bar  .
-  **Network Interface Card (NIC):** When available, a NIC is shown with its IP address and any processes connected to it.
-  **Listeners:** Process families that are listening on or connected to more than one network interface appear multiple times in the Network Map.
-  **Network Device Port**
-  **Virtual LAN**


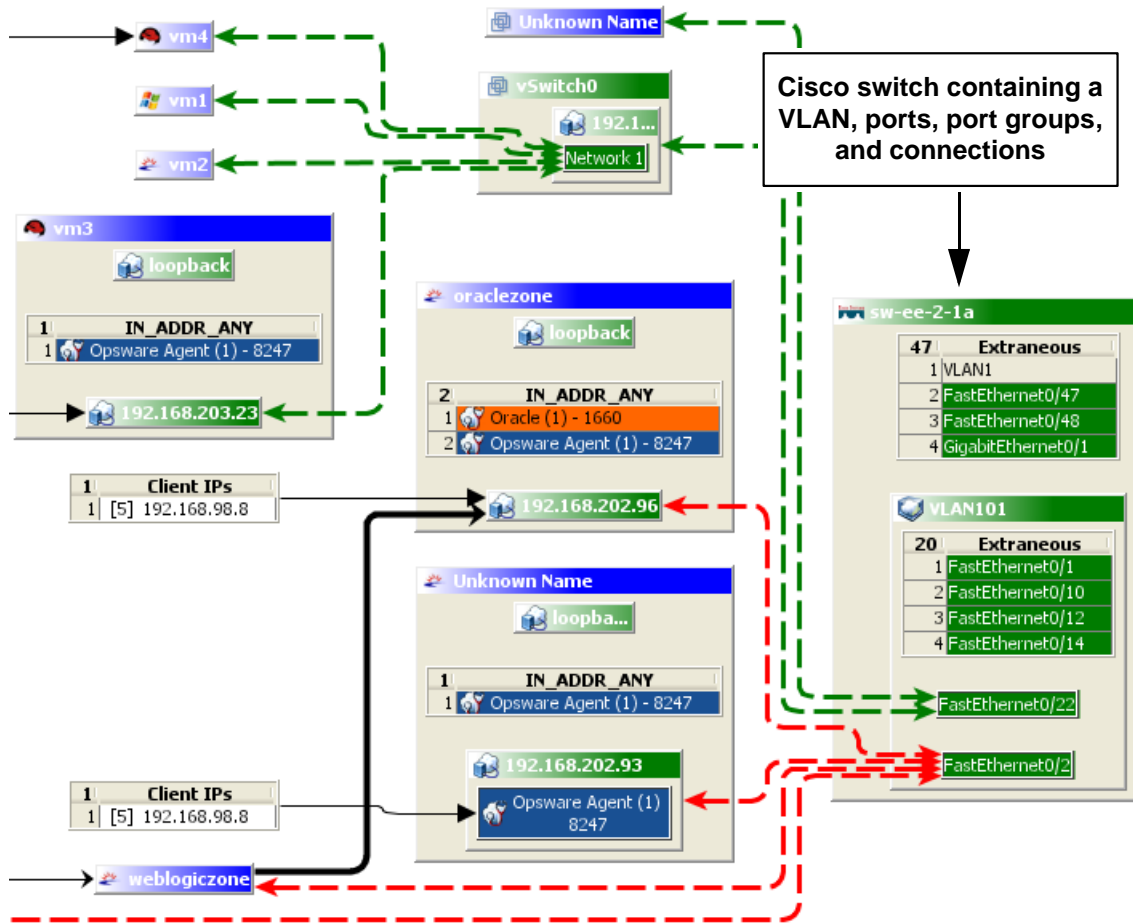
Network devices without connections are also shown. If it is unknown whether the network device is connected to a server or another network device, a warning icon  appears next to it in the Devices tree. See "VAM Scan Error Messages" on page 131.

Figure 1-9 illustrates network devices (green) as shown in the Network Map, with switches and virtual switches, ports and port groups, and network connections (green and red lines).

Figure 1-9: Network Map

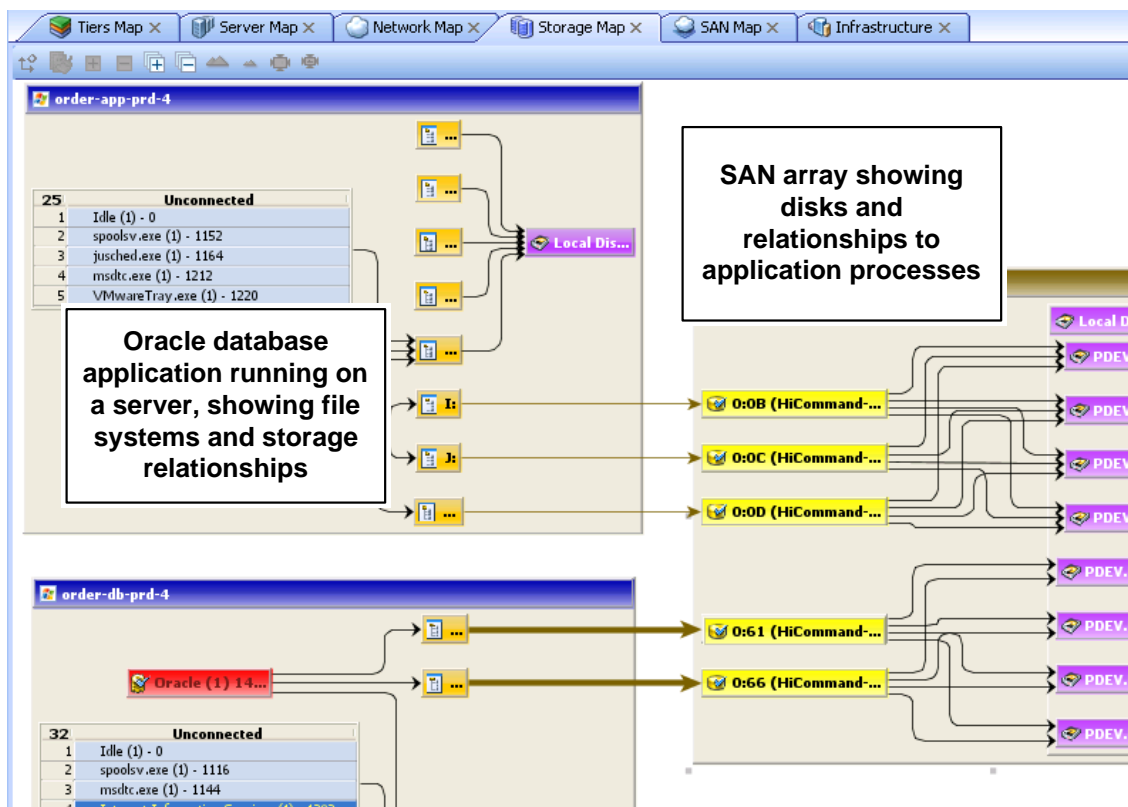


## Storage Map

The Storage Map displays storage dependencies for all of the logical storage elements used by servers you have scanned with VAM.

This map provides a graphical view that displays servers, the process families running on them, including the file systems and local or remote storage devices these files are stored on and served from, either through local disks, NAS filers, or fibre channel disk arrays. The map shows connections between process families and their open files and where they files are stored, with broken connections displayed in red lines.

Figure 1-10: Storage Map Showing Oracle Database Mapping to SAN Array



In this map you will see the following data:

- **Servers:** Shows all process families and their relationship to file systems (through all open files) on each server. These file systems will be served from NAS filers, local discs, and/or remote volumes from storage arrays. In the Storage Map you can see how specific applications running on servers interact with files systems and where data involved with these process families are stored – locally or remotely.

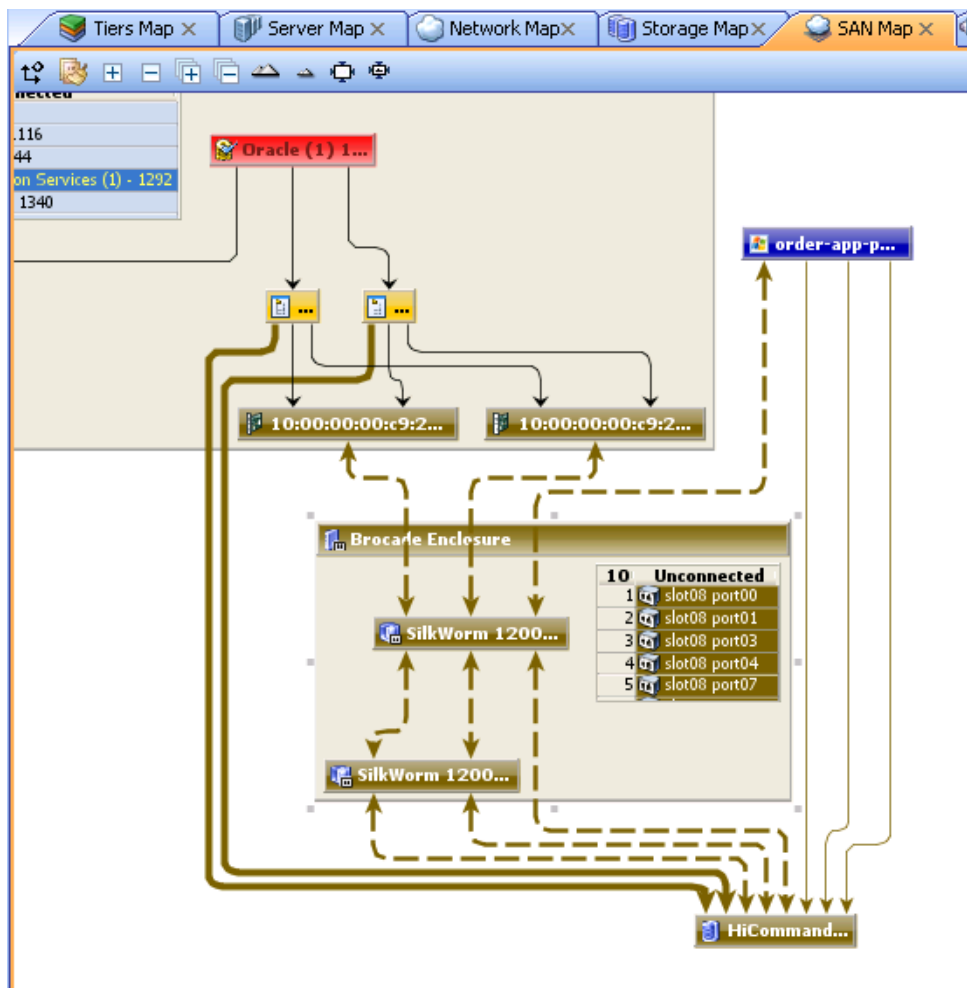


- **File Systems:** Shows file system pathname in the title bar and displays a single list of all open files that live on it. Each file system also shows a connection line between the itself and either local or remote storage.
- **Connection Lines:** Shows connection lines that display how process families connect to the open files list on each file system on which they have open files. Connection lines also tie NFS-mounted file systems to the exported volumes on their hosting NAS filers, as well as the remote storage array volumes that they are hosted on.
- **LUN Mapping:** Displays mapping from a server to a disk array. Multipath link (shown in brown) will have thicker lines based on the number of paths to the storage device (thicker lines mean more paths). For links that are multipath, where at least one of the paths is down, the line will be colored red.
- **Local Disks:** Groups all local disks into a Disk element, which can be expanded to show other disks.
- **SAN Arrays:** Shows storage disk arrays that are mapped and in use by the servers in the snapshot. Storage arrays that are scanned but are not in use by any of the applications show a special icon to distinguish them from those arrays in use. If you expand the SAN Array, you only see volumes in the array that are LUN mapped to the server. SAN Array backups are also displayed.
- **NAS Filers:** Displays all exported file systems stored on NAS Filers (which appear in a brown box) that are in use by servers included in the snapshot. Those filers that are scanned but are not in use by any of the applications have a special icon to distinguish them from those files that are in use. Other servers and devices that use disk on a filer are displayed by way of an exported file system, and a scroll list of other consumers is shown which points to that disk.

## SAN Map

The SAN Map shows a superset of the Storage Map, including a graphical view of the fibre channel storage area network involved in a VAM snapshot, including all servers and their Fibre Channel Adapter (including fibre channel ports), and each adapter's connections to switches in the SAN.


Figure 1-11: SAN Map with Fibre Channel Switches Routing Storage Traffic from Server to SAN Array



The SAN Map displays the following data:

- Servers and their Fibre Channel Adapters (including fibre channel ports), and each adapter's connections to switches in the SAN.
- All SAN arrays and NAS filers included in the snapshot, with physical connections shown between servers, switches, and storage devices.

- File systems (and all open files) and how they rely on remote or local storage.
- LUN mapping between file systems and storage devices (shown in brown solid lines)
- SAN Switches. Each switch, when expanded, shows all ports being used. Any ports not in use in this snapshot are collected in a scrollable list of Extraneous ports.
- Virtual SAN Switches and Ports. Click **Show Virtual/Physical Containment**

**Relationships**  on the VAM toolbar to show virtual switches inside of their physical switch parent.

- Logical connections between devices are represented by solid lines. If a connection is down or has a mismatch, the line is red. If the connection is working and has no mismatch, the line is black. Thickness of line indicates the number of connections.
- Actual physical SAN connections (Fibre Channel cables) are shown in brown dashed lines, while LUN mapping connections between a server and a disk array is shown with solid brown lines.

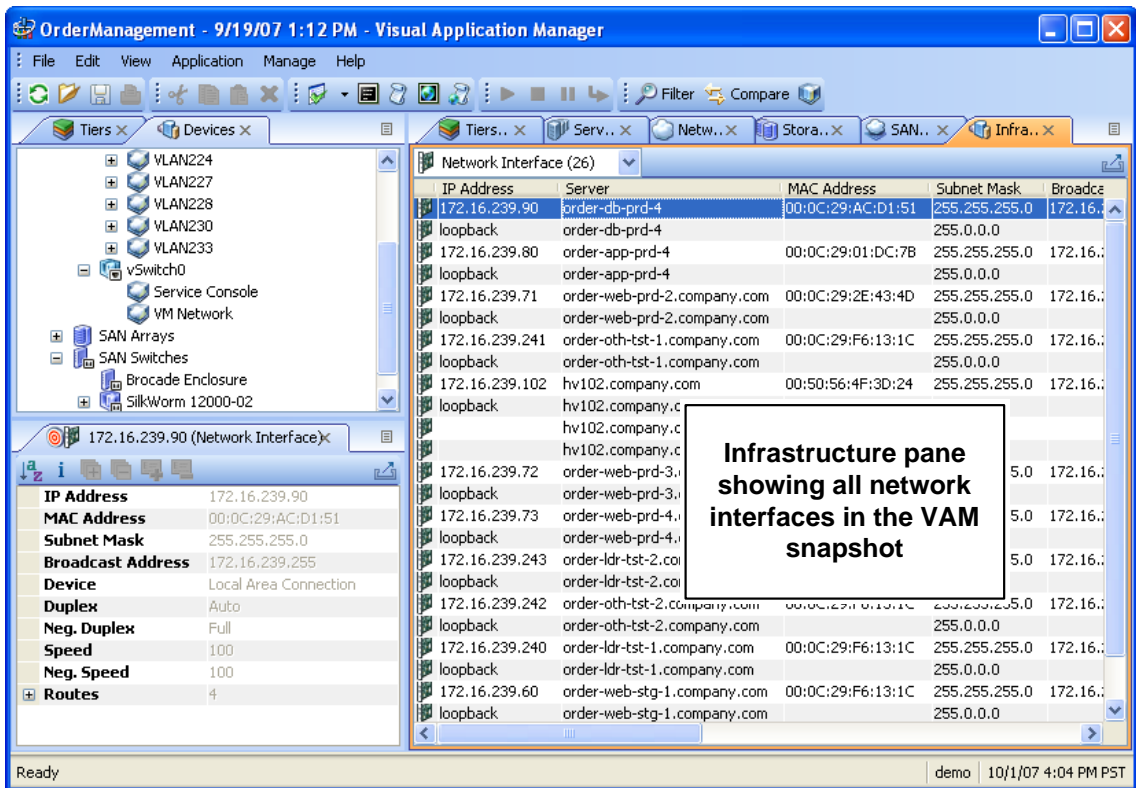
### **Infrastructure Pane**

The Infrastructure pane provides detailed inventory and infrastructure information related to devices you have scanned in VAM. This pane provides a flat list of various objects found in the VAM snapshot and lists detailed properties for each.

Using the drop-down list the Infrastructure pane to filter categories of objects and their properties, to view and sort objects such as, all servers in the snapshot, all SAN arrays in the snapshot, and so on. You can then sort the columns to compare attributes of each item in the list. For example, you can select Servers and look at the load average for all

servers discovered in the snapshot, or sort load average between servers. Or, you can view compliance policies are in use by all the servers or network devices in your business application.

Figure 1-12: Infrastructure Pane in VAM Showing all Network Interfaces



Depending upon what you have scanned in VAM, the Infrastructure pane displays the following types of object categories:

- Disks
- Fibre Channel Adapters
- Fibre Channel Ports
- File Systems
- NAS Filers
- LUN Volumes
- NFS Exported File Systems
- Network Devices

- Network Interfaces
- Network Port
- Compliance Policy
- Process Family
- SAN Arrays
- SAN Switches
- SAN Zone
- Servers
- VLAN

### **Symbols Used in Maps**


VAM uses a variety of symbols in the Network Map, Server Map, and Tiers Map, such as lines, arrows, diamonds, and so on. This section explains VAM map symbols in the following topics:




- Process and Links Connection Symbols
- Map Process and Network Connection Symbols
- Exporting a Map to .gif, .jpg, .pdf, or .svg
- Printing a Map
- “Source” and “Comparison” Snapshot

### **Process and Links Connection Symbols**

In the VAM maps, lines and arrows represent connections between process families. In some cases, VAM will know both the source of the connection and the destination. In other cases, VAM may only know either the source or the destination of the connection.

To represent these process connection relationships, VAM uses the following lines and arrows:

- **Source Unknown – Diamond:**  An arrow with a diamond at its source indicates that VAM does not know the source of the connection. If there is a solid line from a process connection source (in other words, it doesn't show a diamond) then VAM knows the connection source.

- **Destination Unknown – Hollow Arrow:**  A hollow arrow represents an inbound connection to a process family destination that is unknown.
- **From Remote IP – Solid Arrow:**  A solid arrow represents an inbound client connection from a remote IP, such as TCP or UDP, where the destination process family is known.
- **From Process – Lined Arrow:**  A lined arrow represents a connection from a process family where the destination process family is known.


### **Map Process and Network Connection Symbols**


In the VAM maps, lines represent the following connections between devices:


- **Client link:** An internal connection that is labeled by the client IP address.
- **Process link:** A collection of TCP or UDP connections between process families. This link displays processes that provide a network service, such as listening for network connections, and processes that have a connection to another process or server.
- **Layer 1 connection:** A physical link between a server's network or storage interfaces and switch ports/switches. Layer 1 connections are indicated by colored, dashed lines in the map:

 A green dashed line indicates that there is no duplex mismatch.

 A red dashed line indicates that there is a duplex mismatch.






 A gray dashed line indicates that there may or may not be a duplex mismatch because at least one value is unknown.

 A brown solid line shows LUN mapping connections from a server to a disk array (in the Storage Map)

 A brown dashed line shows physical fibre channel cable connections (SAN and Storage Maps only).

- **Line Thickness:** The thinness or thickness of the line represents the number of connections associated with the link. A smaller number is indicated by a thinner line and a larger number is indicated by a thicker line, as illustrated in Figure 1-13


Figure 1-13: Line Thickness and Process Connection Relationship

	1-4
	5-16
	17-64
	65-256
	257+

### Exporting a Map to .gif, .jpg, .pdf, or .svg

You can export a map to a .gif, .jpg, .pdf, or .svg file for use in other applications where you can annotate the drawing or map the exported file in a web browser.

To export a map to an image file, perform the following steps:



- 1 From the **File** menu, select **Save As Image**. (Or, click the Export View  button at the upper right corner of the map.)
- 2 Select a directory where you want the file to be located.
- 3 Enter a file name that includes either .gif, .jpg, .pdf, or .svg as the file name extension.
- 4 Click **Save As Image**.

### Printing a Map


You can print a map on single and multiple sheets of paper, and you can also title the map for better presentation.

If the map you want to print is very dense and complex, you can make adjustments by zooming in and zooming out, and by creating rows and columns that will break the map up over several pages. Doing this enables you to print the map on multiple sheets of paper, thus increasing the map's readability.


To adjust the map before you print:

- Click Zoom In  or Zoom Out  to increase or decrease the size of the map before you print.
- Enter a title for the map

To print a map, perform the following steps:

- 1** From the **File** menu, select **Print** or select the  toolbar icon.
- 2** (Optional) In the Print window, specify page setup and printer options, including a title that you want to appear on the printed map.
- 3** Click **Print**.

### **Showing IPC Service Names in Maps**

If you want to see the type of service being used for connections between processes (and the devices they run on) in the Network or Server map, such as http, ssh, telnet, click the Show IPC Service Names in Maps  on the VAM toolbar. The service name will appear on each connection line.

To show IPC service names in the Server or Network map, perform the following steps:

- 1** Inside of VAM, select either the Server or Network map.
- 2** From the **View** menu, select **Show IPC Service Names in Maps**.

Or

- 3** Click Show IPC Service Names in Maps  on the VAM toolbar

## **VAM Properties**

VAM displays detailed properties for components selected in the Devices tree, Tiers tree, the Infrastructure pane, and any of the maps.

The information displayed in the Properties pane varies depending on the component type, such as a server, network device, storage device, process family, tier, application or storage signature, and links. Depending on the type of object you select, this page



includes the number of users logged in, load average, swap usage, memory usage, application components, network devices, network ports, VLANs, tiers, links, and so on. It shows the MAC addresses for each network interface, as shown in Figure 1-14.


Figure 1-14: Properties Pane for a Server

<b>Name</b>	order-web-prd-4.company.c...
<b>Description</b>	
<b>Opware ID</b>	370001
<b>Operating System</b>	Red Hat Enterprise Linux AS 3
<b>Kernel</b>	Linux 2.4.21-47.EL i686
<b>Agent Version</b>	34.0.0.89
<b>Last Reported Time</b>	8/30/07 11:05 PM
<b>Uptime (1/1000 sec)</b>	2 days 1 hour 7 minutes
<b>Codeset</b>	UTF-8
<b>Users</b>	1
<b>1 Min Load Avg</b>	5.62
<b>5 Min Load Avg</b>	5.78
<b>15 Min Load Avg</b>	5.73
<b>Free Memory (bytes)</b>	27,852 MB
<b>Total Memory (bytes)</b>	248,184 MB
<b>Swap Used (bytes)</b>	189,231 MB
<b>Swap Total (bytes)</b>	1,004 GB
<b>Swap In (bytes)</b>	268 bytes
<b>Swap Out (bytes)</b>	232 bytes
<b>Process Families</b>	40
<b>Extended Process Fa...</b>	0
<b>VM Name</b>	gs2-1-apache-4
<b>VM State</b>	Powered On
<b>Host</b>	
<b>DNS Servers</b>	2

### Exporting Properties Information to .csv

You can export the information contained in any Properties pane to the .csv format, which enables you to view VAM object properties inside spreadsheet applications.

To export a Properties pane to .csv, perform the following steps:

- 1** From inside of a VAM map or tier, select an object.
- 2** From inside the Properties pane for the selected object, click the Export View  button.

- 3** In the Export to CSV window, choose a location and enter a file name (with the .csv extension), and then click **Save**.

### **Tiers Tree: Tiers, Process Family, Signature Properties**

VAM displays property information about the following VAM elements in the Tiers Tree and maps:

- Process Family Properties
- Tier Properties
- Application Signature Properties
- Storage Signature Properties

### **Process Family Properties**

The Properties Pane for a process family displays the following information:

- **Name:** The name of the controlling process of the family.
- **Family ID:** The unique ID given to the process family.
- **Extended Family:** The name of the extended process family, if the selected process family belongs to an extended process family.
- **Max. Resident Memory:** The maximum permanent memory used by the process family, in bytes.
- **Max. Virtual Memory:** The maximum permanent memory used by the process family, in bytes.
- **Max. Run Time:** The length of time the process has been running.
- **Total CPU Time:** The total length of time the process used CPU resources.
- **Max CPU Utilization:** The total amount of CPU resources used by the process.
- **Group ID:** The group ID of the process family on Unix and the session ID on Windows.
- **Listeners:** The interface and port for each listener.
- **Incoming connections:** The connections incoming to the process family, grouped by process family (if known, the IP address otherwise) and interface.
- **Outgoing connections:** The connections outgoing from the process family, grouped by process family (if known, IP address otherwise) and interface.

- **Modules:** The shared libraries associated with the process family. These include DLLs on Windows and shared object files on Unix.
- **Open files:** The files that the process family currently has open.
- **Software Packages:** The packages associated with the files that the process family has open.
- **Processes:** The number of individual processes in the process family. For each process, the following information is displayed:
  - **PID:** The process ID.
  - **User:** The user ID the process is running as.
  - **Command line:** The command line used to start the process.
  - **Path:** The path to the process binary.
  - **Memory statistics:** The percentage of physical memory consumed by the process, the resident size (in bytes) of the process and the virtual size (in bytes) of the process.
  - **Run time:** The time (in milliseconds) that the process has been running.
  - **CPU Statistics:** The CPU time accumulated by the process and the percentage of CPU consumed by the process since it began.
  - **Environment:** The name and value of each environment variable in the process environment.

### ***Tier Properties***

The Properties Pane for a tier displays the following information:

- **Name:** The name of the application tier as displayed in the Tiers tree.
- **Application Tiers:** The number of subtiers that are currently recognized in the tier.
- **Application or signatures:** The number of application signatures that are currently recognized in the tier.
- **Server Filter:** The servers that are associated with this tier. Only matching servers are filtered for matching application components.

### ***Application Signature Properties***

The Properties Pane for an application displays the following information:

- **Name:** The name of the application component as displayed in the Tiers tree.
- **Alias** (Optional): The name of the application component as displayed in the different views. The name will be shown as an alias if not defined.
- **Families:** The number of process families recognized as this application component.
- **Process Name:** The process name filter used to recognize this application component.
- **Command Line:** The command line filter used to recognize this application component.
- **Connected To Port:** The port that the server is connected to.
- **Listener Port:** The listen port used to recognize this application component.
- **Executable Path:** The executable path filter used to recognize this application component.
- **Open Files:** The open file filter used to recognize this application component.
- **Modules:** Shared libraries that are associated with the process family, DLL files on Windows operating systems and shared object files on Unix operating systems.
- **Background:** The background color displayed in the different maps.
- **Foreground:** The foreground text color displayed in the different maps.

### ***Storage Signature Properties***

The Properties Pane for a storage signature displays the following information:

- **Name:** The name of the storage component as displayed in the Tiers tree.
- **Alias** (Optional): The name of the storage component as displayed in the different views. The name will be shown as an alias if it is not defined.
- **Remote Volumes:** Number of remote storage volumes
- **LUN ID:** The port that the server is connected to.
- **Listener Port:** The listen port used to recognize this application component.
- **Executable Path:** The executable path filter used to recognize this application component.
- **Open Files:** The open file filter used to recognize this application component.

- **Modules:** Shared libraries that are associated with the process family, DLL files on Windows operating systems and shared object files on Unix operating systems.
- **Background:** The background color displayed in the different maps.
- **Foreground:** The foreground text color displayed in the different maps.

### **Devices Tree: Server and Network Device Properties**

VAM displays property information about servers, network devices, and the connections between them:

- Devices Tree: Server and Network Device Properties
- Properties for Servers and Devices with Compliance Policies
- Virtual Server Properties
- Link Properties for Servers and Network Devices
- Network Devices Properties
- Virtual Switch Properties
- Port Group Properties
- Network Interface Properties Pane



---

If your core has been enabled to display storage devices with ASAS, then a server's properties will also include related storage and more detailed file system information.

---

### **Server Properties**

The Properties Pane for a server displays the following information:







- **Name:** The host name of the server.
- **Opsware ID:** The Opsware unique identifier for the server.
- **Operating System:** The operating system of the server.
- **Kernel:** The kernel version of the operating system (when applicable).
- **Agent Version:** The version of the Opsware Agent that enables the server to be managed and scanned.
- **Last Reported Time:** The most recent time that the Opsware Agent communicated with the Opsware core.

- **Uptime:** The length of time the server has been powered on.
- **Codeset:** The character encoding for the server's locale.
- **Users:** The number of users that are currently logged in.
- **Load averages:** 1-minute, 5-minute, and 15-minute load averages. The load average for servers running a Windows operating system will display unknown because it is not supported by Microsoft.
- **Memory usage:** The total free memory.
- **Swap usage:** The total used swap and swap in/out activity.
- **DNS Servers:** All configured DNS servers for the selected server.
- **Virtual Machines/Zones:** If the selected server is a hypervisor (Solaris Global Zone or VMware ESX server), you can expand the list and view all zones (Solaris) virtual machines (VMware ESX). Each virtual machine or zone will display its own server properties. For more information, see "Virtual Server Properties" on page 90.
- **Routes:** All configured static routes on the selected server.
- **Interfaces:** The number of network interfaces. For each interface on a server, the following information is displayed:
  - MAC address
  - Broadcast address
  - Subnet mask
  - Device
- **FCAs:** All Fibre Channel Adapters (HBAs) installed on the selected server. (For ASAS-enabled cores only.)
- **File systems:** All files systems in use on the selected server. For each file system, properties include: drive letter, mount point, mount options, type of file system, logical block device used by the file system, amount of free space, percent used, and associated device for each file system.
- **Disks:** All physical disks installed on the selected server.

### **Properties for Servers and Devices with Compliance Policies**

For servers or network devices that have compliance policies associated with them (Software, AppConfig, Patch, Audit, Duplex), the server's properties will show a rollup compliance status for all attached policies. You can expand the compliance list to view each individual compliance policy attached to the server.

Each compliance category will display one of the following compliance statuses:

- **Compliant** : The compliance scan ran successfully and the actual server or device configuration matches the criteria defined in the policy.
- **Partial** : The compliance scan ran successfully, but the server or device configuration did not fully pass the compliance criteria defined in the policy.
- **Noncompliant** : The compliance scan ran and the actual server or device configuration did not match the criteria defined in the policy.
- **Scan Failure** : The compliance scan was unable to run.
- **Scan Needed** : The results are unavailable, perhaps because a compliance scan was never run (for example, on a new installation), or the configuration on the server changed since the last time information was reported to the Compliance Dashboard.
- **Scanning** : The compliance scan currently being run.


You can launch the Device Explorer or remote terminal in the SAS Client to view and remediate any compliance discrepancies by clicking on a compliance status link in the properties window. For NAS-enabled cores, clicking a compliance status link launches the NAS Web interface.

For information on launching a Device Explorer, remote terminal, or global shell, see "Adding and Removing Devices in VAM" on page 60.





The actual compliance state you are viewing in VAM may have changed since you last scanned the server or device. To get the most current information, click **Refresh**

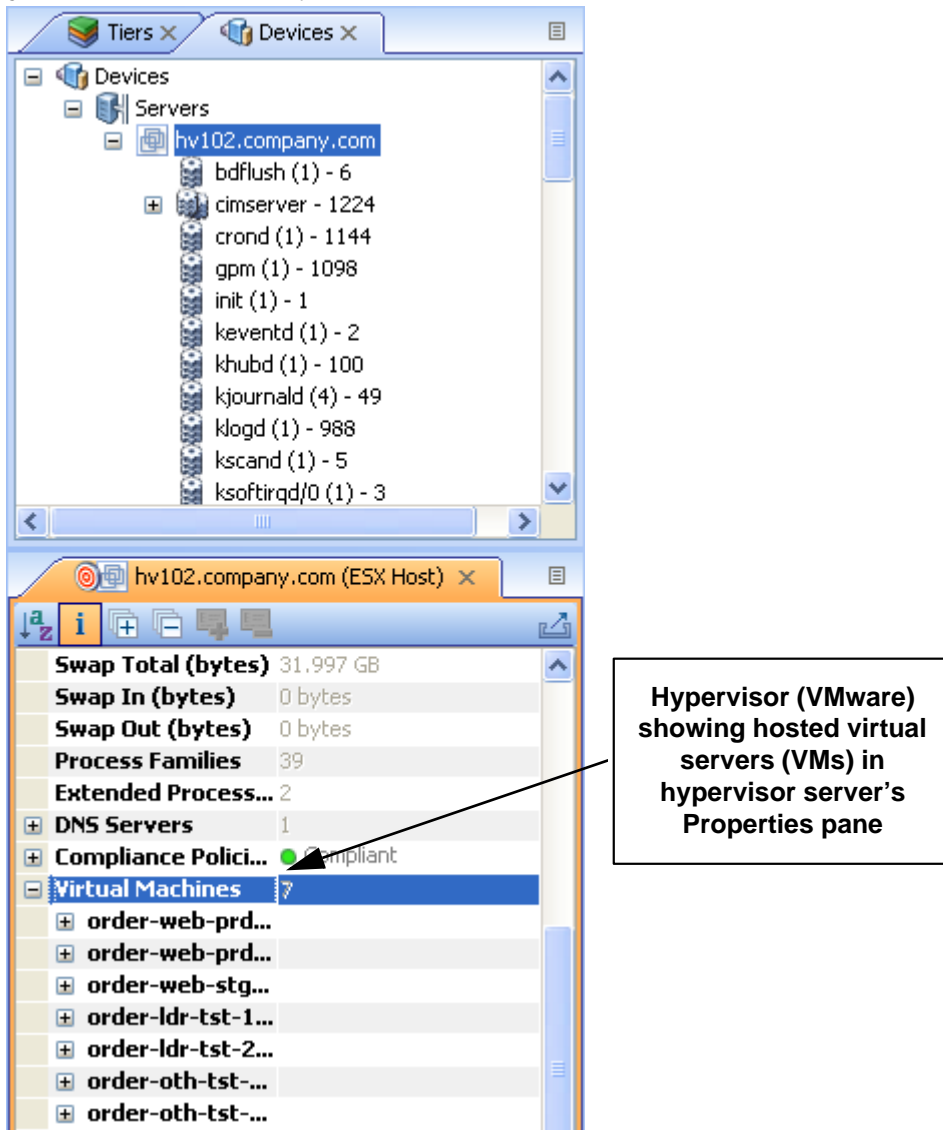
**Snapshot**  on the VAM toolbar. For more information on the Compliance feature, see the *Opware® SAS User's Guide: Application Automation*.

---

### Virtual Server Properties

The properties of a virtual server display all the same information as a physical server except that VMware and Solaris 10 hypervisors will show all hosted virtual servers. You can expand each hosted virtual server and view its properties. Conversely, each virtual server will contain in its properties the hypervisor that is hosting it. Figure 1-16 illustrates this feature.

Figure 1-16: Virtual Server Properties



### **Link Properties for Servers and Network Devices**

The Properties Pane for links between servers, external connections, and network devices displays the following information:

- **Protocol:** The TCP or UDP.
- **Port:** The destination port that is associated with this link.
- **Connections:** The number of connections associated with this link. For each connection, the following information is displayed:
  - **End points:** The process families (if known). IP addresses (if unknown).
  - **Ephemeral port number:** A random port that is assigned by the operating system.

### **Network Devices Properties**

The Properties Pane for network devices (routers, switches, firewalls, load balancers, and so on) displays the following information:

- **Name:** The name of the network device.
- **Last Reported Time:** The date of the last successful snapshot of the network device by Opsware NAS.
- **Manufacturer:** The Vendor that manufactures the network device.
- **Model:** The model number of the network device.
- **Operating System:** The operating system running on the network device.
- **Firmware Version:** The firmware version number for the device.
- **Asset Tag:** The assigned number used for tracking the network device.
- **VLANs:** The total number of VLANs that this network device has.
- **Ports:** The total number of ports that this network device has.
- **ACL Config:** Displays the current list of ACLs configured for the device. You can view or compare ACL configurations by selecting **View ACL Configuration** from the **Manage** menu. This opens the ACL Configuration dialog for the selected network device(s) that have ACLs.

If you select two network devices, right-click, and select **Compare ACL Configuration**, this opens the Compare ACL Configuration dialog for the selected network device(s) that have ACLs. This allows you to compare ACL configurations between two devices in the same snapshot, or when in comparison mode, between the same device in the two

snapshots. For more information, see “ACLs and Server Pool Configurations” on page 115

- **Server Pool:** Displays all sever pool members for the selected device. For load balancers, you can view Server Pool members by clicking the Server Pools link in the Properties pane or selecting **View Server Pool Configuration** from the **Manage** menu.

If you select two load balancers in the Devices tree, right-click, and select **Compare Server Pool Configuration**, you can view the Compare Server Pool Configuration window to see any differences in configuration. For more information, see “ACLs and Server Pool Configurations” on page 115

- **Compliance:** For network devices that have compliance policies associated with them (Duplex), the properties will display its compliance status. For information on Compliance statuses, see “Properties for Servers and Devices with Compliance Policies” on page 87.

### **Virtual Switch Properties**

Virtual Switch properties will display the following information:

- **Port Groups:** These can be expanded to view port groups configured for the selected vSwitch.
- **Network Interfaces:** These can be expanded to view network interfaces assigned to the selected vSwitch.

### **Port Group Properties**

Port group properties will display the following information:

- **Port Group Name:** The name of port group
- **VLAN ID:** The VLAN ID of port group. This is optional in the VMware management user interface.

### **Network Interface Properties Pane**

The Properties Pane for a network interface displays the following information:

- **IP Address:** The IP address that is associated with a network interface.
- **MAC Address:** The Media Access Control ID that is associated with a network interface.
- **Subnet Mask:** The subnet that is associated with a network interface.
- **Broadcast:** The broadcast address that is associated with a network interface.

- **Device:** The device that is associated with a network interface.
- **Duplex:** The configured duplex (if it can be collected).
- **Negotiated Duplex:** The negotiated duplex (if it can be collected).
- **Speed:** The configured speed in Mbps (if it can be collected).
- **Negotiated Speed:** The negotiated speed in Mbps (if it can be collected).

### **Storage and SAN Properties**

VAM displays property information about the following storage- and SAN-related elements:

- Server Properties with ASAS
- File Systems
- LUN ID
- SAN Array Properties
- NAS Filer
- Fibre Channel Switch
- SAN Array Disk Volumes
- Fibre Channel Adapter (Host Bus Adapter)
- Fibre Channel Port
- Storage Signature

### **Server Properties with ASAS**

Servers display the following attribute information in the Properties tab (these elements are visible only in the Storage or SAN Maps):

- **FCA:** Any server that is attached to a SAN storage device will have an HBA (Host Bus Adapter) which VAM labels FCAs (Fibre Channel Adapters). You can expand category in the Properties pane and see a list of all FCAs, each of which can be expanded to display all their ports.
- **File systems:** VAM displays details for ASAS-enabled cores, such as mount options, Type (ext3, NFS), and so on.
- **Disks:** Number of and names of each disk on the server, including type, manufacturer, and disk capacity information.

## **File Systems**

File systems can exist on both servers and NAS Filers (NetApps). VAM displays the following attributes for file systems in the Properties pane:

Mount point/drive letter/export path:

- Capacity
- Free capacity
- Percent used

For server file systems, the following information is displayed:

- Device
- Mount options
- Type (e.g. ext3, NFS, etc)

For NetApp file systems/Qtrees, the following information is displayed:

- Aggregate
- Plexes
- RAID groups

## **Disk Properties**

Disks from SAN arrays, NAS filers, and servers will display the following attributes in the Properties pane:

- Capacity
- Manufacturer
- Model
- Type
- Status
- Serial number
- Firmware version
- Device (only on servers)

### **SAN Array Properties**

SAN Arrays display the following attribute information in the Properties pane:

- Name
- Description
- Last Reported time
- Manufacturer
- Opware ID
- Disks
- Volumes
- Ports

### **NAS Filer**

NAS Filers display the following attribute information in the Properties pane:

- Name
- Description
- Hostname
- Last Reported Time
- Manufacturer
- Model
- Opware ID
- Serial number
- Operating system version
- Hardware version
- Disks
- Filer volumes/Qtrees
- LUN volumes
- Ports

### **Fibre Channel Switch**

Fibre Channel Switches will have the following attribute information in the Properties pane:

- Name
- Description
- Last Reported Time
- Manufacturer
- Model
- Opware ID
- Serial number
- Firmware version
- Hardware version
- Ports
- Virtual switches



---

If a switch is a director-class switch, it could possibly have virtual switch children – these can be expanded as well.

---

### **SAN Array Disk Volumes**

Volumes on SAN arrays (and NAS Filers acting as arrays) will have the following attributes:

- Name
- Description
- Capacity
- LUN ID

### **Fibre Channel Adapter (Host Bus Adapter)**

Fibre Channel Adapters (known generically as HBAs) display the following attributes in the Properties pane:

- Node WWN



- Manufacturer
- Model
- Serial number
- Driver version
- Firmware version
- Hardware version
- Ports

### **Fibre Channel Port**

Fibre Channel Ports display the following attributes information in the Properties pane:

- World wide name
- Description
- Port number
- Status
- Zonesets
- Zones
- Fabric

### **Storage Signature**

Storage signatures match remote storage volumes and display the following attribute information in the Properties pane:

- LUN Name
- LUN ID
- Exported path
- Storage System Manufacturer
- Storage System Model

## **SAN Link Properties**

VAM displays the following links or connections between storage signatures in the SAN Map to illustrate how these signatures are related and connected:

- Fibre Channel Link
- LUN Mapping Link
- NFS Mount Link

### ***Fibre Channel Link***

Fibre channel links represent physical fiber cables connecting two ports. The Properties pane displays attributes information for endpoint ports.

### ***LUN Mapping Link***

LUN mapping links represent the logical connection between a host LUN volume and the corresponding SAN array volume. LUN mapping links display the following attributes information in the Properties pane:

- Number of paths (indicated by thickness)
- Whether any paths are down (in red)
- Server file system consumer
- Array target volume

### ***NFS Mount Link***

NFS mount links represent a dependency of a server file system on a remote NAS Filer) export and display the following attributes information in the Properties tab:

- Exported directory or qtree
- Mount point directory

## **VAM Options**

For VAM, you can specify the following options:

- Virtualization Settings
- Scan Time-Out Preference
- Discovery Settings

- Reset All Settings

### **Virtualization Settings**

You can configure SAS Client options that allow you to choose whether or not you want to perform a scan on any virtual servers or hypervisors related to the virtual server you want to open in VAM.

For example, if you want to visualize a VMware virtual machine (VM) or Solaris zone in VAM, by default you will be asked if you also want to scan any virtualization relationships – in other words, the system asks if you want VAM to also scan the hypervisor that is hosting the selected virtual server. Depending upon the virtual server you select, VAM might have to scan several related virtual servers in order to visualize a single virtual server in VAM.

Conversely, if you select a hypervisor to open in VAM, you will also be asked if you want to scan any virtualization relationships – in this case, VAM would need to scan all of the hosted virtual servers, which could take a long time to perform.

By default, VAM will always ask you if you want to scan virtual relationships, but you can set your own default behavior for scanning related virtual servers with the following virtualization options:

- Ask each time if you want to scan related virtual and host servers.
- Always scan related virtual and host servers.
- Never scan related virtual and hypervisor servers.

To change the virtualization settings, perform the following steps:

- 1** From the **Edit** menu, select **Options**.
- 2** In the Set Options window, in the Views pane, select Visual Application Manager.
- 3** Specify your desired Virtualization Settings, then click **OK** when you are finished.

### **Scan Time-Out Preference**

VAM is optimized to scan a maximum of 50 servers. A number of factors affect the time it takes for a scan to complete, including the load on the scanned servers and the load on Opsware. The default scan time-out is set to 300 seconds. You can reset this time-out value to a minimum of 30 seconds or to a maximum of 3600 seconds.

To change the scan time-out, perform the following steps:

- 1** From the **Edit** menu, select **Options**.
- 2** In the Set Options window, in the Views pane, select Visual Application Manager.
- 3** In the Scan Timeout section, move the slider to increase or decrease the number of seconds at which you want the scanning process to stop.
- 4** Click **OK** to save your changes or click **Cancel** to close the window without saving your changes.

### Discovery Settings

If servers are scanned and it is determined that they are dependent on external IP addresses, when this option is selected VAM will attempt to determine which servers or network devices those IP addresses refer to.

Keep in mind that this could cause scan time to increase, depending on the numbers of servers you selected for the scan and how many remote dependencies are discovered.

For recurring background business application snapshots, this detection will always be done and cannot be turned off.

### Reset All Settings

Restores all VAM settings to their defaults.

---

You can also access these options from inside the SAS Client by selecting **Options** from the **Tools** menu.

---

## Accessing Servers and Devices From VAM

To help you troubleshoot and take action for server and application errors, VAM gives you easy access to servers and devices through the following methods:

- Opening a Device Explorer
- Opening a Remote Terminal
- Opening a Global Shell

## Opening a Device Explorer

To view detailed information about a server or a device (network or storage device) using the Device Explorer, perform the following steps:

- 1** In one of the VAM maps, select one or more servers.
- 2** Right-click and then select **Open with Opsware SAS Client** to open a Server Explorer for each selected server.

See the *Opsware® SAS User's Guide: Server Automation* for information about how to use the Device Explorer.

## Opening a Remote Terminal

The Remote Terminal enables you to log into devices (servers and network devices) and run native commands.

To open a Remote Terminal from VAM, perform one of the following tasks:

- 1** In one of the VAM maps, select one or more servers.
- 2** Right-click and then select **Open Remote Terminal** to open the Select Remote Login window.
- 3** In the Login column, select a login ID from the drop-down list, such as root or LocalSystem, or any of the user logins that might be configured.
- 4** Click **OK** to open a Remote Terminal for each selection.

See the *Opsware® SAS User's Guide: Server Automation* for information about using utilities in a Remote Terminal.

## Opening a Global Shell


You can use the Global Shell feature to navigate between servers and connected network devices by tracing their layer 1 connections in the `/opsw/Servers/@` and `/opsw/Network/@` directories in the OGFS.

In the OGFS, you can also run scripts to perform the following tasks:

- Find servers and network devices.
- Find all servers that are connected to a certain switch.
- Display the network interfaces of a certain server.
- Get the IP addresses of all devices.

- Compare two files to identify changes made, such as what changes were made to a device configuration (.conf) file.
- Change device details, such as the snmp-location.

To launch the Global Shell, perform one of the following tasks:

- From the **File** Menu, select **Global Shell**.
- Select the  toolbar icon.

See the *Opware® SAS User's Guide: Server Automation* for information about how to use Global Shell.

## Creating Business Application Definitions

A business application definition allows you to transform a data display that contains extraneous and hard-to-understand information into a focused and easy-to-understand view of the relevant data. Based on business application tiers and application and storage signatures that you create, VAM recognizes actual application processes and storage devices and displays them in the Tiers Map according to any visual customizations you make to them.

You create business application definitions in order to recognize processes by giving them meaningful names and appearances (colors). You also use business application definitions to define the logical tiers of an application and display application and storage signatures according to the tier in which they reside.

See “Signature Evaluation Order” on page 107 for information on the order in which VAM scans application signatures and matches them to processes and process families on servers.

For information on understanding and creating application definitions, see the following topics:

- “Business Application Tiers” on page 104
- “Creating a Tier” on page 105
- “Application and Storage Signatures” on page 106
- “Creating an Application or Storage Signature” on page 109

## Business Application Templates

When you first scan servers with VAM and visualize them, the Tiers pane is empty –it has no business application definitions until you create them. (There are, however, some predefined commonly used default applications built into the product, such as Apache, WebSphere, and so on, contained in the Default Signatures folder.) Once you create and define an application definition with tiers and signatures, you can save the application definition as a template, which can be reused by yourself or others on your team.

You can also set an application definition to use as the default template, so that whenever you open VAM, it always opens using the application definitions saved in the default template. If you make changes to an application that is based upon a template, and do not wish to save the changes, you can restore the default template.

### ***Setting a Default Application Template***

If you have made changes to the application definition and want to set this as the default, select **Set as Default Template** from the **File** menu.

### ***Resetting the Default Application Template***

If you have made changes to the application definition and want to restore the previously saved default application, select **Reset Default Template** from the **File** menu.

### ***Importing an Application Template***

If you would like to import an application template that has already been saved, from the **File** menu, select **Import Template** and select the template to import.



Importing an application template will replace any existing application definitions in your current VAM session.

---

## Creating Business Application Contacts




In VAM, you can create a list of email contacts – and send emails to contacts on this list – by adding email and contact information to the top level tier of a business application.

You can create groups of contacts, and add to each contact such information as email, instant messenger IDs, phone number, and so on. From the **File** menu, select **Send Email**, and you can email any of the contacts you have added to the Business Application

VAM also will display any email addresses configured on network devices scanned by VAM. (Most network devices have an internal configuration setting such as “sysContact” that allows them to associate an email address for the owner of the device.)

You can add contacts through the properties of the Business Application, and then launch emails for all email contacts listed.

To create a new contact for your business application, perform the following steps:

- 1** From the Tiers pane, select the top-level Business Application  icon.
- 2** Select the Properties for the business application, and then click **Add Contact**  (at the top of the Properties pane).
- 3** To enter information for a contact, double-click in the field to the right of each entry. After an entry line is filled, press Return to enter the information. If you want to be able to send emails to a contact, be sure to enter the contact's email address.
- 4** To delete a contact, select the contact in the Properties pane and click **Remove Contact** .

### **Sending Email to Business Application Contacts**

You can send email to any business application contact that has a well-formed and valid email address.

To send an email to a contact, perform the following steps:

- 1** From the **File** menu, select **Send Email**.
- 2** In the Email Contacts window, expand the business application. Each contact that has its name check marked will be added to the email. If you do not want to send an email to one of the contacts, select the check mark next to the name.
- 3** Click **Compose** to write and send the email. (VAM will launch whatever email client you have installed and configured on your local system.)

### **Business Application Tiers**

Business application tiers provide an architectural framework to organize and display application and storage signatures. You can add, edit, delete, cut, copy, and paste tiers in the Tiers tree. You can paste a tier before or after a selected position in the Tiers tree to



rearrange the order. The order of tiers (and the signatures they contain) is significant because it affects the order that the process families are assigned to signatures. (For more information, see “Signature Evaluation Order” on page 107.)

Tiers that do not have any signatures (including sub-tiers that do not contain any recognized process families) are not drawn in the map. If any tiers have application signatures that do not recognize any process families, they and their ancestors are represented with warning icons in the tree and by yellow title bars in the view. This allows you to quickly identify signatures that should be running but are not.

### **Creating a Tier**

To create an tier in the Tiers tree, perform the following steps:

- 1** In the Tiers tree, select either the top-level business application node or a tier, right-click, and select **New Tier**.
- 2** The Properties pane for the tier becomes active and you can edit the tier’s properties, such as, give the tier a name.

### **Deleting a Tier**



To delete an tier from the Tiers tree, perform the following steps:

- 1** In the Tiers tree, select a tier.
- 2** From the **Edit** menu, select **Delete** or right-click and then select **Delete** (or, click the delete toolbar button ).

### **Cutting and Copying a Tier**


You can cut and copy a tier to the clipboard. After you do this, you can paste the tier before or after a selected position in the Tiers tree to rearrange the order. The order of application tiers (and the signatures they contain) is significant because it affects the order that the process families are assigned to signatures.

To cut and copy an application tier in the Tiers tree, perform the following steps:

- 1** In the Tiers tree, select a tier.
- 2** From the toolbar select either the  icon or the  icon, or right-click and select **Cut** or **Copy**.

## Pasting a Tier

To paste a tier in the Tiers tree, perform the following steps:

- 1 Select a tier in the Tiers tree and then select the Paste icon . The tiers that you cut or copied to the clipboard will be appended to the selected tier's children. When you select a signature in the Devices tree, the Paste icon will be disabled.

## Application and Storage Signatures

Application and storage signatures are organized and displayed in the Tiers pane inside of the VAM application window. A signature always lives inside of a tier.

Application or storage signature contains the following data:

- A name.
- A signature, which is a set of rules that users provide and that VAM uses to identify a process family or storage mapping. For application signatures, these rules use data such as process name, command line, listen port, executable path, and so on. For storage signatures, these rules include either a LUN volume or an NFS File System.
- Object properties such as name, color, and whether this object is used by default each time the user opens up an application.

You can add, edit, delete, cut, copy, and paste signatures in the Tiers tree. You can paste a signature before or after a selected position in the Tiers tree to rearrange the order. The order of signatures (and the tiers that contain them) is significant because it affects the order that the process families and storage mappings are assigned to signatures.

VAM comes with a set of predefined default signatures that recognize a variety of commonly used application process families, such as Apache HTTP, Microsoft IIS, WebLogic, JBoss, Oracle, and so on. So, if your server has any of these applications installed, VAM will be able to recognize and display them in the Tiers tree and the maps.

VAM also includes a set of Opsware signatures, such as the Opsware Agent, Opsware Build Manager, NAS Syslog Server, Opsware Command Engine, and so on. Many of these signatures will appear only if you use VAM to scan the server or servers that the Opsware core is installed on, while others, like the Opsware Agent, will appear on all reachable managed servers.

## List of Application Signature Discovery Properties

The full list of application signature properties that can be used to find applications and their process families include:

- Process Name
- Command Line
- Connected To Port
- Listener Port
- Executable Path
- Open Files
- Open Modules
- Environment Variable Name and/or Value

The means by which an application signature discovers applications and process families for environment variables allows you to match application signatures by the name of an environment variable, its value, or both.

“Environment variable” is an application signature property that can be added using the following syntax:

```
NAME=VALUE
```

where NAME is the name of the environment variable, and VALUE is its value. You can type either name, the value, or both to find matching process families. However, if you want to find an exact match, you must use both NAME=VALUE.

### **Signature Evaluation Order**

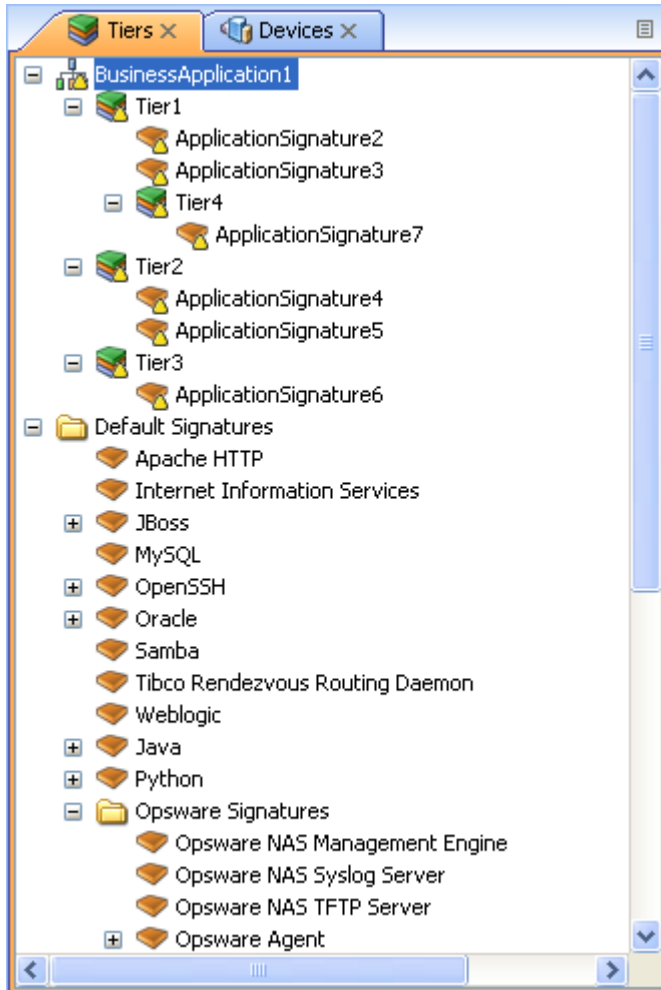
The order that signatures are recognized in VAM is important because a process family or storage mapping is associated with the first signature that it matches in the Tiers tree. Evaluation order is significant especially when the recognition criteria for a signature matches the same process family or storage mapping found in multiple signatures.

Signatures are evaluated in a depth-first, top to bottom order: signatures in a tier's sub-tiers are evaluated before the tier's signatures (depth -first), and tiers in the Tiers tree are evaluated from top to bottom. Signatures are applied in the order in which they appear in each tier.

After all user-created tiers and signature hierarchies are evaluated, then all of the default Opsware signatures are evaluated; for example, Opsware NAS Management System, Opsware NAS Syslog Server, and so on. After the Opsware signatures are evaluated, then all of the default predefined signatures are evaluated; for example, Apache HTTP, Internet Information Server (IIS), and so on.

Consider an application definition that has the structure shown in Figure 1-17. In this image, no processes or process families match the signature definitions, and so the signatures are represented with the ⚠ icon.

Figure 1-17: Tiers Signature Evaluation Order



In this application definition example, the signatures are evaluated in the following order:

1. ApplicationSignature7
2. ApplicationSignature2
3. ApplicationSignature3
4. ApplicationSignature4

5. ApplicationSignature5
6. ApplicationSignature6
7. Opsware NAS Management System
8. Opsware NAS Syslog Server
9. <All remaining Opsware signatures, top to bottom>
10. Apache HTTP
11. Internet Information Services
12. <All remaining default signatures, top to bottom>

### Creating an Application or Storage Signature

To create an application or storage signature in the Tiers tree, perform the following steps:


- 1** Make sure you have already created a Tier into which you want to create an application or storage signature.
- 1** Select the Tiers tree.
- 2** From the **Application** menu, select **New Application Signature** or **New Storage Signature**. The Properties pane becomes active and ready to be filled out for the new signature.
- 3** In the Properties pane for an *application* signature, enter any of the following information:
  - **Process Name:** The name of the process family.
  - **Command Line:** The command line that a signature was started with.
  - **Executable Path:** The path to the executable file of this application component.
  - **Open Files:** The name of an open file.
  - **Modules:** The shared libraries associated with the process family. These include DLLs on Windows and shared object files on Unix.
  - **Environ Variables:** For environment variables, enter a name, the value, or both to find matching process families, where NAME is the name of the environment variable, and VALUE is its value. If you want to find an exact match, you must use both NAME=VALUE.
  - **Connected to Port:** The port the signature is connected to.

- **Listener Port:** The port on which the signatures are listening.
- **Alias:** The name of the application component as displayed in the different views.
- **Background color:** Click to change the background color displayed in the different maps.
- **Foreground color:** Click to change the foreground text color displayed in the different maps.

**4** If you created a *storage* signature, enter any of the following information:

- **Name:** The name of the storage device; for example: HiCommand.
- **Alias:** Alias for the storage device (if any).
- **Remote Volumes:** Number of remote volumes on the storage device.
- **LUN ID:** LUN ID number.
- **LUN Name:** Name of the LUN.
- **Exported Path:** Remote exported path for NAS filers.
- **Manufacturer:** Maker of the storage device.
- **Model:** Model number.
- **Background color:** Click to change the background color displayed in the different maps.
- **Foreground color:** Click to change the foreground text color displayed in the different views.

**5** After each entry, press Return on your keyboard to enter the property.

**6** When you have finished defining the signature, click **Refresh Snapshot**  on the VAM toolbar so VAM can update the snapshot and scan your data center to find matching process families and storage devices.

## Editing Signatures

To edit a signature in the Tiers tree, perform the following steps:

- 1** In the Tiers tree, select a signature.
- 2** From the Properties pane of the selected signature, double-click in the right-side of the property entry and edit the text.

- 3 Press Return on your keyboard to enter the changes.

### Deleting Signatures

To delete signatures from the Tiers tree, perform the following steps:

- 1 In the Tiers tree, select an application component.
- 2 From the **Edit** menu, select **Delete** or right-click and then select **Delete**.



### Cutting and Copying Signatures

You can cut and copy signatures to the clipboard. After you do this, you can paste the signature before or after a selected position in the Tiers tree to rearrange the order.



Default signatures and Opware signatures can be copied and pasted into user-created application tiers, but cannot be deleted or overwritten.

To cut and copy a signature tier in the Tiers tree, perform the following steps:


- 1 In the Tiers tree, select a signature.
- 2 From the toolbar, select the  icon or the  icon, or right-click and select **Cut** or **Copy**.
- 3 Then, click **Paste** from the **Edit** menu, or press Control + V.

Or

- 4 You can press the Control button on your keyboard, and then select and drag a signature from one tier to another, creating a copy of the selected signature.

### Pasting a Signature

You can perform the following paste actions if one or more signatures have been cut or copied to the clipboard:

- Select a signature in the Devices tree and then select the Paste icon . The signatures that you cut or copied to the clipboard will be appended to the selected tier's signatures.

- Select a signature in the Devices tree and then select **Paste** from the **Edit** menu. The signatures that you cut or copied to the clipboard will be inserted into the selected signature's parent tier *before* the selected signature.
- Select a signature in the Devices tree and then select the **Paste** from the **Edit** menu. The signatures that you cut or copied to the clipboard will be inserted into the selected signature's parent tier *after* (below) the selected signature.



---

For default Opware signatures, you can copy and paste them into user-created application tiers, but you cannot delete or overwrite them.

---

## VAM Business Application Management

In VAM, a snapshot represents the state of a set of network and storage devices and managed servers, the process families running on those servers, the connections among those process families, local and remote storage devices, file systems, and any external clients and dependencies. Snapshots are saved as part of a Business Application. A Business Application can contain any number of Snapshots, each of which can be used for Snapshot comparisons.


You can open, edit, and save a Business Application from inside of VAM, from the SAS Client's Library, or from the Opware Global File System (OGFS).

However, in order to save Business Applications to the Library or OGFS, your user account needs permissions to be able to write to those directories. To obtain the necessary permissions, contact your Opware administrator.

### Opening a Business Application

After you have launched the VAM, you can open a previously saved Business Application.

To open a Business Application, perform the following steps:

- 1** In the Opware VAM window, select the  toolbar icon or select the **File** menu and then select **Open** to display the Open window.
- 2** In the Look in drop-down list, select the directory on your computer, in the Opware Library, or the OGFS where the Business Application was saved.
- 3** Click **Open**.



## Saving a Business Application

To save a Business Application, perform the following steps:

- 1** From the **File** menu, select **Save** or **Save As** to open the Save window. (Note that by default, all scan results are selected to be saved.)
- 2** If you chose Save As, in the Save in drop-down list, select the your local computer, the Opsware Library, or the OGFS and choose where you want to save the Business Application.
- 3** Click **Save**.



If you exit VAM before saving your changes (either application definition changes or topology changes), you will be prompted to choose whether you want to save your changes and then exit or exit without saving your changes.



## Saving a Business Application as an Application Template

If you would like to save the current application definition as a template, so it can be reused or set to open VAM using that definition, see “Business Application Templates” on page 103.

## Running Scripts on Devices

From inside VAM you can run a script, either directly on a selected server or network device (but not on SAN devices), or on the Opsware Global File System (OGFS) using the Global Shell – given that your user account has sufficient permissions to run the Global Shell and to perform any operations on servers under Opsware management.


There are three possible scenarios in which you can run a script in VAM:

- By selecting a server and clicking **Run Script**  on the VAM toolbar, or selecting **Run Script** from the **Manage** menu. This launches the Run Script Task window.
- By selecting Run Global Shell Script , which launches the Opsware Global Shell, and which gives you access to the OGFS.


- Selecting a network device and click **Run Script** on the VAM toolbar, or selecting **Run Script** from the **Manage** menu, This launches the NAS interface, where you can log in to NAS and run the script on the selected network device.

For more detailed information about the script execution process and how it works, see Chapter 8, “Script Execution” on page 471 of this guide. For information on running scripts on network devices, consult the NAS online documentation.


To run a script on a server, perform the following steps:

- 1** From inside VAM, select a server from the Devices pane or one of the maps.
- 2** From the **Manage** menu, select **Run Script**, or click Run Script  from the VAM toolbar.
- 3** In the Run Script window, fill out the necessary information and perform the steps to execute the script. For more information on running a script on a server, see “Running a Server Script (Saved Script or Ad-Hoc Script)” on page 484.

To run a global shell script in VAM, perform the following steps:

- 1** From inside VAM, from the **Manage** menu, select **Run Global Shell Script**, or click Run Global Shell Script  from the VAM toolbar.
- 2** In the Run Global Shell Script window, fill out the necessary information and perform the steps to execute or schedule the script execution. For more information on running a global shell script on the OGFS, see “Running an OGFS Script” on page 490.

To run a script on a network device, perform the following steps:

- 1** From inside VAM, select a network device from the Devices pane or the Network Map.
- 2** From the **Manage** menu, select **Run Script**, or click Run Script  from the VAM toolbar. This launches the NAS web application interface.
- 3** Log in to NAS, and on the New Task - Run Command Script page, fill out the necessary information to run or schedule the script execution. For more information on running a script on a network device through NAS, consult the NAS online documentation by clicking the Help link in the upper right corner of the page.



You must have proper permissions to run global shell scripts and scripts on a device. To obtain these permissions, contact your Opsware administrator. See the *Opware® SAS Administration Guide* for more information.

## ACLs and Server Pool Configurations

For network devices – such as firewalls, load balancers, routers, switches – you can view Access Control List (ACL) configuration information. For load balancers, you can view server pool configuration (in addition to ACLs).

You also can compare ACL and server pool configurations between two devices in the same snapshot, or when in comparison mode, between the *same* device in the two *different* snapshots.

This section shows you how to perform the following tasks:

- Viewing ACLs
- Comparing ACLs
- Viewing Server Pool Configuration
- Comparing Server Pool Configuration

### Viewing ACLs

In VAM, you can view ACLs for network devices such as firewalls, load balancers, routers, switches.

To view ACL configuration information, perform the following steps:

- 1** In the Network Map or Devices Tree, select a network device that has ACLs configured for it, right-click, and select View ACL Configuration. The Access Control Lists window opens.

You can also access the Access Control Lists window by:

- From the Properties pane of the selected device, click the View link in the ACLs property.
- Select View ACL Configuration from the **Manage** menu.

- 2** In the Access Control List window, you can select and copy the text information. Also, from the drop-down list at the top of the window, you can select other devices in the snapshot that contain ACLs and view them. You can also search for specific text strings in the current ACL configuration, highlight strings, and perform other search functions.
- 3** When you are finished viewing ACLs, click **Close**.

### Comparing ACLs

You can compare ACL configurations between any two devices in the same snapshot.

(To compare ACLs between the same device in two different snapshots, see “Comparing Snapshots” on page 118.)

To compare ACL configurations for two devices in the same snapshot, perform the following steps:

- 1** From the Devices Tree, expand the Network Devices node.
- 2** Press and hold the Control button on your keyboard, and then select two network devices that have ACLs configured, such as a load balancer, a firewall, router, or LAN switch. (To see if a network device has ACLs configured, select the device and look for ACLs in the Properties pane.)
- 3** From the **Manage** menu (or, right-click), select **Compare ACL Configuration**. In the Compare window, you see two panes side by side, each representing one of the devices you are comparing. At the bottom of each pane is the name of the device.

To indicate the differences for each ACL configuration, the Comparison window uses the following colors:

- **Blue**: Indicates that information has been added modified.
- **Red**: Indicates that information is different between both configurations.
- **Black**: Indicates no changes.

To move through differences between the two configurations, click the arrow buttons at the right top of the window.

- 4** When you are finished viewing the differences, click **Close**.

## Viewing Server Pool Configuration

To view server pool configuration for a load balancer, perform the following steps:

- 1 In the Network Map or Devices Tree, select a load balancer that has server pool configurations, right-click, and select **View Server Pool Configuration**. The Server Pools window opens.

You can also access the Server Pools window by:

- From the Properties pane of the selected device, click the View Server Pools in the load balancer properties.
  - Select View Server Pool Configuration from the **Manage** menu.
- 2 In the Server Pools window, you can select and copy the text information. Also, from the drop-down list at the top of the window, you can select other devices in the snapshot that contain server pool configurations and view them. You can also search for specific text strings in the current server pool configurations, highlight strings, and perform other search functions.
  - 3 When you are finished viewing server pool configuration, click **Close**.

## Comparing Server Pool Configuration

You can compare server pool configurations between any two load balancers in the same snapshot.

(To compare server pool configurations between the same load balancer in two different snapshots, see “Comparing Snapshots” on page 118 for more information.)

To compare server pool configurations for two load balancers in the same snapshot, perform the following steps:

- 1 From the Devices Tree, expand the Network Devices node.
- 2 Press and hold the Control button on your keyboard, and then select two load balancers that have a server pool configuration. (To see if a load balancer has a server pool configured, select the device and look for Server Pools in the Properties pane.)
- 3 From the **Manage** menu (or, right-click), select **Compare Server Pool Configuration**. In the Compare window, you see two panes side by side, each representing one of the devices you are comparing. At the bottom of each pane is the name of the device.


To indicate the differences for each server pool configuration, the Comparison window 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.

To move through differences between the two configurations, click the arrow buttons at the right top of the window.

- 4** When you are finished viewing the differences, click **Close**.

## Comparing Snapshots

When you click **Refresh Snapshot**  on the VAM toolbar and then click **Save**, VAM captures and saves all information related to your application in the currently open Business Application in a VAM *snapshot*. A snapshot including all servers and processes associated with the business application, the current state of all running processes, all local and remote storage devices, and all values and signature definitions you have created in the Tiers tree.

Each time you refresh a snapshot and then save the Business Application, snapshot results are saved within the currently loaded Business Application. You can also schedule snapshots to occur at later time on a one time or recurring schedule.


Each saved snapshot can be used in a one to one comparison between the currently loaded snapshot and a saved one. You can take snapshots from the currently loaded Business Application or from another Business Application to help you determine if any changes have occurred between the current state of the business application and its state as captured in a previously saved snapshot.

When you compare scan results, VAM evaluates certain key objects and their attributes on a one to one basis, and displays any differences in value between those objects. The results of the comparison are displayed in the Compare Results pane at the bottom of the VAM window.

## Creating a Snapshot

Create a snapshot in VAM any time you want to capture the current state of your business application. Because a data center and all the devices and elements within it are constantly changing, it is a good idea to capture the current state so you can compare the current state of a business application with one you captured in the past.

To create a snapshot, perform the following steps:

- 1** Click **Refresh Snapshot**  on the VAM toolbar.
- 2** From the **File** menu, select **Save**. (Or, click **Save** on the VAM toolbar) A new snapshot has been created.
- 3** To see the snapshot and give it a name, from the **Application** menu, select **Show Snapshots**.
- 4** In the Snapshots window displays all saved snapshots. To rename a snapshot, click the name cell of the list and type a name.

## Opening a Snapshot

To view a previous state of a business application, you can load and view a saved snapshot.

To open a saved snapshot, perform the following steps:

- 1** From the **Application** menu, select **Show Snapshots**.
- 2** In the Snapshots window, select a saved snapshot and click **Open**.
- 3** The business application snapshot opens inside of the VAM application window. To delete a snapshot, select it and click **Delete**.

## Scheduling a Snapshot

You can automate snapshot creation by scheduling a snapshot at a future point in time, or you can schedule a recurring snapshot to regularly capture the state of your business application.

To schedule a snapshot, perform the following steps:

- 1** From the **Application** menu, select **Scheduled Snapshots**.
- 2** In the Scheduled Snapshots window, click **New Schedule**.
- 3** Type a name for the snapshot schedule in the Name field.

**4** In the Scheduled Frequency section, select one of the following snapshot frequency options:

- **Daily:** Choose to run the snapshot on a daily basis.
- **Weekly:** Choose a day of the week to run the snapshot.
- **Monthly:** Choose the months to run the snapshot specification job.
- **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

```
0 0 * * 1-5
```


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 more information about crontab entry formats, consult the Unix man pages.

**5** In the Time and Duration section, select a start and end time and day of the month.

**6** When you are finished filling out the schedule, click **Close**.

### “Source” and “Comparison” Snapshot

The currently loaded snapshot in VAM is referred to as the *source* snapshot, while the set of scan results you are comparing against the currently loaded snapshot is called the *comparison* snapshot. When you compare snapshots, you are always comparing the currently loaded scan result (*source*) with another saved snapshot result (*comparison*).

(Remember that to create a new snapshot, you need to click **Refresh Snapshot**  on the VAM toolbar and then click **Save**.)

### Comparison Types

VAM will display comparison results based on the following criteria:

- Object Existence Comparison
- Object Attribute Difference
- “Significant” Object Attribute Differences



### Object Existence Comparison

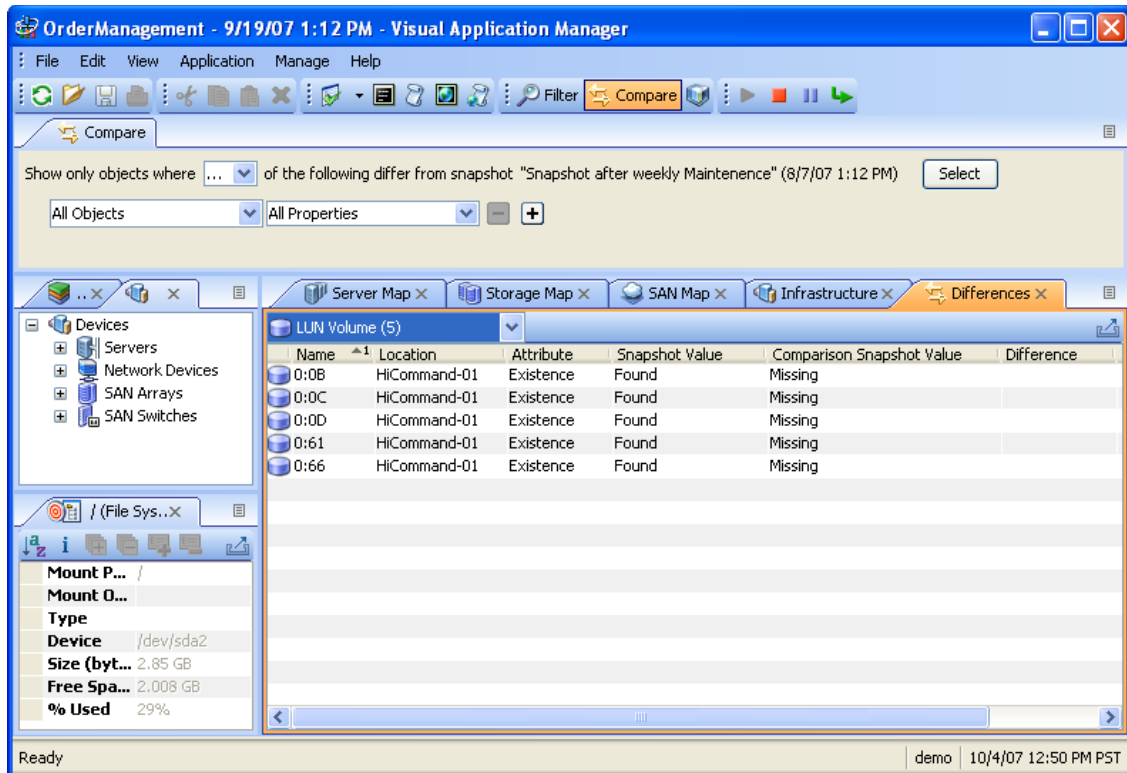
Comparing two snapshots helps determine whether or not an object exists between the them. If an object exists in one snapshot, but does not exist in the other, the comparison results display the object with an attribute named “existence” and describes it as either “found” or “missing” on either the source or comparison snapshot. Figure 7-8 describes the types of objects compared during a VAM snapshot comparison.

Table 1-2: Objects Checked for Existence in a VAM Scan Results Comparison

OBJECT	UNIQUE IDENTIFIER
Compliance Policy	Opware policy ID
Disk	Opware server ID plus device Opware ASAS array ID plus serial number Opware ASAS filer ID plus shelf plus bay
Fibre Channel Adapter	Opware server/ASAS device ID plus device
Fibre Channel Port	Opware server/ASAS device ID + WWN
File Systems on servers	Opware server ID plus device
LUN Volume	Opware ASAS device ID plus Opware ASAS volume ID
NAS Filer	Opware ASAS device ID
NFS Export	opware ASAS device ID plus export
Network Device	Opware NAS device ID
Network Interface	Opware server ID plus device
Network Device Port	Opware NAS device ID plus port name
Processes on servers	Opware server ID plus process id plus process start time (~one hour)
Process families on servers	Must contain at least one process in common
SAN Array	Opware ASAS device ID
SAN Switch	Opware ASAS device ID
Servers	Opware server ID

A process family could be running (and thus, “exist”) when you refresh a snapshot and save the Business Application. But if the process is no longer running, and you refresh the snapshot again and save the Business Application, the results change. When you compare the saved snapshot (the “comparison” snapshot) with the currently loaded snapshot (source), the results of the comparison will display in the Differences pane, as shown in Figure 1-18.

Figure 1-18: Snapshot Comparison Showing Difference in Existence of LUN Volumes



The results of the selected row show that on the target snapshot, all the listed LUN volumes are missing, meaning they did not exist in the comparison snapshot, but they exist now in the current snapshot.

### Object Attribute Difference

The VAM compare feature also evaluates two snapshots to determine any differences in the value of an object attribute. If the same attribute does not match between the two snapshots, then it will be marked as a “difference” and displayed in the comparison results. For attribute values with numerical differences, the comparison results will display both the numerical difference and percentage of change.

For example, if you scan a server in VAM and the server shows that it has 2 gigabytes of RAM, and then at a later point in time, one gigabyte of RAM is removed from the server, then when you perform a comparison, the results will show the server's total memory as having a difference of one gigabyte. The results will also indicate that the target (the earlier saved scan results) had a value of two gigabytes, and the source (currently loaded scan results) has a value of one gigabyte.

Table 1-3 lists all object attributes evaluated during a scan results comparison.

Table 1-3: Object Attributes Checked for Difference in Snapshot Comparison

OBJECT CATEGORY	OBJECT ATTRIBUTES COMPARED
Compliance Policy	Existence, compliance status (compliant, partial, noncompliant, scan failure, scan needed, scanning)
Disk (applies to servers, SAN arrays, and NAS filers)	Existence, firmware version, status
Fibre Channel Adapter	Driver version
Fibre Channel Port (applies to servers, SAN arrays, and NAS filers)	Connected port, existence, fabric, name, zone
File System	Existence, mount options, mount point, significant change in % free space, size (bytes), type
LUN Volume	Capacity (bytes), existence, LUN ID, name
NAS Filer	Existence, hostname, name, operating system version
NFS Export	Existence, export path, significant change in % of free space, size (bytes)
Network Device	ACLs, existence, firmware version, name, operating system, server pools
Network Interface	Broadcast address, connected switch, connected switch port, connected VLAN, duplex, existence, IP address, MAC address, neg. duplex, neg. speed, subnet mask
Network Port	Duplex, existence, MAC address, neg. duplex, neg. speed, speed, VLAN
Process Family	Max CPU utilization, significant change in # of connections, significant change in # of open files

Table 1-3: Object Attributes Checked for Difference in Snapshot Comparison

OBJECT CATEGORY	OBJECT ATTRIBUTES COMPARED
SAN Array	Existence, firmware version, name
SAN Switch	Existence, firmware version, name
Server	Boot time, codeset, DNS server, existence, kernel, name, operating system
VLAN	Existence

**“Significant” Object Attribute Differences**

VAM will also compare a set of attributes by using special heuristics specific to certain attributes, so that differences VAM considers “significant” will be shown.

If an attribute value in one of the snapshots exceeds a minimum (or maximum) threshold and its value changes by at least a certain percentage between the snapshots being compared, then VAM will present this in the comparison results.

Table 1-4 shows special object attribute differences.


Table 1-4: “Significant” Object Attribute Differences in VAM

OBJECT	OBJECT ATTRIBUTES COMPARED
Server	Load average, percentage of free memory on a server, percentage of free swap memory
Server’s file system	Percentage of free space
Process families	Number of open files, total number of all related connections, total count of process family member connections
NFS Export on NAS Filer	Percentage of free space

For more information on the heuristics used to calculate what is considered a significant difference, see “Scan Results Comparison Heuristics” on page 134.

## Comparing Snapshots

In order to compare two snapshots, you must have at least one saved snapshot. If you select **Save** from the **File** menu, this will save the currently loaded snapshot. If you click

**Refresh Snapshot**  on the VAM toolbar, and then save again, this will create and save a new snapshot.

To compare snapshots in VAM, perform the following steps:




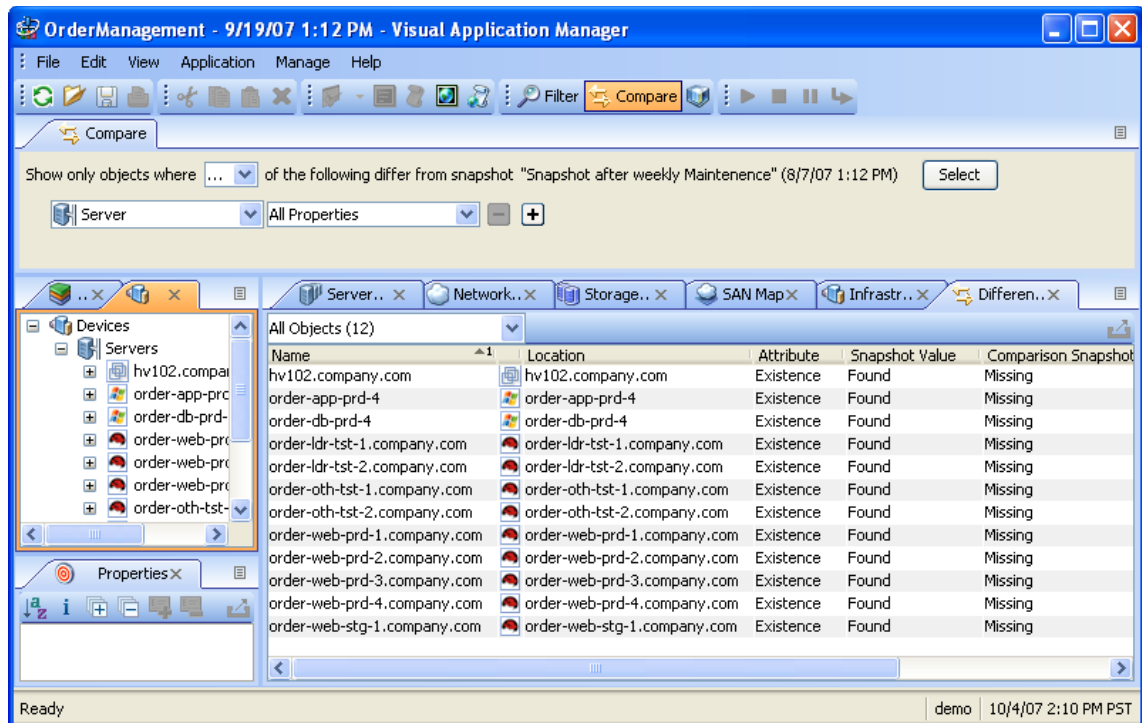
- 1** From the **View** menu, select **Compare**. (Or, click Compare  on the VAM toolbar). The Compare pane appears.
- 2** Click **Select**. The Select Comparison Snapshot window opens.
- 3** Select a snapshot you want to compare against the currently loaded snapshot.
- 4** Next, select if you want to show either ANY or ALL of the selected objects (plus their attributes) that are different to show when you compare snapshots.
- 5** From the drop-down list, select an object category to compare and create a comparison rule. You can select any of the object categories, such as file systems, or select All Categories.
- 6** To select another comparison rule, click Add  to add another criteria selector. To remove a comparison rule, click Remove .
- 7** As you create comparison rules, VAM automatically displays any or all differences in the Differences pane.

Figure 1-19 displays differences between two snapshots where several servers were found to exist in the currently loaded snapshot, but where the comparison snapshot showed none of the servers to exist.

Figure 1-19: Compare Results Showing Server Differences



Notice that the column named Snapshot Value lists the servers as found, which means that all of these server are listed as “Found,” which means all the listed servers exist in the currently loaded snapshot. The column next to that named Comparison Snapshot Value lists all the servers as “Missing,” which means that the none of the listed servers exist in the comparison snapshot.

- 8 To view differences in other maps, you can select a map, and all objects that are found to exist will appear normally. All objects that are listed as missing will appear grayed out. You can also select other criteria in the Difference pane drop-down list to filter the results of the comparison in more granular detail.
- 9 To close the Compare pane, from the **View** menu, select **Compare** again. (Or, click

**Compare**  on the VAM toolbar.

## Filtering VAM Data


When you select Filter from the **View** menu, or, click **Filter**  on the VAM toolbar, a search control appears above the maps and tabs that allows you to filter according to the following objects (if they were captured as part of the scan) and their attributes, listed in Table 1-5.

Table 1-5: Objects and their attributes you can filter in VAM

OBJECT	ATTRIBUTES FILTERED
Compliance Policy	Name, compliance status (compliant, partial, noncompliant, scan failure, scan needed, scanning)
Disk (applies to servers, SAN arrays, and NAS filers)	Device, manufacturer, model, serial number, size (bytes), type
Fibre Channel Adapter	Driver version, firmware version, hardware version, model, node world wide name, serial number
Fibre Channel Port (applies to servers, SAN arrays, and NAS filers)	Fabric, name, port number, status, world wide name
File System	% used, device, free space (bytes), mount options, mount point, size (bytes), type
LUN Volume	Capacity (bytes), name
NAS Filer	Hardware version, hostname, manufacturer, model, name, operating system version, Opsware ID, serial number
NFS Export	% used, export path, free space (bytes), size (bytes)
Network Device	Asset tag, firmware version, manufacturer, model, name, operating system, Opsware ID, processor
Network Interface	Broadcast address, device, duplex, IP address, MAC address, neg. duplex, neg. speed, speed
Network Port	Duplex, MAC address, neg. duplex, neg. speed, speed
Process Family	CPU utilization, command line, connected port, environment variables, listener port, modules, name, open files

Table 1-5: Objects and their attributes you can filter in VAM

OBJECT	ATTRIBUTES FILTERED
SAN Array	Firmware version, manufacturer, model, name, Opware ID, serial number
SAN Switch	Firmware version, hardware version, manufacturer, model, name, Opware ID, serial number
SAN Zone	Name
Server	1 minute load average, 15 minute load average, 5 minute load average, codeset, free memory (bytes), kernel, name, operating system
VLAN	Descriptions, ports, VLAN ID


You can filter the current snapshot in according to one or several of the objects in the list, as well as applying operators and attributes, according to certain attributes, and then apply operators to the attributes, depending if the object is a string or a number.

Filter results appear in the Tabs and Maps. In the maps, while all other objects that do not meet the filtering criteria appear grayed out in the maps.



For more information on filtering criteria and regular expressions, see “Filter Criteria” on page 129.

### Creating a Data Filter in VAM

To filter data that was collected in the currently loaded VAM snapshot, perform the following steps:

- 1** From the **View** menu, select **Filter**. (Or, from the VAM toolbar, click **Filter** ,.)  
The Filter pane appears above the maps.
- 2** In the Filter pane rule criteria, choose if you want to show either ANY or ALL of the selected objects (plus their attributes) that meet your filtering criteria.
- 3** From the drop-down list, select an object category to filter in the current snapshot and add criteria to narrow the filter, such as Compliance Policy, File System, Process Family, SAN Array, and so on.



- 4 Using the criteria drop down list, create a meaningful expression. For example, if you chose Disk as a category, you could set Size (bytes) is greater than (>) 5000 (bytes). For more information on expressions, see “Filter Criteria” on page 129 and “Examples of Regular Expressions” on page 130.
- 5 To select another filter criteria rule, click Add . To remove a comparison rule, click Remove .
- 6 As you create filter rules, VAM automatically displays any or all results in any of the maps or the Infrastructure pane. All results that meet the criteria appear normally in the maps and the Infrastructure pane. Any results in the snapshot that do not meet the criteria are show in the maps grayed out.
- 7 From any of the results related to server, you can select the server on which the results were found, right-click, and select **Open Remote Terminal** or **Open Device Explorer** to browse the server.

## Filter Criteria

In the filter criteria text boxes in the Filter pane, enter Perl 5 compatible regular expressions as filtering criteria. You can filter by using standard text matching and also by adding any regular expression patterns.

### Strings Operators

- Contains (default)
- Does Not Contain
- Is
- Is Not
- Starts With
- Ends With
- Matches Regular Expression

### Numbers

- == is equal to
- != not
- < less than

- > Greater than
- <= less than or equal to
- >= greater than or equal to

By default, all filtering is performed in case-sensitive mode. To filter in a case-insensitive mode, select the corresponding check box for the filter criteria.

In paired text boxes ([ ] to [ ]), enter numbers only to filter within a numeric range. If either the beginning or ending number in the range is left blank, then that part of the range will not be included in the filter. For example, to find all file systems that are almost full, you would filter for file systems that are using at least 80% of their capacity. Enter 80 to specify the beginning of the range and leave the end of the range blank, such as [80] to [ ]. To filter for an exact numerical match, enter the same number in the beginning and ending range positions, such as [80] to [80].

The units of measure for filtered items should match what is shown in the Filter Results and Properties Pane, such as:

- **Memory:** Bytes
- **Uptime:** Days
- **Percentages:** A number from 0 to 100, (such as disk space used and CPU utilization)
- **Disk space:** Bytes




### **Examples of Regular Expressions**

The following examples show how to use regular expressions in filter text boxes:

- **Operating System:** To find all servers that are not running a Windows operating system, look for servers whose operating system does not begin with an "M" (for Microsoft Windows). For example, enter `^[^M]` in this text box.
- **Kernel:** To find servers whose kernel is one of 2.6.5, 2.6.6 or 2.6.7, enter `2.6.[5-7]` in this text box.
- **Mount Point:** To find all mounted Unix file systems other than /, enter `/.+` in this text box.

## VAM Scan Error Messages

VAM indicates when an error occurred on a managed server by displaying the following server icons when you move your mouse pointer over the icon:

- **Server Error Icon** : There was an error in gathering information from the server when VAM scanned it (see Table 1-6 for possible causes for the error).
- **Server Unreachable Error Icon** : The Opsware core was not able to communicate with the Opsware Agent installed on the server.
- **Server Unknown** : VAM is unable to scan the server at all, possibly because the server is no longer in the core and under Opsware management.

It will show these icons before the server name in the Devices tree, Network Map, Virtualization Map, and Server Map. You can move your cursor over the server name to display the detailed error message.



Scan failures and scan time-outs typically occur when the Opsware managed server is very busy, or when network traffic is very heavy or running over a low bandwidth connection. If these types of errors occur too frequently, please contact your Opsware administrator for assistance.

### Server Scan Errors

Table 1-6 describes server scan errors and recommended actions.

Table 1-6: Server Error Messages in VAM

ERROR	DESCRIPTION	ACTION
Not Enough Disk Space	A selected managed server does not have enough disk space to perform a scan.	Free up disk space.

Table 1-6: Server Error Messages in VAM (continued)

ERROR	DESCRIPTION	ACTION
Remediation Failed	The Runtime State Server Module failed to remediate on the selected server.	Click the Remediation job number link and the job window from the SAS Client opens. Or, select the server, right-click, and select <b>Open Device Explorer</b> to troubleshoot the error.
Scan Timed Out	The scan process has exceeded the time-out limit.	See "Scan Time-Out Preference" on page 99.
Server Access Denied	By using the OGFS, you are unable to access the server's file system as root (on a Unix server) or as LocalSystem (on a Windows server).	Contact your Opware administrator for the required permissions.
Server Capture Failed	The remote capture of data or the transfer of data back to the Opware core failed.	Review the log file that is in /tmp/.sitemap/<number> for details in your global shell session.
Server ID Invalid	The server's directory was not found in the OGFS, which means that Opware SAS does not know the server exists.	
Server Scan Agent Failed	The driver used to collect data could not be correctly copied to the managed server. This could be caused by a checksum mismatch.	Contact Opware Support and provide the log file.
Server Unreachable	The managed server is unreachable by Opware SAS. This could be caused if the Opware core cannot communicate with the server's Opware Agent.	Try again later. If this condition persists, contact your Opware administrator.

Table 1-6: Server Error Messages in VAM (continued)

ERROR	DESCRIPTION	ACTION
Unknown Scan Error	An unknown error occurred during the scanning process.	Try again later. If this condition persists, contact your Opware administrator.
Unsupported Agent for Scan	The VAM does not support the Opware Agent version running on a selected managed server.	Opware Agent 5.1 or higher is required.
Unsupported OS for Scan	The VAM does not support the operating system running on a selected managed server.	See "Supported Operating Systems" on page 35.

### **Network Device Scan Errors**

Table 1-7 describes network device scan errors and recommended actions.

Table 1-7: Network Device Scan Error Messages in VAM

ERROR	DESCRIPTION	ACTION
NAS Scan Timed Out	The time needed to gather NAS data exceeded the timeout	Scan fewer devices or wait until the NAS server can handle this request.
NAS Scan Failed	Gathering NAS data failed.	Export this snapshot to a Business Application and contact your Opware administrator.

## Storage Scan Errors

Table 1-8 describes network device scan errors and recommended actions.

Table 1-8: Storage Scan Error Messages in VAM

ERROR	DESCRIPTION	ACTION
ASAS Scan Timed Out	The time needed to gather ASAS data exceeded the timeout	Scan fewer devices or wait until the SAS core server can handle this request.
NAS Scan Failed	Gathering ASAS data failed.	Export this snapshot to a Business Application and contact your Opware administrator.

## Scan Results Comparison Heuristics

When you compare scan results in VAM, specific objects and their attributes are evaluated between each scan result and any differences (and non-existence of objects) will be displayed in the comparison results. (For information on the basic set of objects and attributes compared in a scan results comparison, see “Comparing Snapshots” on page 118.)

In addition to the basic set of attributes evaluated in a comparison, VAM also applies a certain set of heuristics to some attributes in order to discover unique differences that VAM has determined to be interesting or useful.

Specifically, if an attribute value in one of the scan results exceeds a minimum (or maximum) threshold and its value changes by at least a certain percentage between the scan results, then VAM will present this in the comparison results.

The following special object attribute differences are compared:

- **Server:** Load average, percentage of free memory on a server, percentage of free swap memory.
- **Server's File System:** Percentage of free space.
- **Process Families:** Number of open files, total number of all related connections, total count of process family member connections.

The heuristics applied to certain attributes in scan results during a comparison are listed in Table 1-9.

The following variables are used in the expressions:

- X = The maximum value of the attribute between the two scan results.
- N = The minimum value of the attribute between the two scan results.
- P = The percentage change in value of the attribute between scan results.

Table 1-9: Scan Results Comparison Heuristics

OBJECT ATTRIBUTE	EQUATION
server – 15 minutes load average	$X > 0.8 * \text{cpu count AND } P > 20\%$ OR $X > \max(1, 0.25 * \text{cpu count}) \text{ AND } P > 100\%$
server – percentage memory free (%)	$N < 0.1 * \text{total mem AND } P > 25\%$
filesystem – percentage free (%)	$N < 0.2 * \text{size AND } P > 10\%$
process family – open file count (on any member process)	$X > 50 \text{ AND } P > 50\%$
process family – connection count (aggregate across all member processes)	$X > 50 * \text{process count AND } P > 30\%$
process family - connection count (on any member process)	$X > 50 \text{ AND } P > 50\%$





# Chapter 2: Audit and Remediation

## IN THIS CHAPTER

This section discusses the following topics:

- Overview of Audit and Remediation
- Audits
- Creating an Audit
- Viewing Server Audit and Snapshot Usage
- Configuring an Audit
- Audit and Remediation Rules
- Configuring Specific Audit and Snapshot Rules
- Audit Rule Exceptions
- Audit Policies
- Running an Audit
- Scheduling an Audit
- Viewing and Remediating Audit Results
- Snapshots
- Creating a Snapshot Specification
- Configuring a Snapshot Specification
- Scheduling Snapshot Jobs
- Locating Snapshots
- Viewing Snapshot Contents
- Copying Objects from a Snapshot to a Server

## Overview of Audit and Remediation

Opsware's Audit and Remediation feature allows you to define server configuration policies and ensure that servers in your facilities meet policy standards. When servers are found to be out of compliance (not configured the way you want them to be), you can remediate the differing server configurations.

With Audit and Remediation, you can audit server configuration values based on a live server (or server snapshot), or based on your own custom values (or both). Opsware Audit and Remediation also allows you to take snapshots of a server to capture the current state of a system, so you can perform server comparisons against a baseline, or use the snapshot inside of an audit. You can also create custom audit policies that define company or industry server configuration compliance standards, which can be used inside of audits or snapshot specifications.

If you subscribe to The Opsware Network (TON), you can be kept up to date on the latest industry compliance standards based on the needs of your data center. For example, subscribing to TON Essential Content gives you access to regularly updated security best practices, such as the Center for Internet Security (CIS), NSA, and so on, as well as the Opsware patch supplement for Microsoft Windows. The TON Subscription Service enables you to access the most current regulatory compliance policies (FISMA, Sarbanes-Oxley, etc.) and daily vulnerability alerts. You can also join the TON content developer communities to share and access custom-created audit policies and rules. And much more.



---

For information about subscribing to TON, contact your Opsware sales representative.

---

### Audit and Remediation Examples

The following examples illustrate ways the Opsware Audit and Remediation feature helps you manage server configurations in your facility:

- Capturing Golden Server Configurations
- Enforcing Security Policies

#### ***Capturing Golden Server Configurations***

Sometimes a server becomes configured in such a way that it represents the ideal state of server configuration for some purpose in your facility. For example, if you want to set up a collection of servers that handle web traffic, you might configure a single server that

represents a perfect configuration – a golden server configuration – for a group of Web servers. After you configure this golden server, you can duplicate the golden server configuration across a group of servers.

For example, you have a Red Hat Linux server with a unique configuration of Apache Web Servers, and you want to duplicate this exact configuration across several other servers. With Opsware Audit and Remediation, you can create an audit that uses the golden server as the source. In the audit, you select those configurations to use to audit other servers, such as an application policy and specific application configuration rules.

Then, select those servers as the target of the audit to be configured like the golden server. After you run the audit, you can remediate any target server's configurations that do not match the golden source. Then, you can schedule the audit to run on a regular basis, so if any of the servers become non-compliant, you can remediate them when they deviate from the golden standard.

### **Enforcing Security Policies**

Your IT organization likely has security policies you want to enforce, to make sure servers are configured properly and are safe from security attacks. Your organization can use Opsware Audit and Remediation to enforce these policies.

For example, your organization wants to ensure that a collection of Windows 2003 servers has a recent Microsoft security patch, regardless of the applications installed on the servers. By being subscribed to The Opsware Network (TON), you can access the rules for this patch in an audit policy to define this security configuration. This policy would let systems administrators who directly configure and manage those servers know that this policy exists. You can create an audit and link it to the audit policy that contains the patch, and then set the Windows 2003 servers as targets of the audit. The audit can be scheduled to run regularly. If the audit results show that any of the target servers do not contain the new security patch, those servers can be remediated to have the patch installed. If new patches come out and need to be installed on the target servers, you can update the audit policy with the new patch, and the audit that runs against the target servers is automatically updated to reflect the new patch definition.

### **Audits**

An audit is the tool you use to define the desired configuration values for a server, compare expected configurations against live servers, and remediate any differences found by the audit. Using audit rules, you can define the audit to look for such configurations as IIS Metabase, Windows Services, file system checks, hardware

configurations, application configurations, event logging, COM+, and so on. You can define what the audit should look for, what values you expect to find on the server, and what value to use to fix when differences are found.

For more information on audits, see “Audits” on page 143.

### **Audit Policies**

An audit policy is used to define rules for checking the configuration of a server and can be reused by other people in your organization. An audit policy contains a set of ideal server configuration rules that help define compliance best practices for others to use for running audits. Audit policies can be linked to audits or snapshot specifications, which maintain the latest changes made to the audit policy.

For more information on audit policies, see “Audit Policies” on page 208.

### **Audits and the Compliance View**

The Compliance View allows users to view the overall compliance levels for servers in their facility and helps them remediate compliance problems. For more information, see “Server Compliance” on page 257.

### **Snapshots**

Snapshots differ from audits in that snapshots allow you to take a picture of the current state of configuration of a server. Snapshots are useful for capturing the configuration of a golden or baseline server that you would like to compare against other servers in your facility. You can use the snapshot as the source of an audit if any servers do not match the configuration captured in the snapshot, then you can remediate those servers after the audit has run from the Audit Results window.

For more information on snapshots, see “Snapshots” on page 235.

## **Terms and Concepts**

The following list defines key Audit and Remediation terms and concepts:

- **Archived Audit Result/Snapshot:** Archiving audit results and snapshots allows you to move them from the audit result or snapshot list but keep them available for historical purposes.
- **Audit:** A set of rules that expresses the desired state of a managed server's configuration objects – for example, a server's file system directory structure or files,

a server's Windows Registry, application configuration, and so on. An audit also contains sources (servers, snapshots, snapshot specifications), targets (servers or snapshots), rule exceptions, and a schedule.

An audit's rules can be linked to an audit policy. An audit can be run to compare server configuration object values against a baseline server, a server snapshot, or user-defined values, to determine how values differ. When an audit reveals a difference between servers or user-entered values, the user can install software and server objects to remediate the variance.

- **Audit Job:** The process that occurs when you run an audit. An audit job can be run immediately one time, or on a recurring basis by scheduling the job. When an audit job is finished, it produces an Audit Result.
- **Audit rule types:** An audit can contain both types of the following rules:
  - **Comparison:** A rule that compares a server's or snapshot's configurations of a server with other servers or snapshots.
  - **Value-based (user-specified):** A rule that compares one or more set of user-defined values. This type of audit includes an audit that links to an audit policy.
  - **Existence:** A rule that checks for the non-existence of an object to determine if it exists on the target server. If the object exists on the target server, then the user or group rule is out of compliance.
- **Audit policy:** A collection of rules that defines a desired configuration for a server. A policy can be used by an audit in the following ways:
  - **Link:** A linked policy maintains a persistent connection between the audit and the policy. This means that the rules in the audit are exactly those of the audit policy, and if any updates are made to the policy, then the latest changes are also reflected in the audit to which the policy is linked.
  - **Import (replace, non-linked):** When a user imports a policy into an audit, then the connection between the audit and the audit policy is no longer maintained, and the user can make changes to the audit without affecting the policy. Conversely, any changes or updates made to the policy will not be reflected in the audit.
  - **Import (merge):** When an audit policy is imported and merged into an audit, the audit policy's rules are added to the rules already present in the audit. No persistent link between the audit and the audit policy is maintained. During the merge, if rules are found to conflict, the newly imported rules from the audit policy

will replace the rules in the audit policy.

- **Audit Result:** The results of running an audit. This shows how a target server or a group of servers' configuration object values match or mismatch the values as defined in the audit.
- **Exception:** A server and specific rules that has been excepted, or disabled, so that when the audit is run, the rule exception is not checked on the selected server thus not considered when determining audit compliance.
- **Compliance:** Denotes the degree to which a server object conforms to a test. Compliance in Opware Audit and Remediation is defined by the audit's or snapshot's rules, which specify the values expected of the target servers. If the values are different than specified, then the server is out of compliance.
- **Policy Setter:** A person in an organization who is responsible for defining server configuration compliance standards – the way a server should be configured – and who defines audit policies.
- **Rule:** A check on a particular server configuration object along with a desired value, and optional remediation value. Rules come in two types: server-based, which derive directly from a source server, and user-defined, which are created by a user.

If you are subscribed to The Opware Network (TON), you can access pre-created rules that define a wide range of industry compliance standards, such as the latest patch supplement for Microsoft Windows, current regulatory compliance policies (for example, FISMA, Sarbanes-Oxley), user-created rules from the TON developer community, daily vulnerability content updates, and so on.

- **Server Object:** An object from a server to which an audit or snapshot specification rule can be applied. This can be a value (such as minimum password length) or an object, such as a file or directory, registry entry, Windows Services hardware configuration, and so on. For more information on servers objects used in audits and snapshot specifications, see "Server Objects Used in Audits and Snapshots" on page 160.
- **Snapshot:** Shows a picture of how an Opware-managed server is configured at a certain point in time. A snapshot is the result of a snapshot specification job that has been run.
- **Snapshot Specification Job:** The process that occurs when you run a snapshot specification. A snapshot job can be run once, or on a recurring basis by scheduling the job. When a snapshot specification job is completed, it produces a snapshot.

- **Snapshot Specification:** An object window that allows you to define and create a snapshot. In other words, you can define the rules and servers to take a snapshot of.
- **Target:** The server or servers that you run an audit against or take a snapshot of. The target for an audit can be a server, several servers, a group of servers, or a snapshot. The target for a snapshot can also be other servers.

## Audits

An audit consists of a collection of rules that enable you to define what should be or what should not be for a server's configuration. An audit contains rules, a source, target servers, and a schedule that defines when and how often the audit will run.

Audit rules allow you to define and check the state of various objects on a server, such as the state of server's file system, registry settings, installed and registered software (patches and packages), events, software, application configurations, operating system settings, and so on. If the configuration of the object on the target server is different than the state you defined in the audit rules, the rule is considered Non-Compliant. When you view an audit's results, you can remediate the object configuration to make sure the target server's configuration is in compliance with the desired configuration.

You can audit server configuration values for a single server, groups of servers, or another server snapshot. You can also schedule audits to run immediately, or on a recurring schedule, and send email notifications when the audit has finished.

## Audit Comparison Types

In general, an audit can contain the two following types of comparisons, based on the source of the audit:

- **Comparison:** An audit based on configuration values from a source server or source snapshot specified at the time the audit is created. The source server or server snapshot is also known as a "golden" or reference server. For example, you might want to compare file directories or file contents, registry structures, IIS Metabase entries, or user group settings among servers. Using a snapshot as the source of an audit, you can compare the snapshot with other servers in your facility.

Comparison audits can perform the following types of comparisons:

- **Property:** Checks the property of a selected object or object configuration. For example, you could check the release version of a patch on a target server or group of servers, to make sure it matches what you expect to be installed on the

targets. You can select this version number based upon a source server or snapshot, or add your own value.

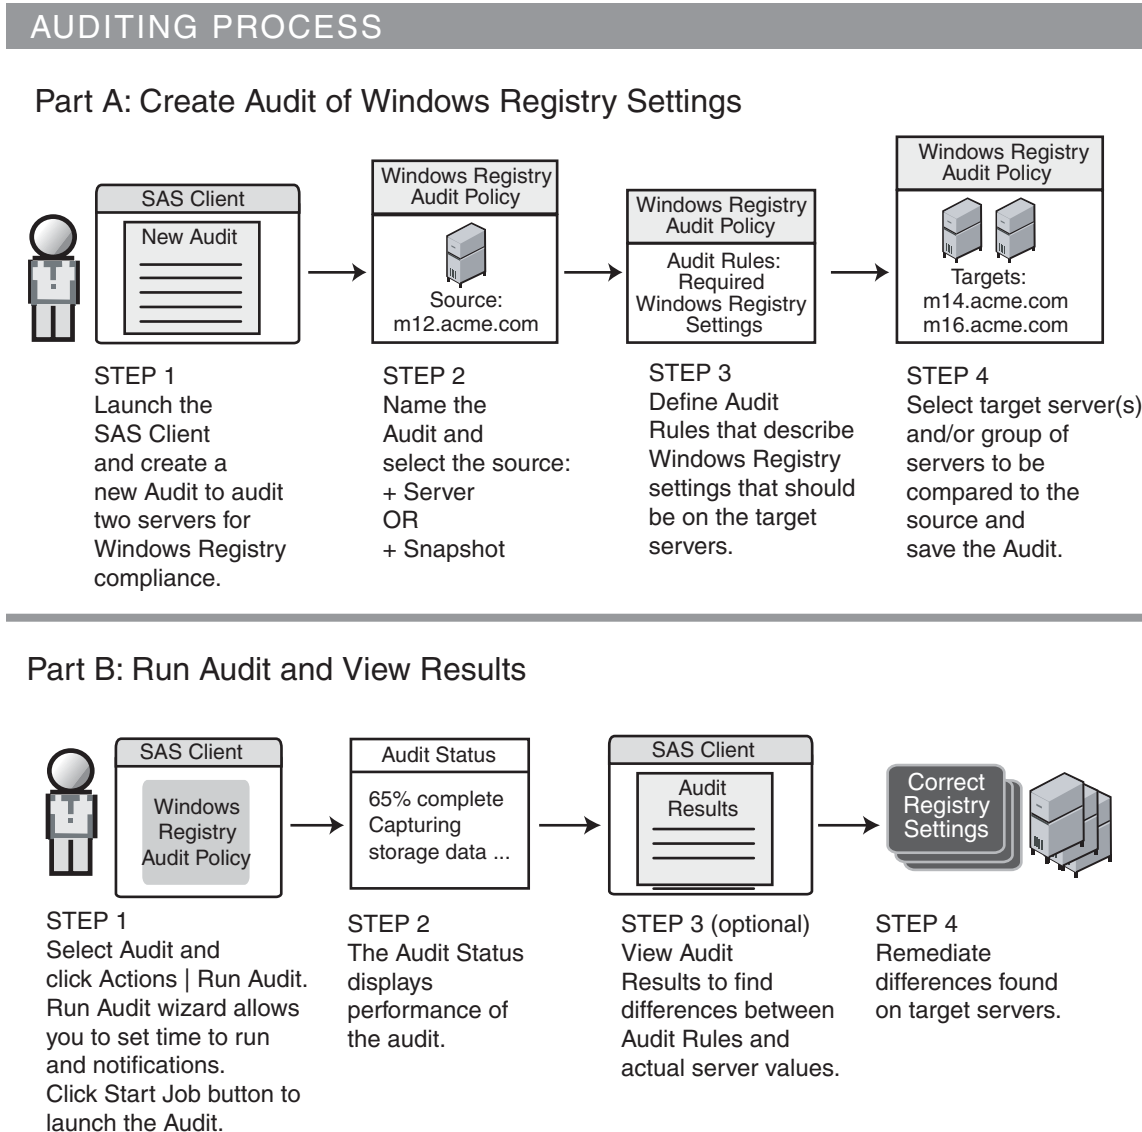
- **Equivalence:** Checks to determine that a target server configuration is the same between the source server or snapshot of the audit. For example, you could check to see if the target of the audit has the same user group as a group you selected from a source server.
- **Non-existence:** Checks the target server to determine the non-existence of a server object or configuration. For example, you could check a server to make sure it does not contain a specific COM+ object.
- **User-Defined Value Comparison:** An audit based on custom, user-defined values for each server object (file system, windows services, IIS Metabase, users and groups, and so on). These values can be derived from a source server, or from Opware attributes or custom attributes. This type of audit includes those based on an audit policy. In an audit policy, a user (known as a “policy setter”) pre-defines values for each configuration object based on company or industry compliance standards.



## The Auditing Process

The following diagram illustrates a basic example of creating and running an audit.

Figure 2-1: The Auditing Process



## Audit Elements

An audit consists of the following elements:

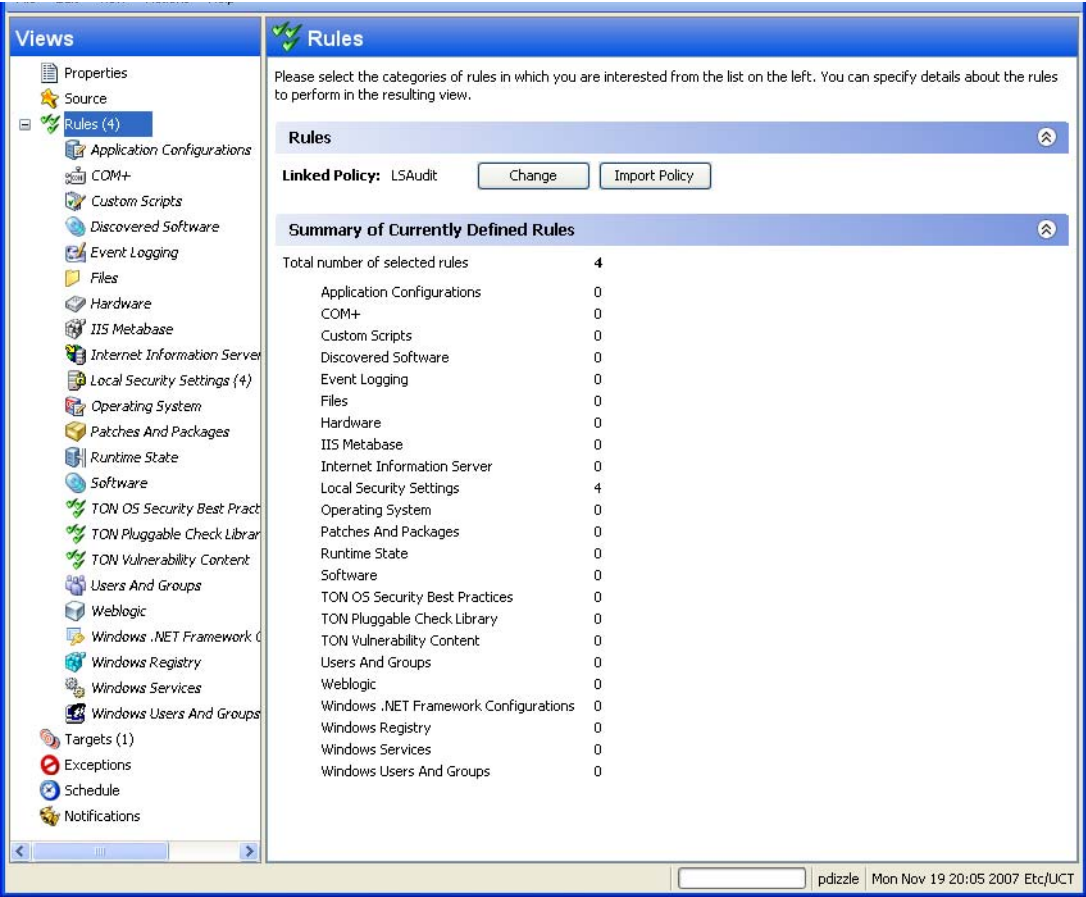
- **Properties:** The name and description of the audit.
- **Source:** The source of an audit can be a server, a snapshot, or no source at all. (However, some rules require a source.) Choosing a server as the source for an audit allows you to select server objects from that server as the basis of your audit. Choosing a snapshot as the source of an audit allows you to use the configuration values of the snapshot. Choose a snapshot specification as the source allows you to audit a server against itself over time.

For example, if you took a snapshot of a server, then used that snapshot specification as the source of the audit, every time you run the audit, you can compare the original state of the server against the server's actual configuration over time (using a recurring audit schedule). If you choose no source, then you can define only your own custom values for the audit or snapshot.

- **Rules:** A check on a particular server object with a desired value and an optional remediation value. For example, you might check to see if this server contains a specific Windows Service, and if found, determine if the service is turned off. For a description of server objects and rules, see "Server Objects Used in Audits and Snapshots" on page 160.
- **Targets:** The servers that the audit will check for compliance. You can choose as many servers and groups of servers as needed for an audit or snapshot.
- **Exceptions:** Servers and specific rules that will not be checked for compliance when the audit is run.
- **Schedule:** You can run an audit on a onetime basis, or on a recurring schedule. Audits that run on a recurring schedule appear as a single compliance column in the Compliance Dashboard.
- **Notifications:** You can send emails when the audit has finished running, and base the notification on the success, failure, or the completion of an audit job.

To configure an audit, select server configuration objects and then apply rules to those objects in order to define their desired configuration state. For example, Figure 2-2 shows an audit that has defined four rules. These rules will determine if target server configurations match the rules in the audit.

Figure 2-2: Audit Window Showing Elements of an Audit



## Creating an Audit

You can create an audit from several locations inside the SAS Client. You can choose to audit a specific server by selecting it from the server list, you can audit a group of servers, you can an audit from a snapshot, and so on.

You can create an audit from the following locations inside the SAS Client:

- From a managed server, using the selected server as the source of the audit. You can choose to run the audit on a single server or a group of servers.
- From the Device Groups list, choosing a group of servers as the target at the audit.
- From the Library, by creating a new audit.
- From a snapshot, by creating an audit based on the snapshot.
- From an audit policy, by creating an audit based on the audit policy.

### **Creating an Audit from a Server**

When you create a new audit from a managed server, the audit will use the selected server as the source of the audit. You can choose another server or snapshot for the audit source, if you want, or choose no source at all and define your own custom rules.



---

To audit a managed server, the server must be reachable and you must have access to the server.

---

To create an audit from a server, perform the following steps:

- 1** From the Navigation pane, select **Devices > Servers > All Managed Servers**.
- 2** Select a server, and then from the **Actions** menu, select **Create > Audit**.

For instructions on how to create an audit, see “Creating an Audit” on page 148.

### **Creating an Audit from a Group of Servers**

If you create an audit from a group of servers, then the audit will evaluate all the servers in that group. However, the audit will only evaluate those servers in a group to which your user has access.

To audit a group of servers, perform the following steps:

- 1** From the Navigation pane, select **Devices ► Device Groups**.
- 2** In the Navigation pane, browse until you see the group of servers (public or private) you want to audit.
- 3** Select the group of servers from inside the Content pane, right-click, and select **Create ► Audit**.
- 4** When you perform an audit by selecting a group of servers, the group of servers becomes the target. If the audit rule requires a source, you must supply one.

### **Creating an Audit from the Library**

To create a new audit from the SAS Client Library, perform the following steps:

- 1** From the Navigation pane, select **Library ► By Type ► Audit and Remediation**.
- 2** In the Navigation pane, select Audits, and then Windows or Unix.
- 3** Right-click inside the Content pane and from the **Actions** menu, select **New**.

### **Creating an Audit from a Snapshot**

You can select any snapshot in the Library and create an audit based on the server configuration captured in the snapshot. The snapshot will serve as the source of the audit, but you can also select another snapshot or server as the source after you create the new audit from the snapshot.

- 1** From the Navigation pane, select **Library ► By Type ► Audit and Remediation**.
- 2** In the Navigation pane, select Snapshots, then Windows or Unix.
- 3** From the Content pane, select a snapshot to create an audit from, right-click, and select **Create Audit**.

### **Creating an Audit from an Audit Policy**

Audit policies are designed to be used by audits. When you create an audit from an audit policy, the audit policy is linked to the audit. So, if any updates are made to the audit policy, those changes are automatically reflected in the audit.

- 1** From the Navigation pane, select **Library ► By Type ► Audit and Remediation**.
- 2** In the Navigation pane, select Audit Policies, and then Windows or Unix.
- 3** From the **Actions** menu, select **Create Audit**.

### **Saving an Audit as Audit Policy**

You can choose to save an audit as an audit policy, which will save only the rules from the audit and create a new audit policy. You can also save an audit using the “Save As” function to create a new audit with a new name.

All audit policies you create must be saved to the Library in a folder. You must have permissions to write to the folder you want to save the audit policy to. For more information on folder permissions, see *Opsware® SAS Policy Setter's Guide*, or contact your Opsware administrator.



---

For more information on creating, using, linking, and importing audit policies, see “Audit Policies” on page 140.

---

To use Save as to create an audit policy from an existing audit (or create a new audit), perform the following steps:

- 1** From inside the Audit or Snapshot Specification window, from the File menu, select Save As.
- 2** In the Save As window, enter a name. If you are renaming an audit or snapshot specification, you must use a unique name.
- 3** (Optional) Enter a description.
- 4** From the Type drop-down list, select either Audit or Audit Policy.
- 5** If you selected Audit Policy, from the Location section, click Select.
- 6** Select a folder in the Opsware SAS Client library to save the audit policy to. (You must have write permissions on the folder to save the audit policy.)

- 7 Click **OK**.

## Viewing Server Audit and Snapshot Usage

After you create and run an audit, you can view it from the All Managed Servers list or from the Device Explorer, and see all audits that are associated with a specific server.

### **Viewing a Server's Audit/Snapshot Usages from All Managed Servers**

To view a server's audit usage from the All Managed Servers list, perform the following steps:

- 1 From the Navigation pane, select **Devices > Servers > All Managed Servers**.
- 2 In the Content pane, select a server.
- 3 From the View drop-down list, select Audits or Snapshot Specifications. Notice that the lower Details pane shows information about audit and snapshot usage.
- 4 In the Details pane, if you selected Audits., you can choose one of the following options:
  - **Audit - Server is Target:** Shows all audits where the selected server is the target of the audit.
  - **Audit - Server is Source:** Shows all audits where the selected server is used as the source of the audit.
- 5 From any one of these views, you can select an audit or audit results, and perform actions from the Actions menu. For example, you can open an audit, re-run an audit, and so on.
- 6 If you selected Snapshot Specifications, then the Details pane shows all snapshot specifications that target the selected server.

### **Viewing a Server's Audit Usage from Device Explorer**

To view a server's audit usage from the Device Explorer, perform the following steps:

- 1 From the Navigation pane, select **Devices > All Managed Servers**.
- 2 In the Content pane, select a server, right-click, and select **Open**.
- 3 In the Device Explorer, from the Views pane, select Management Policies > Audits.
- 4 In the Content pane, from the Show drop-down list, select one of the following options:

- **Audit - Server is Target:** Shows all audits where the selected server is the target of the audit.
  - **Audit - Server is Source:** Shows all audits where the selected server is used as the source of the audit.
- 5** From any one of these views, you can select an audit and perform actions from the Actions menu. For example, you can open an audit, re-run an audit, and so on.
  - 6** Next, from the Views pane you can select Archived Audit Results to see all audit results associated with this server that have been archived. For more information, see “Archiving Audit Results” on page 235.

## Configuring an Audit

Configuring an audit consists of performing the following general steps:

- Name and describe the audit
- Select a source for the audit: a server, a snapshot, snapshot specification, or none
- Configure the audit rules
- Choose a target server, group of servers, or snapshot to audit
- Add audit rule exceptions (optional)
- Schedule the audit
- Set the Email Notification (optional)
- Save the audit

To configure an audit, perform the following steps:

- 1** Create the new audit from one of the methods described in “Creating an Audit” on page 148. The Audit window opens.
- 2** Enter the following information for the audit:
  - **Properties:** Enter a name and description for the audit.
  - **Source:** Every audit can use a server, snapshot, or snapshot specification as its source. (Or, you can choose no source and define your own rules.) If you use a server as the source, you can browse the server for values to define the audit’s rules. If you choose a snapshot, you will be limited to the rules in the snapshot and the snapshot results when you define the audit rules. If you choose a snapshot



specification, then the audit will compare the snapshot taken of the targets of the snapshot specification, and compare those against the targets of the audit. When you choose snapshot specification as the source, the rules in the snapshot are not editable. If you choose no source, you must define your own rules, or choose to link to an audit policy in the rules section. Some rules, however, require a source in order to be defined.

- **Rules:** Choose a rule category from the list to begin configuring your audit's rules. Each audit rule is unique and requires its own instructions. For information on how to configure individual audit rules, see "Audit and Remediation Rules" on page 163. If you want to use an audit policy to define the rules of your audit, click either Link Policy or Import Policy. When you link an audit policy, the audit maintains a direct connection with the audit policy. So if any changes are made to the policy, the audit will update with the new changes. If you import an audit policy, the audit will use all the rules defined in the policy but will not maintain a link to the audit policy. For information about audit policies, see "Audit Policies" on page 208.
- **Targets:** Choose the Targets of the audit. These are servers, groups of servers, or snapshots that you want the configured audit rules to evaluate and compare. To add a server or group of servers, click **Add**. To add a snapshot target, in the Snapshot Targets section, click **Add**.
- **Exceptions:** Click **Add** to add exceptions to the rules in your audit. In the Add Exception window, select a server or multiple servers (or device groups), and then select one or more rules you want to except from the chosen servers. You can except any of the rules in the audit from any of the target servers or snapshots. You can optionally add an explanation, a ticket ID, and an expiration date for the exception.
- **Schedule (Optional):** Choose whether you want to run the audit once, daily, weekly, monthly, or on a custom schedule. Parameters include:
  - **None:** No schedule will be set. If you want to run the audit immediately, or on a onetime basis, you have to select the audit, right-click, and select **Run Audit**.
  - **Daily:** Choose this option to run the audit on a daily basis.
  - **Weekly:** Choose the day of the week that you want the audit to run.
  - **Monthly:** Choose the months that you want the audit run.

- **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

```
0 0 * * 1-5
```

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 more information about crontab entry formats, consult the Unix man pages.

- **Time and Duration:** For each type of schedule, specify the hour, minute, day of the week, and month for the schedule to start. Unless you specify an end time, the audit will keep running indefinitely. To choose an end date, select End. From the calendar selector, choose an end date. The Time Zone is set according to the time zone set in your user profile.
- **Notifications:** Enter email addresses to notify people when the audit job finishes running. You can choose to send the email on both the success and the failure of the audit job (not the success of the audit rules). To add an email address, click Add Notification rule. (This is only relevant if the audit is set to run on a recurring schedule.)

**3** When you have finished configuring the audit, from the **File** menu, select **Save**.

### **Audit Sources: Server, Snapshot, or Snapshot Specification**

You have two options for choosing a source for an audit or snapshot specification: a server, a snapshot, or a snapshot specification. The source of an audit determines what rules you are able to select from and configure in your audit or snapshot specification. Choosing a source depends on the purpose of your audit or snapshot specification:

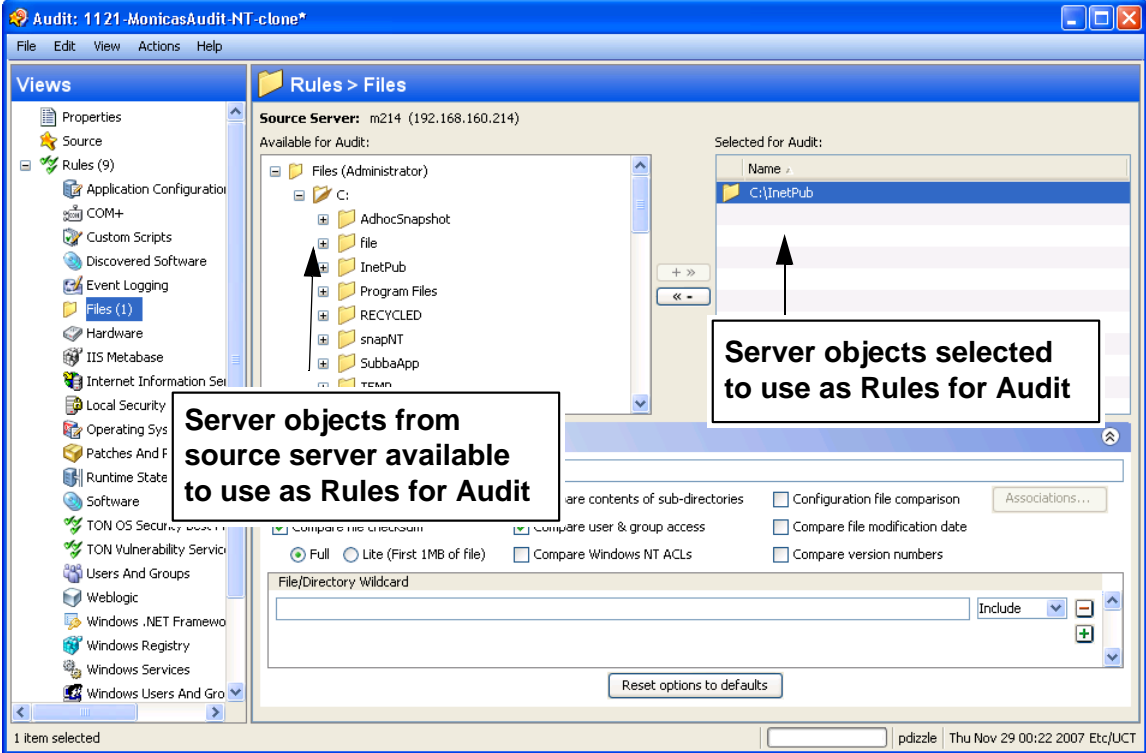
#### **Server as Source for an Audit or Snapshot Specification**

Choose a server as the source of an audit if you know that specific server contains the desired server objects that you want to add to the audit or snapshot specification. For example, if you are interested in auditing or taking a snapshot of application configuration files for an Apache Web Server (for example, httpd.conf) on some target servers, choose as the source of your audit – a server that you know has Apache installed on it and that is configured correctly.

Remember that you can choose several different source servers as you build your audit or snapshot specification rules. In fact, you can choose a different source for each server object rule.

When you choose a server as the source for an audit, Figure 2-3 shows what you see in the audit or snapshot specification window's Content pane (right side of window):

Figure 2-3: Server as Source of Audit for Building Audit Rules

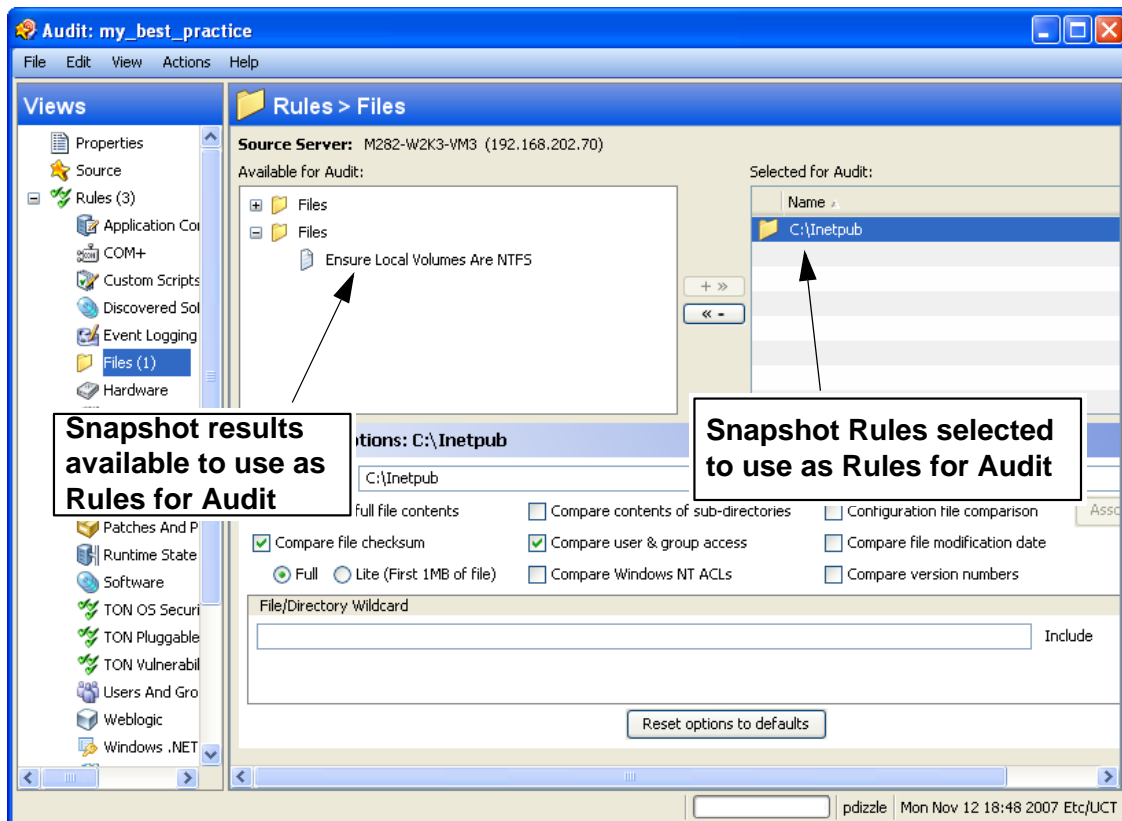


### Snapshot as Source for an Audit or Snapshot Specification

Choose this option if you have a snapshot of a server that was in a known good state (a “golden” server configuration), and you would like to compare that snapshot with other servers in an audit. Or, choose this option to use the captured server values to take a snapshot of another server. Using a snapshot as the source for an audit or snapshot specification allows you to choose both the results and the rules of the original snapshot specification that the snapshot was based on.

Figure 2-4 displays the choices you have for building audit or snapshot specification rules when you use a snapshot as the source. You can choose from the snapshot's results and the snapshot's rules.

Figure 2-4: Snapshot as Source of Audit: Available Server Objects to Build Audit Rules



### Snapshot Specification as Source of Audit – Reflexive Auditing

Choose this option if you want to keep track of a server's configuration over time and monitor any changes that occur. For example, you might want to keep track of a specific application to make sure that its configuration remains correct over a period of time. If this application runs on several servers, you can create a snapshot specification that defines a desired state of server configuration, and then run the snapshot.

Next, you can create an audit and use the original snapshot specification as the source for your audit. Each server that was targeted by the snapshot are now also included as targets of the audit. Next, when you run the audit (either on-demand or on a scheduled

basis), each server's current configuration will be compared with the state originally captured when you took the initial snapshot. Any changes are displayed in the audit results window.

### **Rules That Use a Source Value From Source Server**

Most rules require a source in order to define them, except the following rules:

- Any of the pre-configured rules that you do not set the value to derive from a source (server or snapshot or snapshot specification)
- Custom Scripts rules that you do not set the compare value to derive from a source (server or snapshot or snapshot specification)

You cannot save rule without giving a source if the rules specified require a source. You must select a source for all comparison checks and for rules that compare against a source value.

### **Opware Audit and Remediation Rules**

An audit enables you to determine how your servers are configured, and whether or not those servers are configured correctly – that is, as defined in an audit or audit policy. You achieve this goal by creating rules about server objects. You can gather information about the server objects listed in Table 2-1 and either take a picture of their current state – in a snapshot – or define the desired configuration state for these objects – in an audit or an audit policy. (For a list of all server objects you can configure for an audit, see “Server Objects Used in Audits and Snapshots” on page 160.)

In an audit and audit policy, you can also define what, if any, remediation value you would like the object to have. This is used only if a server object is found to be different than the desired state. The remediation value is not implemented automatically, but rather manually after the audit has been run.

An audit rule consists of a server object (file system, IIS Metabase entry, and so on), the specific thing about the object that you want to check (the specific files or directories you want to check), the desired state of the object, and a remediation value should the server configuration differ from the audit rule (optional).

An audit rule consists of the following components:


- **Server Object:** This is a specific server configuration category that an audit can evaluate, such as a server's file system, application configurations, hardware, software (installed patches and packages), Windows Registry entries, and so on. A server object usually consists of several other things that you can check. For

example, for the windows services server object, you might want to know if a specific service exists on target servers and whether or not the service is enabled or disabled.

- **Target Value:** This is the specific server configuration object element. For example, you might want to determine if a specific directory exists on a server, or if an application is configured properly, and so on.
- **Remediation Value:** This is the value that you want to change for the server object, if it is found to be different than desired. This value is not implemented automatically; you can make the remediation change after the audit has run.

Figure 2-5 illustrates an audit rule defined for a Windows Service named File Replication.

Figure 2-5: Windows Services Audit Rule



In this example, the audit rule has been configured in the following manner:

- **Available for Audit:** Lists all services from the source server available to be added to the audit, plus specific Windows services rules provided by Opsware.
- **Selected for Audit:** The service name File Replication has been chosen.
- **Description:** Describes what is being checked on the target server. In this case, the audit will check to see if the service is enabled or disabled.

- **Target Value:** This is the value compared against the target server. In this example, the user has set Service Disabled. This means that the audit will check to see if this service is disabled. If the service is in fact enabled, the audit results will indicate the variance, and the configuration would be considered out of compliance with CIS standards.

Depending on the type of check being done on a server, the target value can contain an operator (equals, greater than, and so on), a reference (use from the source of the audit), your own or a preset list of Values (for predefined rules, these values are built in), or a custom attribute that was exists on the target server.

- **Remediation Value:** The remediation value determines the action to take if the service on the target server does not match the value you defined in the audit. Remediation values can be derived from a prebuilt or user-entered Value, a Server Attribute on the target server, or a custom attribute that exists on the target server.

### Server Objects Used in Audits and Snapshots

Table 2-1 lists all server objects that you can create rules for inside an audit or a snapshot specification. Some server object values are captured and audited live and some objects are captured from the Model Repository.

Table 2-1: Audit and Remediation Server Objects

SERVER OBJECT	DESCRIPTION	CAPTURED LIVE AND/OR FROM MODEL REPOSITORY
<b>Application Configurations</b>	Contents of application configuration files and their values.	Live
<b>Windows COM+</b>	COM+ objects and component categories.	Live
<b>Custom Scripts</b>	Write your own custom scripts to retrieve information from a server and compare contents. For example, you can run a script to gather output from a custom application and evaluate returned output against values set in the audit. (Python 1.5.2 only for python scripts.)	Live



Table 2-1: Audit and Remediation Server Objects (continued)

SERVER OBJECT	DESCRIPTION	CAPTURED LIVE AND/OR FROM MODEL REPOSITORY
<b>Files</b>	Contents of files and directories (and subdirectories), user and group access, checksum for files, file modification date, and Windows ACLs (Windows only).	Live
<b>Hardware</b>	CPU, storage devices, and memory.	Model Repository
<b>IIS Metabase</b>	IIS Metabase objects and configuration values to snapshot or audit.	Live
<b>Internet Information Server</b>	Real time information about IIS for a Windows server, such as server name, server type, server state, log file path, document file path, and so on.	Live
<b>Local Security Settings</b>	Real time information about security settings, including security settings such as password policy, audit policy, user rights, and security options.	Live
<b>Registered Software</b>	All installed packages or patches actually installed on a source server, whether or not they have been registered by the Opsware model repository.	Live
<b>Runtime State</b>	The Runtime State window rule allows you to use time information about run time data for an audit rule, such as DNS servers, Routes, and Processes for every managed server.	Live
<b>Software</b>	All patched and packages that have been registered with the Opsware model repository.	Model Repository
<b>Storage</b>	Information related to storage devices and SAN devices and connections in your data center (if your core is ASAS-enabled).	Live

Table 2-1: Audit and Remediation Server Objects (continued)

SERVER OBJECT	DESCRIPTION	CAPTURED LIVE AND/OR FROM MODEL REPOSITORY
<b>TON (The Opware Network) Rules</b>	If you are subscribed to TON, you have access to many different types of audit rules (called "pluggable checks". The exact kind of rules you have access to depend on your subscription, but can include such rules as the latest patch supplements for Microsoft Windows, current regulatory compliance policies (for example, FISMA, Sarbanes-Oxley), user-created rules from the TON developer community, daily updated vulnerability content, and so on.	Live
<b>Users and Groups</b>	Compare information about users and groups on servers, such as user name for last login, whether or not CTRL + ALT + DELETE is enabled, and so on.	Live
<b>Windows .NET Framework Configuration</b>	Real time information about Assembly Cache and Configured Assembly List, such as assembly name, version, locale, public key token, cache file (GAC or ZAP), processor architecture, custom, and file name.  For every Configured Assembly List, you can use information such as assembly name, public key token, codebases, binding policy, file name, file data.	Live
<b>Windows Registry</b>	Select Windows Registry directories or registry key values to capture and compare.	Live
<b>Windows Services</b>	Select Windows services.	Live
<b>Windows Users and Groups</b>	Users and groups information on a Windows Unix servers.	Live



---

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 Device Explorer.

---



---

Opsware Audit and Remediation does not support device files or sockets.

---

## Audit and Remediation Rules

Creating an audit (or snapshot specification) requires configuring Opsware Audit and Remediation rules, which define:

- The type of server object to snapshot or audit and compare – objects such as the server's file system, hardware information, application configurations, installed patches or software, users and groups, and so on.
- Information about that object to audit or snapshot. For example, for a server's file system, you can capture Windows NT file's Access Level Controls. For an application, you can capture the application configuration values you want to snapshot or audit, plus any remediation values to specify if differences are discovered between the rule and the actual value on the target server.

A rule can contain a custom script that seeks to determine if all the passwords stored in a file match a certain character length, or a rule can include a check to determine if a particular Windows Service is running or disabled on a server. For some rules, you can also specify the remediation value for the server object if the value defined in the audit or snapshot differs from the server's value after the audit has run. For example, if a Windows Service is disabled, you can specify that the Remediation value should restart the service.

Remediation values are implemented manually, after the audit has run, from the Audit Results window. For more information on how to remediate audit results, see "Viewing and Remediating Audit Results" on page 220.

### Configuration Rules: Expected (Target) and Remediation Values

Some rules are a very simple to configure and define and do not require anything more than selecting the server objects that you want to snapshot or audit. Some rules might check to determine if a value or property exists on a configuration file on a server, without the need for any advanced parameters. For example, Opsware Audit and Remediation


rules for the Software server object evaluate the patches or packages that are installed on the target servers. The Hardware rule allows you to check the CPU, memory, or storage values that exist on target servers. In this case, no extra rule parameters are necessary. Other rules are more complex and require more advanced configuration, such as specifying an expression that looks for a range of values and specifies a remediation that replaces undesired values.

**Example Rule: Event Logging**

Event Logging requires that operators and reference values (user-entered values, custom attributes from the source server, or server attributes) be defined. For example, you can choose to configure an Event Logging rule that will check the maximum application event log size. CIS standards and Microsoft recommends that this value to 16MB. You can

define the audit rule to determine if this value is no more than 16MB on your target servers. You can also set the remediation value to be 16MB, if the value found on a target server is greater than 16MB.

Figure 2-6: Example Audit Rule for Event Logging Server Object



In the example shown in Figure 2-6, the user has chosen to audit the Event Logging setting of “Maximum Application Event Log Size.” (This audit rule is one of the many predefined rules that come as part of the SAS Client product distribution.)

The top left side of the rules pane, Available for Audit, shows all Event logging objects available from the source server to add to the audit. The top right, Selected for Audit, shows all Event Logging rules that have been selected for the audit.

This rule consists of the following parameters:

- **Rule Details:** Describes this setting and the CIS and Microsoft recommended value, which is 16MB.
- **Target Value:** Allows you to define a target value, which is the value you expect to find on the server.
- **Operator:** Uses an operator to set the expression. Operators include equals (=), less than (<), greater than (>), and so on.
- **Reference:** Choose the source of the script output. You can choose from the following options:
  - **Source:** Takes the value of this setting from the source of the audit, either a server or a snapshot, or from the source of the snapshot specification, a server.  
If you choose a server as source for an audit or snapshot specification, then you can select from all the objects available on that server.  
If you choose a snapshot as your source for an audit, then you will only be able to select the snapshot rules and snapshot results for the audit. (You can only choose a server as the source for a snapshot specification.)
  - **Value:** Allows you to enter your own value.
  - **Server Attributes:** Common server attributes from the Opsware model.
  - **Custom Attributes:** Derives from the target server. (For the application configuration and custom script rule, if you choose a custom attribute for the rule definition, this custom attribute must also exist on the target servers.)
- **Value:** Either a user-entered value, a server attribute from the Opsware model, or a custom attribute from a target server.
- **Remediation Value:** The value that will replace those found on the target server that do not match the target value specified. The remediation value will not be implemented automatically. Rather, you must manually choose to remediate the value from the Audit Results window after the audit has run.

After you select parameters for the rule, the Value field will show the desired value for the selected configuration file. If the value set in the rule does not match the value on the target of the audit, then you can specify in the Remediate section.

## Configuring Specific Audit and Snapshot Rules

For information on rules you can set for each type of server object, see the section for the specific server object that you want to configure a rule for, listed below:

- Configuring Application Configuration Rule
- Configuring COM+ Rule
- Configuring Custom Scripts Rule
- Configuring Files Rule
- Configuring Hardware Rule
- Configuring IIS Metabase Rule
- Configuring Internet Information Server Rule
- Configuring Local Security Settings Rule
- Configuring Registered Software Rule
- Configuring Runtime State Rule
- Configuring Software Rule
- Configuring Storage Rule
- Configuring Windows .NET Framework Configurations Rule
- Configuring Windows Registry Rule
- Configuring Windows Services Rule
- Configuring Windows/UNIX Users and Groups Rule
- Configuring The Opsware Network (TON) Rule



You must have permissions to create and configure Opsware Audit and Remediation rules. To obtain these permissions, contact your Opsware administrator. See the *Opsware<sup>®</sup> SAS Policy Setter's Guide* for more information.

---



Some Opsware SAS cores may contain legacy content, specifically, Event Logging, Operating System, and Users and Groups rules with custom checks. These checks have been integrated into the CIS policies available in TON.

---

## Configuring Application Configuration Rule

The application configuration rule inside an audit, snapshot specification, or audit policy allows you to configure values for application configuration files on a target server. You can define rules you want to capture and the configuration file values that you want to check.

You can choose from a list of predefined application configurations for the configuration file you would like to audit or take a snapshot of. You can also choose from custom application configurations that a user in your organization has created and made available for usage in an audit, snapshot specification, or audit policy.

An application configuration is a template (or collection of templates) that models the information found inside of a configuration file for an application. When you choose an application configuration inside an audit, snapshot specification, or audit policy and click **View**, you will see the contents of the configuration file from the source of the audit. All key-value pairs that you are able to add to the audit rule will display.

The information displayed inside an audit windows depends on the source of the audit or audit policy (or the target for a snapshot specification):

- If you choose a server as the source of the audit or audit policy, then the application configuration values displayed in the audit rule will be those of the configuration file on the server, as filtered through the application configuration template.
- If you choose a snapshot as the source of the audit or audit policy, then you will only be able to modify the values that were captured at the time the snapshot was taken.
- If you do not choose any source, then you will not be able to configure a rule for the application configuration file.
- If you choose to configure an application configuration in a snapshot specification, then the values of the configuration will derive from the target server.



You will only see values of the source configuration file that have been modelled in the application configuration. If the application configuration is customized and has no name-value pair defined (but the value exists in the source configuration file), you will not see it in the audit or audit policy.

---



After you view the contents of the source application configuration file, you can select values and define target values – the values that the audit should look for on the target server. You can also define remediation values should the audit find a difference.

### **Creating an Application Configuration Rule**

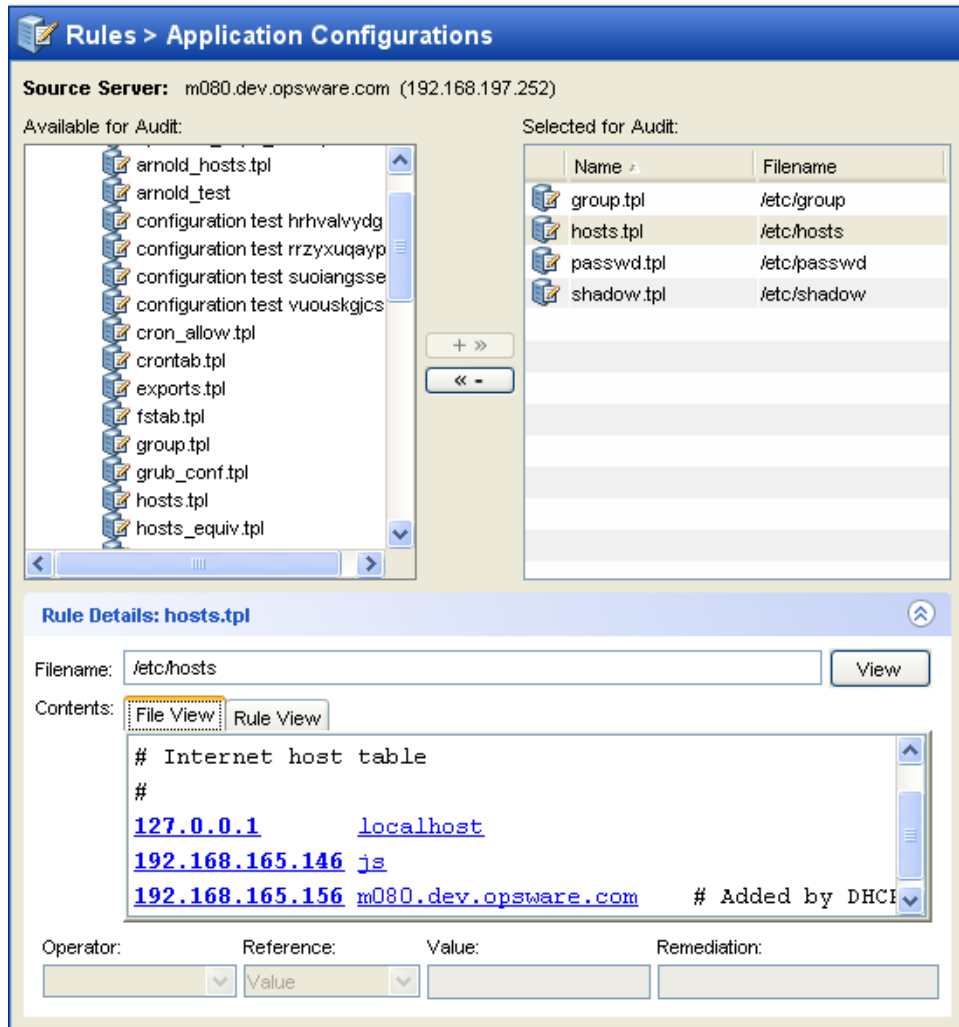
A useful way to understand how to configure an application configuration rule is to look at an example. Let's say you want to create an audit rule for a UNIX hosts file (/tmp/hosts). You know that the UNIX hosts file on a specific server represents the ideal state of the hosts file configuration, so you would choose that server as the source for your audit. This means that you will use the hosts file on that server as the basis of your application configuration audit rule.

To create an application configuration rule, perform the following tasks:

- 1** Create the new audit from any one of the methods for creating an audit listed at "Creating an Audit" on page 148.
- 2** Select a source for the audit – No Source, Server, Snapshot, or Snapshot Specification. The source selected for the audit will determine what types of rules, if any, you can create for an application configuration. You must choose a source, or you will not be able to configure the application configuration rule.
- 3** In the Audit window, from the View pane, select Rules ► Application Configurations.
- 4** In the content pane of the audit or snapshot specification window, expand the top level node in the Available section and select an application configuration.
- 5** Click the right arrow button to move the application configuration into the Selected for Audit section.
- 6** In the Selected for Audit or Snapshot Specification section, select the application configuration.
- 7** Click **View**. (If you cannot view the contents of the configuration file, you might need to enter the correct path in the Filename section.) You see the contents of the configuration file in the File View tab.

For example, if you view a UNIX hosts file, you would see something similar to that shown in Figure 2-5:

Figure 2-7: Application Configuration Audit Rule for hosts File



You can see the contents – the IP address/host name pairs – from the source hosts file, highlighted in blue text.

- 8** In order to create an audit rule for this configuration file, you need to choose a key-value pair from the hosts file on the source server (the server you choose as the source for the audit).

- 9** To create this rule, first select an IP addresses in the File View tab content area. In the example in Figure 2-5, you would select an IP address such as 127.0.0.1. After you select the IP address, the element becomes highlighted in dark blue. This means that the element is selected but has not yet had a rule created from it.

(For more information on the color scheme used when configuring an application configuration audit rule, see Table 2-2 on page 172.)

To create a rule that will look for this IP address (127.0.0.1) in the hosts file on the target server, select the IP address in the contents area. Notice that the value in the Operator field below is set to blank. This means that the value has not yet been added to the rule. To add the value to the rule, you can either double-click it, or enter the following parameters in the rule expression area below the contents:

- **Operator:** Choose = (equals). When you change the operator to =, then the value immediately becomes added to the rule. If you change the operator back to no selection, then the value is immediately removed from the rule.
- **Reference:** Choose Value.
- **Value:** Enter 127.0.0.1.
- **Remediation:** Enter 127.0.0.1.

This expresses that you want to look for an IP address with the value of 127.0.0.1. If this is not found, then the remediation should be 127.0.0.1, so you can add this to any host files on the target servers that do not contain this IP address.

- 10** Next, select the host name in the File View tab area, and in the Rule section, set the following parameters:

- **Operator:** Choose = (equals).
- **Reference:** Choose Value. (If you choose a custom attribute here for the rule definition, this custom attribute must also exist on the target servers or the audit for this rule will fail.)
- **Value:** Choose host.
- **Remediation:** Choose host. This adds the final part of the rule that will check the target server for the key-value pair of IP address 127.0.0.1 matched with host.

**11** Now, select the Rules View tab. The rule will be expressed as:

“Check that there is an entry where IP address is equal to value 127.0.0.1 and Hostnames contains an entry equal to value host.”

This rule is what will be used to audit the hosts file on the target server or snapshot specification.

**12** To configure more application configuration rules, select more application configurations from the Available for Audit section.

**13** To finish configuring the audit, define other rules and set the target servers, schedule, and notification for the audit.

**14** Save the audit.

**15** To run the audit, from the **Actions** menu, select **Run audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

### **Application Configuration Audit Rule Color Scheme**

When you first view an application configuration, all elements that can be used to build an audit rule will appear in blue underlined text. After you start selecting and building rules, then the colors will change. Table 2-2 describes the color scheme used for configuring application configuration audit rules.

Table 2-2: Application Configuration Audit Rule Color Scheme

TEXT COLOR	DESCRIPTION
Blue underlined	This shows all elements in the source configuration file that can be used in a rule.
Dark blue	This shows an element is selected but has no rule has been associated with it.
Light blue	This shows all that you add an element to a rule.
Medium blue	This shows all that an element is both selected and has a rule associated with it.

Table 2-2: Application Configuration Audit Rule Color Scheme (continued)

TEXT COLOR	DESCRIPTION
Green	<p>This shows all that the element is a primary key and is related to the current selected element. This means that the element will be used in the same rule that the current selected element will be used in.</p> <p>If the currently selected element is given a comparison value (=, contains, matches...) then the other elements with the green text will automatically be given a comparison value of “=”.</p> <p>An example of this would be:</p> <pre>127.0.0.1    localhost</pre> <p>If localhost is selected, then 127.0.0.1 would be green. If localhost is given a comparison value, then 127.0.0.1 will also be given an automatic comparison value, giving you a rule such as:</p> <p>There is an entry where ip is equal to 127.0.0.1 AND hostname is equal to localhost.</p>
Bold	This represents a primary key.
Italicized	This shows a custom attribute or Opsware attribute.

## Configuring COM+ Rule

To configure a Windows COM+ object rule, select the source COM+ objects or COM+ object categories that you want to audit or snapshot on a target server. The COM+ rule also checks Access Control Levels (ACLs) for the selected object as well as any that are inherited.

COM+ objects are categorized based on attributes of the object, where the COM+ object specifies zero or more categories. The audit or snapshot window displays all COM+ objects in one node in the Rules section of the COM+ object tree. To add a COM+ rule to the audit or snapshot, select it and click the right arrow button.

If you would like to be able to remediate COM+ rules in your audit or snapshot results, select the "Archive all associated files" option when you select the COM+ object or category.

Selecting the "Archive all associated files" option will also include all AccessPermissions and LaunchPermissions associated with the COM+ object in the audit or snapshot rule, including those that are inherited parent COM+ objects.

To configure a COM+ rule, perform the following steps:

- 1** Create the new audit using one of the methods for creating an audit listed in "Creating an Audit" on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► COM+.
- 4** In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select a COM+ object or object category.
- 5** Click the right arrow button to move the COM+ object or object category into the Selected for Audit section. All COM+ object or object categories you select will be audited on the target servers or snapshot specification.
- 6** You can now choose an option from the bottom of the rule window:
  - Select the Archive all associated files option if you want to be able to remediate COM+ rules in your audit or snapshot results.
  - Select Compare only the file name and not the full pathname if you want the COM+ rule to check only the selected filename and not the full path.

- 7** To finish configuring the audit, define any other COM+ object or object category rules you want and set the target servers, schedule, and notification for the audit.
- 8** If you want to be able to
- 9** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 10** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

### Configuring Custom Scripts Rule

The custom script rule allows you to define your own script (batch, Python 1.5.2, or Visual Basic) to get and compare values used in an audit, audit policy, or snapshot specification. You can also write your own remediation scripts.

When you configure a custom script rule, you specify the target value, which is the expected values you want the script to return. The audit can gather this information in two ways:

- **Comparison-Based Audit:** Execute the script on the source server. The return values from the script (exit code or standard output) are compared with the output of the script after it has run on the target server or servers. This option is named: Source.
- **Value-Based Audit:** Specify your own value. This is compared with the output of the script after it has run on the target server. You can enter this value manually, if you know what the expected results of the script should be, or, you can execute the script on the source server and use those return values. When the audit is run, this value is compared with the returned results from the script after it has executed on the target server or servers. The option is named “Value.”

For an audit, you can also configure a remediation script, which can be used if differences are found between the rule and the value returned after the script has run on the target server.

For a snapshot, the script results will be generated by running the script (as defined in the rule detail) on target servers, and then captured in the snapshot. When you set up a snapshot specification, you can also add a remediation script. This type of script can be used to force remediation on target servers. You can execute the snapshot’s remediation script on target servers on an individual server basis from the Snapshot window.

To configure a custom script rule, perform the following steps:

- 1** Create the new audit using one of the methods for creating an audit in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** To build a script and define the audit rule, you can choose the following options:

#### **Source**

- **Rules:** Click **Add Rule** to add a new custom script rule.

#### **Rule Details**

- **Name:** Enter a name for the script.
- **Type of Script:** Choose from Batch, Python 1.5.2, PowerShell, or Visual Basic (VBS).
- **Script:** Type or copy and paste the script contents here. Or, click **Import Script** to import a script from your computer.

#### **Success Criteria**


- **Output:** Either Exit Code or Standard Output.
- **Operator:** Choose an Operator, such as equals (=), not equals (<>), less than (<()), greater than (>), and so one.
- **Reference:** Choose the source of the script output.
  - **Source:** Select this option if you want the rule to execute the script on the source when an audit is run, and gets the value that the script requests. It will then compare that value with the value retrieved from the script that was run on the target server.

If you choose this option for a snapshot specification, then the script will run on the target, and the results of the script execution will be captured in the snapshot (results).

If the source of the audit is a snapshot, then the custom script rule will use the custom script definition configured in the snapshot specification.

- **Value:** Enter your own value. This option uses the value you enter and compares it with the value returned from the script after it is run on the target server. Using this option means that the script does not run on the source server at audit



runtime. However, you can get the output from the script immediately from the source server, if you click the eyedropper  icon. The returned value is displayed in the text box, which you can accept as is or edit to your liking.

If the source of the audit is a snapshot, then the custom script rule will use the Custom Script definition configured in the snapshot specification.

- **Server Attribute:** Select this option to compare a server attribute found on the source server with the output from the script that is run on the target server.
- **Custom Attribute:** Select this option to compare a custom attribute found on the target server with the output from the script that is run on the target server. Custom attributes for this option derive from the selected source server for the audit.

If you choose a custom attribute here for the rule definition, this custom attribute must also exist on the target servers or the audit for this rule will fail.

If you do not choose a source for the audit, then this list will be empty.

### Remediation

- **Type of Script:** Choose from Batch, Python 1.5.2, PowerShell, or Visual Basic (VBS).
- **Script:** Type or copy and paste the script contents here. Or, click **Import Script** to import a script from your computer.

- 4 (Optional) You can add a remediation script to run if the audit comparison fails. The remediation will not be applied automatically; you can only run the remediation script from the audit results after the audit has run.

For a snapshot, the remediation script you define here can be executed on target servers on an individual server basis.

- 5 To finish configuring the audit, set the target servers, schedule, and notification for the audit.
- 6 To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 7 To run the audit, from the **Actions** menu, select **Run audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

## Configuring Files Rule

The file system rule allows you to audit a server's file system and directory structure for the following evaluations:

- **Directory name:** Displays the absolute path of the selected file or directory
- **Archive the full file contents:** Archives the entire file. Use this option if you want to be able to remediate and view any files differences found between the rule and the target file.
- **Compare file checksum:** Performs a Checksum on the contents of the selected file or files in a directory. You can choose to audit the entire contents of the file, or just the first 1MB of the file.
- **Compare contents of subdirectories:** Includes contents of all subdirectories for a selected file system folder to the audit.
- **Compare user and group access:** Audits the user and group access related to the file and directories.
- **Windows NT ACLs (Access Control List):** Audits the Windows Access Control List for files and directories.
- **Configuration file comparison:** Allows you to use an application configuration to evaluate configuration files on a target server. Selecting this option (and clicking **Associations**) enables to you to utilize a configuration template to compare any differences in values between a source configuration file and one on a target server.

For example, if you want to manage a Windows hosts file, you can select the target file in the audit rule, then select Config File Comparison and click **Associations**. You can select an AppConfig template that manages a Windows hosts file on a source server that you know represents a desired configuration. When the audit is run, the rule will use the configuration template to compare both files and display any differences in the audit results.

- **Compare file modification date:** Audits the file modification date to use for file or folder comparison.
- **Compare version numbers:** For specific Windows file types – .exe, .dll, .ocx, .olb, .scr, .rll, .sys, .drv, .acm – the author of the file can set a File Version and Product Version. This option, compares these version numbers and if they are different, the rule is considered non-compliant and the actual values on the target file can be

viewed in the audit results. (Note that not all files with these extensions always have a product or file version attribute, but some can.)

- **File/Directory wildcard:** Allows you to specify directories and files in the file system you want included in and excluded from the audit. For more information on how this option works, see “File Inclusion and Exclusion Rules” on page 198.

There are two categories of file system rules that appear in the Available for Audit section of the Audit window. You can define the following specifications in an audit or snapshot:

- **File System:** These are comparison-based rules, which enable you to select a file system file or directory from the source of the audit or snapshot specification and compare these with the target servers. The purpose of this rule is to determine that the file or directory exists and its properties. You cannot set a target or remediation value in the rule.
- **Specific File System Rules:** These are value-based file system rules built into the SAS Client. They allow you to configure expected (target) and remediation values.

To configure file rules, perform the following steps:

- 1** Create the new audit using one of the methods in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Files.
- 4** In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select a folder or file to create a rule for.
- 5** Click the right arrow button to move the folder or file into the Selected for Audit section. All folders or files that you select will be used to audit or snapshot the target server.
- 6** In the Selected for Audit section, select a folder or file to apply a rule to.
- 7** In the Directory Options section, select file system rule options to apply to the selected folder or file. If you would like to reset the original settings of the source file system, select the Reset options to match those of the File System option.
- 8** (Optional) For folders, you can select a File/directory Wildcard option to specify files and directories that you want to include or exclude from the audit.

Click the **plus (+)** button to add a new rule, or click the **minus (-)** button to remove a rule. For more information on how to enter files and directories and how this affects the audit, see "File Inclusion and Exclusion Rules" on page 198.

- 9** (Optional) If you want to use an application configuration to compare configuration files, select Configuration file comparison, and click **Associations**.
- 10** In the Edit AppConfig Associations window, from the Installed AppConfig Templates, select a template you would like to use to compare a source and a target configuration file.
- 11** In the Associated Files section, you can use the default path to the source configuration file, or edit the path. You can click the plus button to add another path to a source configuration file you want to compare with a configuration file on the target.
- 12** When you are finished, click **OK**.
- 13** To finish configuring the audit, set the target servers, schedule, and notification for the audit.
- 14** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see "Saving an Audit as Audit Policy" on page 150.
- 15** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see "Creating an Audit Policy" on page 209.

## Configuring Hardware Rule

Configuring a hardware rule allows you to audit the following information about a server's hardware:

- **Interfaces:** Compares duplex mismatch and all network interfaces on a server.
- **CPU:** Compare CPU type and specification of target server.
- **Memory:** Compare memory of the target server.
- **Storage:** Compare storage capacity on the target server.
- **Interfaces:** Compare all network interfaces attached to the device.

To configure hardware rules, perform the following steps:

- 1** Create the new audit using one of the methods for creating an audit listed in "Creating an Audit" on page 148.

- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Hardware.
- 4** In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select a hardware category to create a rule for.
- 5** Click the right arrow button to move the hardware item into the Selected for Audit section. All items that you select will be used to audit or snapshot the target server.
- 6** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see "Saving an Audit as Audit Policy" on page 150.
- 8** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see "Creating an Audit Policy" on page 209.

### Configuring IIS Metabase Rule

The IIS Metabase audit rule allows you to select IIS Metabase objects and objects folders to compare in your audit. The audit will capture IIS Metabase object property information such as ID, name, path, attributes, and so on.

To configure IIS Metabase rules, perform the following steps:

- 1** Create the new audit using one of the methods for creating an audit listed at "Creating an Audit" on page 148.
- 2** Select an Audit Source: Server, Snapshot, Snapshot Specification, or No source. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► IIS Metabase.
- 4** In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select an IIS Metabase folder or object to create a rule for.
- 5** Click the right arrow button to move the IIS Metabase folder or object into the Selected for Audit section. All items you select will be used to audit or snapshot the target server.
- 6** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.

- 7** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 8** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

## Configuring Internet Information Server Rule


The Internet Information Server rule allow you to use real time information about IIS for your audit, such as a Windows server, such as server name, server type, server state, log file path, document file path, and so on.

To configure the Internet Information Server rule, perform the following steps:

- 1** Create the new audit using one of the methods in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Internet Information Server.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select an Internet Information Server rule that you want to create a rule from.
- 5** Click the right arrow button to move the rule object into the Selected for Audit section. All Internet Information Server rules that you configure will be audited on the target servers or snapshot specification.
- 6** For each rule, select one of the following check types:
  - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array.
  - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.

- **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.

**7** You can also configure a rule based upon a wildcard search by selecting the

Wildcard rule object \*. When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.

For example, you could enter simply \* which would match everything on the target, P\* would match all objects that begin with a capital P, while \*P would match all elements ending with uppercase character 'P'.

After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.

It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.

- 8** To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 10** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.


### Configuring Local Security Settings Rule

The Local Security Settings rule allows you to use real time information about security settings, such as password policy, audit policy, user rights, and security options in your rule.

To configure the Local Security Settings rule, perform the following steps:

- 1** Create the new audit using one of the methods in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User’s and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Local Security Settings.

- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select an Internet Information Server rule that you want to create a rule from.
- 5** Click the right arrow button to move the rule object into the Selected for Audit section. All Internet Information Server rules that you configure will be audited on the target servers or snapshot specification.
- 6** For each rule, select one of the following check types:
  - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array.
  - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.
  - **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.

- 7** You can also configure a rule based upon a wildcard search by selecting the Wildcard rule object \*. When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.

For example, you could enter simply \* which would match everything on the target, P\* would match all objects that begin with a capital P, while \*P would match all elements ending with uppercase character 'P'.

After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.

It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.



- 8** To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 10** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

### Configuring Registered Software Rule


The Registered Software rule allows you to audit use all installed packages or patches actually installed on a source server to build your rule, whether or not the patches or packaged have been registered by the Opsware model repository.

To configure the Registered Software rule, perform the following steps:

- 1** Create the new audit using one of the methods in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Registered Software.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a patch or a package that you want to create a rule from.
- 5** Click the right arrow button to move the rule object into the Selected for Audit section. All rules that you configure will be audited on the target servers or snapshot specification.
- 6** For each rule, select one of the following check types:
  - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of ‘true’ and ‘false’), Date (a date compare, not a time of day compare), or an Array.
  - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.

- **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.

**7** You can also configure a rule based upon a wildcard search by selecting the

Wildcard rule object \*. When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.

For example, you could enter simply \* which would match everything on the target, P\* would match all objects that begin with a capital P, while \*P would match all elements ending with uppercase character 'P'.

After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.

It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.


- 8** To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 10** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

## Configuring Runtime State Rule

The Runtime State rule allows you to use time information about run time data for an audit rule, such as DNS servers, Routes, and Processes for every managed server.

To configure the Runtime State rule, perform the following steps:

- 1** Create the new audit using one of the methods in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Runtime State.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Runtime State rule that you want to create a rule from.

- 5** Click the right arrow button to move the rule object into the Selected for Audit section. All Runtime State rules that you configure will be audited on the target servers or snapshot specification.
- 6** For each rule, select one of the following check types:
- **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array.
  - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.
  - **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.
- 7** You can also configure a rule based upon a wildcard search by selecting the Wildcard rule object \*. When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.
- For example, you could enter simply \* which would match everything on the target, P\* would match all objects that begin with a capital P, while \*P would match all elements ending with uppercase character 'P'.
- After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.
- It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.
- 8** To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.

- 9** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 10** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

### Configuring Software Rule

Configuring a software rule allows you to audit all packages and patches that are installed on the target servers and are registered in the Opware model repository.

Note that this is different from configuring the Registered Software rule, which actually shows the currently installed patches and packages on a managed server, regardless of software registration. For more information, see “Configuring Registered Software Rule” on page 185.

To configure software rules, perform the following steps:

- 1** Create the new audit using one of the methods in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Software.
- 4** In the Content pane of the Audit window, expand the top level node in the Available for Audit section and select Installed Packages or Installed Patches.
- 5** Click the right arrow button to move the items into the Selected for Audit section. All items that you select will be used to audit or snapshot the target server.
- 6** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 8** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

### Configuring Storage Rule

The storage rule allows you audit servers for storage devices and SAN devices and connections in your data center, if your core is ASAS-enabled.

To configure the storage rule, perform the following steps:


- 1** Create the new audit using one of the methods in “Creating an Audit” on page 148.
- 2** Select an Audit Source: Server, Snapshot, Snapshot Specification, or No source. (Some audit rules, such as Application Configuration and Windows User’s and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Storage.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Storage rule that you want to create a rule from, such as LUN Utilization, Multipath, Raid Type, or Unreplicated LUN Count.
- 5** Click the right arrow button to move the rule object into the Selected for Audit section. All storage rules that you configure will be audited on the target servers or snapshot specification.
- 6** For each rule, select one of the following check property:
  - An operator, such as equal to (=), less than (<), less than or equals to (<=), and so on.
  - A value, depending upon the rule type, such as a number.
- 7** To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 8** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 9** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

### Configuring Windows .NET Framework Configurations Rule

The Windows .NET Framework Configuration rule allows you to use time information about Assembly Cache and Configured Assembly List, such as assembly name, version, locale, public key token, cache file (GAC or ZAP), processor architecture, custom, and file name in your audits.

To configure the Windows .NET Framework Configuration rule, perform the following steps:

- 1** Create the new audit using one of the methods in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User’s and Groups, must have a source.)

- 3** In the Audit window, from the View pane, select Rules ► Windows .NET Framework Configuration.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Windows .NET Framework Configuration rule that you want to create a rule from.
- 5** Click the right arrow button to move the rule object into the Selected for Audit section. All Windows .NET Framework Configuration rules that you configure will be audited on the target servers or snapshot specification.
- 6** For each rule, select one of the following check types:
  - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array.
  - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.
  - **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object exists on the target server, then the rule is out of compliance.
- 7** You can also configure a rule based upon a wildcard search by selecting the Wildcard rule object \*. When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.

For example, you could enter simply \* which would match everything on the target, P\* would match all objects that begin with a capital P, while \*P would match all elements ending with uppercase character 'P'.

After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.

It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.

- 8** To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 10** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

## Configuring Windows Registry Rule

The Windows Registry rule allows you to select Windows Registry folders and keys to compare in your audit. The audit compares the selected registry folders and keys and determines if these keys and folders exist on the target servers.

There are two categories of Windows Registry rules that you can define in an audit or snapshot specification. The following categories appear in the Available for Audit section of the Audit window:

- **Windows Registry:** These are comparison-based rules, which enable you to select a Windows Registry key or folder from the source of the audit or snapshot specification and compare these with the target servers. The purpose for this kind of rule is to determine if the Windows Registry key or folder exists and its properties. You cannot set a target or remediation value in the rule.

The Windows Registry object allows you to capture registry keys, values, and subkeys. 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. The Opsware Audit and Remediation feature supports the following Windows Registry keys: HKEY\_CLASSES\_ROOT, HKEY\_CURRENT\_CONFIG, HKEY\_LOCAL\_MACHINE, and HKEY\_USERS.

Valid control characters audited and captured for the contents of the key entry (Data) include: #x9, #xA, [#xD, #x20-#xD7FF], [#xE000-#xFFFFD], and [#x10000-#x10FFFF]. Invalid control characters cannot be stored by the SAS Client 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 the audit or snapshot specification results.

You can also choose to compare Access Control Levels (ACLs) for registry rule.

- **Specific Windows Registry Rules:** These are value-based Windows Registry rules prebuilt into the Opsware SAS and they allow you to configure expected (target) and remediation values.

To configure Windows Registry audit rules, perform the following steps:

- 1** Create the new audit using one of the methods for creating an audit listed in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Windows Registry.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Windows Registry folder or key to create a rule for.
- 5** Click the right arrow button to move the Windows Registry folder or key into the Selected for Audit section. All items that you select will be used to audit or snapshot the target server.
- 6** For each registry entry key rule you create, you can set two options to include when the audit checks the target:
  - Also Compare Contents of Sub-Keys: Will evaluate all subkeys belonging to the selected registry key.
  - Also Compare ACLs: Will also compare ACLs of the selected registry key.
- 7** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 8** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 9** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.



## Configuring Windows Services Rule

The windows service rule allows you to select windows services to compare in your audit or snapshot specification. The audit or snapshot specification compares the selected services with services on the target servers to determine if the services exist and if the services are started, stopped or disabled.

There are two categories of windows services rules that you can define in an audit or snapshot specification. The following rules appear in the Available for Audit section of the Audit window:

- **Windows Services:** These comparison-based rules enable you to select a service from the source of the audit or snapshot specification and compare them with the target servers. The purpose of windows services rule is to determine if the service exists and its settings. You cannot set a target or remediation value with this type of rule.
- **Other Windows Services Rules:** These value-based windows services rules prebuilt into the SAS Client allow you to configure expected (target) and remediation values.

To configure windows services rules, perform the following steps:


- 1** Create the new audit using one of the methods for creating an audit listed in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Windows Services.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a windows services to create a rule for.
- 5** Click the right arrow button to move the selected windows services into the Selected for Audit section. All items that you select will be used to audit or snapshot on the target server.
- 6** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 7** Save the audit.
- 8** To run the audit, from the **Actions** menu select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

## Configuring Windows/UNIX Users and Groups Rule

The Windows or Unix Users and Groups rule allows you to access local users and groups information from Windows and Unix servers.

To configure the Users and Groups rule, perform the following steps:

- 1** Create the new audit using one of the methods in "Creating an Audit" on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules ► Windows/Unix Users and Groups.
- 4** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a Users and Groups rule that you want to create a rule from.
- 5** Click the right arrow button to move the rule object into the Selected for Audit section. All Users and Groups rules that you configure will be audited on the target servers or snapshot specification.
- 6** For each rule, select one of the following check types:
  - **Property Values:** A values-based check that checks individual properties of the target object. For this type of check, each object requires that you build an expression that defines properties related to the object using the drop down lists at the bottom of the rule window. You can specify a unique operator which depending upon the type of object can be a String, a Number (integer or float), Boolean (comparing values of 'true' and 'false'), Date (a date compare, not a time of day compare), or an Array. For some property types you can select the values from the 'value selector box'.
  - **Equivalent to source:** A comparison check that performs a one to one comparison between the object on the source vs. the target servers. In this type of check, the values of each property selected from both the source and target servers must match exactly for the object to be compliant.
  - **Non-existence:** Checks for the non-existence of an object, to determine if it does not exist on the target server. If the object does exist on the target server, then the rule is out of compliance. The rule is considered compliant if no objects are found.

- 7** You can also configure a rule based upon a wildcard search by selecting the Wildcard rule object \*. When you select this object, in the rule configuration section at the bottom of the window displays a Name field, into which you can type a name (primary key) that will be searched on the target server.  
  
For example, you could enter simply \* which would match everything on the target, P\* would match all objects that begin with a capital P, while \*P would match all users with name ending with uppercase character 'P'.  
  
After you enter a name or wildcard string, you can configure the rule parameters as you did in step 6.  
  
It is important to notice that when using wildcard, all matching objects are restricted by the rule configuration. This type of audit rule is considered compliant if all found objects match the rule parameters.
- 8** To finish configuring the audit, set the target servers, any rule exceptions, the schedule, and the notification for the audit.
- 9** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 10** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

### Configuring The Opsware Network (TON) Rule

The Opsware Network (TON) rules (known by content developers as “pluggable checks”) are part of the TON Subscription Service and give you access to many different types of customized audit rules based on industry compliance standards.


The kinds of rules you have access to depend on your subscription, but can include such rules as the latest patch supplements for Microsoft Windows, current regulatory compliance policies (for example, FISMA, Sarbanes-Oxley), user-created rules from the TON developer community, daily updated vulnerability content, and so on.



If you are not subscribed to TON, you will not see any TON rules in your audits, audit policies, or snapshots. If you would like more information on how to subscribe to TON, contact your Opsware sales representative.

While each TON rule is slightly different and requires its own configuration values, the basic parameters for each TON rule require that you define the Target Value – the expected value you want to find on the server – and an optional Remediation Value.

To configure TON rules, perform the following steps:

- 1** Create an audit using one of the methods described in “Creating an Audit” on page 148.
- 2** Select an Audit Source: None, Server, Snapshot, or None. (Some audit rules, such as Application Configuration and Windows User's and Groups, must have a source.)
- 3** In the Audit window, from the View pane, select Rules.
- 4** Depending on the kind of TON subscription you have, select one of the TON rule categories, indicated by the pluggable checks icon . For example, if you are subscribed to it, you would select the TON rule category named TON Compliance Content.
- 5** In the content pane of the Audit window, expand the top level node in the Available for Audit section and select a TON rule that you want to create a rule from. For example, if you are subscribed to TON Compliance Content, you could select the following: SOX ► Windows 2003 ► .NET Framework Support set to disabled.
- 6** Click the right arrow button to move the TON rule object into the Selected for Audit section. All TON rules that you select will be audited on the target servers or snapshot specification when you run the audit or snapshot specification.
- 7** For each rule, define or set the following parameters:

#### **Input Value**


Some TON rules require an input value as part of the configuration of the target value. For those rules, you will need to specify a success or failure which you can set to true or false. The Description section of the audit rule explains the recommended values.

#### **Target Value**

Here you can specify the value that you expect to be on the target server or servers of the audit, or the value you want to capture in a snapshot. You can change the following parameters:

- **Operator:** If you want to build an expression from the output of the script, choose an Operator, such as equals ( $=$ ), not equals ( $\neq$ ), less than ( $<$ ), greater than ( $>$ ),

and so on.

- **Reference:** Choose the source of the script output.
  - **Source:** This will use the value from the source server and compare that value to with the value found on the target server or servers.
  - **Value:** Enter your own value. This option uses the value you enter and compares it with the value returned on the target server. You can get the value from the source server if you click the eyedropper  icon. The returned value is displayed in the text box, which you can accept as is or edit to your liking.
  - **Server Attribute:** Select to compare a server attribute located on the source server.
  - **Custom Attribute:** Select to compare a custom attribute found on the target server.

### Remediation Value

Each remediation value setting will be different depending on the type of rule, so choose accordingly.

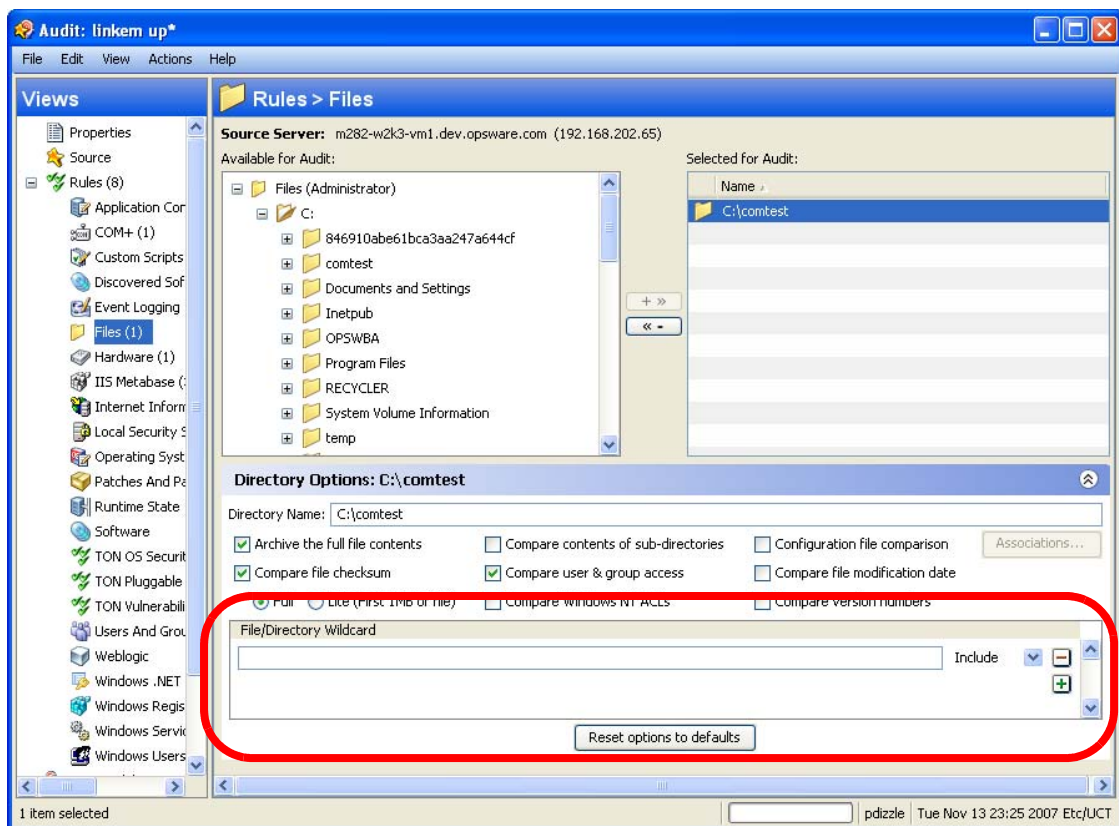
- 8** To finish configuring the audit, set the target servers, the schedule, and the notification for the audit.
- 9** To save the audit, from the **File** menu, select **Save**. You can also save the Audit as a policy. For more information, see “Saving an Audit as Audit Policy” on page 150.
- 10** To run the audit, from the **Actions** menu, select **Run Audit**. For more information about running an audit, see “Creating an Audit Policy” on page 209.

## File Inclusion and Exclusion Rules

When configuring a file rule inside an audit, audit policy, or snapshot specification, you can specify the directories and files that you want included in and excluded from an audit or a snapshot. This section explains what the inclusion and exclusion rules are and how these rules are applied to the relative subset of the absolute path of the file.

Inclusions and exclusion rules inside of an audit's file rule are found at the bottom of the audit or snapshot specification window, as shown in Figure 2-8.

Figure 2-8: File System File/Directory Wildcard Inclusion and Exclusion Rules



When you configure the file rule in an audit or snapshot specification, you can enter inclusion/exclusion rules in the File/Directory Wildcard field. After you enter a rule, you can choose either Include or Exclude from the drop-down list. To add a new inclusion or exclusion rule, click the plus (+) button.

For information on how to create and configure file system rules for an audit or snapshot specification, see “Configuring Files Rule” on page 178.

### **Inclusion and Exclusion Rule Types**

Opsware Audit and Remediation provides the following types of inclusion and exclusion rules configuring a file rule:

- 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, Opsware Audit and Remediation will convert it to a “\” (backslash) to use it as a valid path.
- Environment variable and custom attribute parameterization for filenames and path. For more information, see “Parameterizing Filenames for Opsware/Custom Attributes” on page 204.

Opsware Audit and Remediation 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 might not even be processed because, combined with the exclusion rules (which occur first), they might 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 a path, or one or more characters.

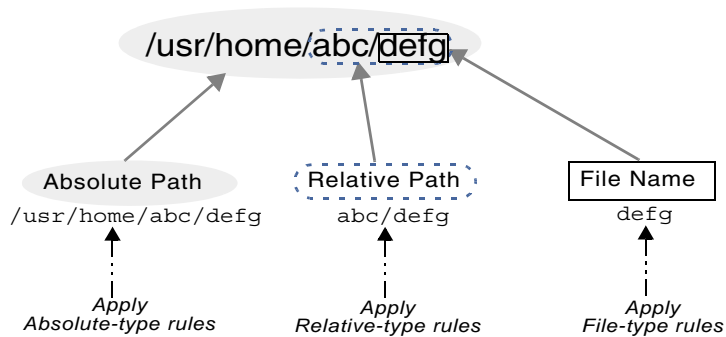
Depending on the type of inclusion and exclusion rule, the rule is applied only to the relevant subset of the absolute path of the file. In Opsware Audit and Remediation, there is one top level for each snapshot or audit. Each file that you compare against the inclusion and exclusion rules has an absolute path. In Figure 2-9, 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 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 2-9 for an illustration of how Opsware Audit and Remediation applies the inclusion and exclusion rules to a relative subset of the path of the file.

Figure 2-9: How Inclusion and Exclusion Rules Apply



To best explain how these rules are applied, the following examples are provided.

A sample file system structure used in “Example: Including all .txt Files in a Snapshot or Audit” on page 200 and “Example: Including last temp.txt file and exclude all else” on page 202 is as follows:

```

/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 .txt Files in a 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 Opsware Audit and Remediation 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 excluded.
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 File a in a Snapshot or Audit**

If you want to include only the file 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 Opsware Audit and Remediation 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 last temp.txt file and exclude all 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 Opsware Audit and Remediation 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.

### **File Rule Overlap**

When you include a parent directory (with options) in a rule and a child directory (with different options) as additional parameters, 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. The following examples explain how audit rules for a parent and child directory overlap.

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 rules:

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 audit rules (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 audit rules (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 audit rules for conf1 override the /cust/app/bin/conf audit rule.

### **Parameterizing Filenames for Opware/Custom Attributes**

When you create a file rule in an audit or snapshot specification, you can also reference environment variables and custom attributes in the file name. In the File/Directory Wildcard area of the rule window, you can edit the file name to add these references.

To add a reference to a Windows environment variable, the syntax is %envVarName% and for Unix, the syntax is \${varName}.

The syntax for specifying custom attributes is @varName@. For example:

```
@/opsware/facility/..@ for getting facility attributes
@/opsware/customer/..@ for getting customer attributes
@/opsware/server/...@ for getting server attributes
@/opsware/groups/...@ for getting server group attributes
```

This allows for auditing relative paths on both source and target servers using a parameterized environment variable or custom attribute in the filename.

### **Examples of Parameterizing Filenames**

For example, on the servers you want to audit you know the relative path to an application, but not necessarily the absolute path for all servers. You can parameterize the path in your audit's File rule so the relative pathname is eliminated and the audit checks the relative path anywhere it exists on the target server.

For example, you want to audit a target servers against a golden source server where '%ProgramFiles%' is '\Program Files' against target servers where %ProgramFiles% is D:\Program Files.

In the File/Directory Wildcard section of the File rule, you can specify the root of the directory rule in the audit to be %ProgramFiles%\Company\MyApp. The audit will remove %ProgramFiles% from the paths of the servers it targets when you run the audit. In other words, C:\Program Files\Company\MyApp\file1.txt on the source server will be compared with D:\Program Files\Company\MyApp\file1.txt on the target servers.

In another example, you may want to audit an application that is installed into two completely different subdirectories on two different servers.

For example, in your audit you choose from a golden source server configuration the installation path of the following:

```
/usr/local/app-version-1232/prog
```

And, your target servers have the application installed anywhere under this path:

```
/usr/local/app
```

In order to audit the target server, you can defines a custom attribute APP\_INSTALL\_LOC with a value of /usr/local/app-version-1232/prog for the golden server and /usr/local/app for the production servers. The File rule in the audit would look something like this:

```
@/customattribute/APP_INSTALL_LOC@/prog
```

This would cause the audit to treat @/opsware/customattribute/APP\_INSTALL\_LOC@ as if it were an environment variable on the target server and do a path replacement.

If you wanted to reference a server attribute, the path would be entered like this:

```
@/server/APP_INSTALL_LOC@/prog
```

## Audit Rule Exceptions

For most audit rules, you can create temporary or permanent rule exceptions on selected target servers (or groups of servers) in the audit. This means you can exclude specific rules on selected targets of the audit when the audit runs.

For example, in an audit that is auditing several servers, you might want to suspend one or more of the rules for a subset of the servers targeted by the audit. You might have a collection of Windows servers that are regularly audited to make sure that the IIS service is disabled, for example, to meet company security standards. Your audit is configured to check each of those servers to make sure IIS is disabled. If IIS is enabled on any of the servers, the audit will fail.

However, for a short period of time you might want to run a business application that requires the IIS service to be enabled in order to run on a few of the servers targeted in the audit. You can create a rule exception for the rule governing the IIS service and associate the exception with the servers that need to run the application. This ensures that the audit can still run and not fail when it encounters the servers that do have the IIS service enabled.

You can set an expiration date for the rule exceptions to make sure that when the rule exception is no longer needed or permitted, the rule will be applied to all servers in the audit. You can also write a reason for the exception and associate a ticket ID with it. Exceptions you create in one audit do not affect rules in any other audits.

## Rules That Cannot Have Exceptions

Most audit rules can have exceptions created for them. However, rule categories that include ALL of a set of rules cannot have exceptions, such as:

- Hardware
- Installed Packages
- Installed Patches

## Considerations When Applying Exceptions to Device Groups

When you set an audit rule exception for a device group, the exception will be applied to all servers in the group. It is possible that one of the servers in the group with the exception also belongs to another device group, which also happens to be the target of an audit that has no exceptions applied to it.

In this situation, the rule exception always applies to the server, even though the server also belongs to a device group with no exceptions. As a rule of thumb, keep in mind any servers in a device group that has a rule exception applied to it will have the audit rule excepted, whether or not the server belongs to another device group that is targeted by an audit and has the same rule applied without an exception.


### Adding a Rule Exception to an Audit

To create an audit rule exception, select any of the rules configured in your audit and using the Add Rule Exception window, associate them with a target server in the audit. When you run the audit, the selected rule and the target servers or snapshots associated with the rule will not be applied.

You can also apply rule exceptions to device groups. You can set the rule exception to run indefinitely, or to expire at some future point in time. You can add a comment to explain why you are creating the exception, and also associate a ticket ID with the exception.

Some audit rules and audit rule collections cannot be excepted. For more information, see “Rules That Cannot Have Exceptions” on page 206.

To add a rule exception to an audit, perform the following steps:

- 1** First, create an audit. For information see “Creating an Audit” on page 148.
- 2** Configure audit rules for the audit. For information on configuring audit rules, see “Audit and Remediation Rules” on page 163.
- 3** From the audit view pane on the left, select the Exception  object.
- 4** Next, from the content pane, click **Add**.



You can also select any rule in the Audit window, right-click, and select **Add Exception**. However, if the audit is referencing a linked audit policy, right-clicking a rule to add an exception will not work.

- 5** In the Add Exception window, from the Select Target Server section, select a server, multiple servers, or device groups to which you want to apply the rule exception.
- 6** Next, from the Select Rule section, select one or more rules you want associated with the servers you selected in the previous step.
- 7** (Optional) In the Reason for Exception section, add an explanation.


- 8** (Optional) In the Ticket ID section, add the ticket ID associated with this exception.
- 9** In the Expires section, either enter a date to indicate when the exception expires, or select a date from the drop down list.
- 10** When you are finished configuring the exception, click **Add**.
- 11** You now see a list of rule exceptions that will be applied when you run the audit.

### Editing or Deleting a Rule Exception


You can edit an exception in one of two ways:

- Double-click the exception to modify the reason for the exception, the ticket ID, and the exception expiration date
- Click the **Add** to edit a rule (overwrite the existing rule)

To edit an exception, perform the following steps:

- 1** Open an audit window.
- 2** From the audit's View pane on the left, select the Exception  icon.
- 3** From Contents pane, double-click an exception.
- 4** In the Edit Exception window, you can edit any of the exceptions and servers or device groups they are assigned to. When you have edited the exception, click **Add**.
- 5** If you want to completely change and the rule, click the **Add** button and then in the Add Exception window, change the rule by selecting target server and one or more rules. When you are finished, click **Add** to change the exception.

To delete an exception, perform the following steps:

- 1** Open an audit window.
- 2** From the audit view pane on the left, select the Exception  object.
- 3** From the Contents pane, select the exception you want to select, and then click **Delete**.

## Audit Policies

Audit and Remediation allows you to create audit policies, which are a collection of rules that define a desired state of a server's configuration. An audit policy can be used inside an audit or snapshot specification, either through linking or importing. An audit policy is



very similar – in fact, nearly identical – to an audit, but differs from an audit in that it does not contain any information about target servers or scheduling or notification. All audit policies are available in the SAS Client Library, and must be saved to folders.

An audit policy is like a reusable template that represents an ideal state of server configuration and defines specific compliance standard for servers in your facility. An audit policy is useful because it allows a policy setter to define server configuration compliance values, which can then be used by others in the context of an audit or snapshot specification.

You can create an audit policy from scratch, or you can save an existing audit as an audit policy, which extracts only the rules defined in an audit so it can be reused in other audits or snapshots. An audit policy can *link* into an audit or snapshot specification so whenever a change is made the audit using the policy will have the latest changes. An audit policy can also be *imported* into an audit or snapshot specification, without keeping the link to the source audit policy. When you import an audit policy into an audit, you can choose to replace any current values in the audit or merge rules from the audit policy with those in the audit or snapshot specification.

If you subscribe to The Opsware Network (TON), some of the latest industry compliance standards are defined as rules inside each new audit policy. For example, subscribing to TON Essential Content gives you access to regularly updated audit policies containing security best practices, such as CIS, NSA, and so on, and the Opsware patch supplement for Microsoft Windows. Subscribing to the TON Subscription Service, you will be able to access the most current regulatory compliance audit policies (FISMA, Sarbanes-Oxley, etc.) and daily vulnerability alerts.

For information on subscribing to TON, contact your Opsware sales representative.

For information on creating rules for an audit policy, see “Audit and Remediation Rules” on page 163.

### **Creating an Audit Policy**

When creating an audit policy, you have the option of creating the rule using either a live server or a snapshot. This allows you to use the rule from a known good server, or a snapshot of a known good server.



All audit policies must be saved to a folder in the SAS Client Library. To save an audit policy to a folder, you must have permissions to write to that folder. For more information about folder permissions, see the *Opsware® SAS Administration Guide*.

---

To create an audit policy, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Audit and Remediation ► Audit Policies, and then select Windows or Unix.
- 2** From the **Actions** menu, select **New**.
- 3** In the Content pane, for the audit policy's Properties, enter a name and description.
- 4** In the Properties, you also need to specify a location in the Library where you want to save the audit policy. Next to the Location field, click **Select**.
- 5** In the Select Folder window, select a location to save the audit policy. You must have permissions to write to the folder where you save the policy. Click **Select** when you have chosen a location.
- 6** Next, from the Views pane on the left side of the Audit Policy window, select Source if you would like use a server to base the audit policy's rules on.
- 7** From the Content pane, select a source for the audit policy, and then click **Select**.
- 8** In the Select a Source window, select either a server or a snapshot, and then click **OK**.
- 9** From the Views pane, select rules to configure. For more information on how to configure specific rules, see "Configuring Specific Audit and Snapshot Rules" on page 167.
- 10** When you are finished configuring the audit, from the **File** menu select **Save**.

### Locating an Audit Policy in the Folder Library

Once you create and save an audit policy to the folder library, you can easily find the audit policy in the library by using the Locate in Folders feature.

To locate an audit policy in folder, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Audit and Remediation ► Audit Policies, and then select Windows or Unix.

- 2 Select an audit, right-click, and select Locate in Folders. The location where the audit policy is saved is now visible.

### Linking and Importing Audit Policies

You can import or save an audit policy into either an audit or snapshot specification in the three following ways:

- Linking an Audit Policy
- Importing an Audit Policy (replace or merge)
- Saving as Audit Policy

#### Linking an Audit Policy

Linking an audit policy into an audit or snapshot specification creates a link that populates an audit's or snapshot specification's rules with those of the audit policy. This is useful if a policy setter wants to define a policy in an audit and have others users link to the audit policy. If the policy setter makes any changes to the audit policy, then the changes will be reflected audit or snapshot specification that is linked to the policy.

When an audit policy is linked into an audit or snapshot specification, the rules cannot be modified in the audit or snapshot specification.

If the audit or snapshot specification you are linking to already has some rules defined, then the linked the an audit policy will overwrite those existing rules.

To link an audit policy in an audit, perform the following steps:

- 1 Open an existing audit from the Library using one of the following methods:
  - From the Navigation pane, select Library ► Audit and Remediation ► Snapshot Specification, and then open the audit.
  - From the Navigation pane, open an existing snapshot specification from Library ► Audit and Remediation ► Snapshot Specification.
- 2 From the **Actions** menu, select **Link to Policy**.
- 3 In the Select a Policy window, select an audit policy to link to the audit or snapshot. When you have selected a policy, click **OK**.
- 4 If you are linking an audit policy into an audit or snapshot specification that already has rules defined, a message window asks if you want overwrite any existing rule definitions. Click **Yes** to import the audit policy.

- 5 To save the audit or snapshot specification, from the **File** menu, select **Save**.

### **Importing an Audit Policy**

Importing an audit policy into an audit or snapshot specification allows you to import (and optionally merge) an audit policy's rules into an audit or a snapshot specification, without keeping a link to the audit policy.

After you import an audit policy, there is no more connection to that audit policy, and any changes made to the source audit policy are not reflected where the audit policy was imported into.

To import an audit policy into an audit, perform the following steps:

- 1 Open an existing audit or snapshot specification from the Library using one of the following methods:
  - From the Navigation pane, select Library ► Audit and Remediation ► Audits, and then open the audit.
  - From the Navigation pane, open an existing snapshot specification from Library ► Audit and Remediation ► Snapshot Specification.
- 2 From the **Actions** menu, select **Link to Policy**.
- 3 If the audit or snapshot specification already has rules defined, choose to either to overwrite the existing rules, or merge the audit policy rules with the existing rules:
  - If you click **Yes**, then the audit policy will overwrite any existing rules in the audit or snapshot specification.
  - If you click **No**, then the audit policy will merge the audit policy rules with any existing rules. If any conflicts are found, then the audit policy rules will overwrite any existing rules.
- 4 To save the audit or snapshot specification, from the **File** menu, select **Save**.

### **Saving as Audit Policy**

You can save an audit or a snapshot specification's rules as an audit policy, which can be then used by others in an audit or snapshot specification.



All audit policies must be saved to a folder in the SAS Client Library. To save an audit policy to a folder, you must have permissions to write to that folder. For more information about folder permissions, see the *Opware® SAS Administration Guide*.

---

To save an audit or snapshot specification as an audit policy, perform the following steps:

- 1** Open an existing audit or snapshot specification from the Library using one of the following methods:
  - From the Navigation pane, select Library ► Audit and Remediation ► Audits, and then open the audit.
  - From the Navigation pane, select Library ► Audit and Remediation ► Snapshot Specification.
- 2** After you have configured the audit's or the snapshot specification's rules, from the **File** menu, select **Save As**.
- 3** In the Save As window, enter a name and description.
- 4** From the Type list, select Audit Policy.
- 5** Click **Select**.
- 6** In the Select Folder window, choose a folder where you want to save the audit policy, and then click **OK**. The audit policy is saved and can be accessed at Library ► Audit and Remediation ► Audit Policies.

## Running an Audit

Running an audit will execute the selected audit on the target server, servers, or snapshot of the audit, and it will evaluate the targets according to the rules defined in the audit. You can run an audit from the following locations in the SAS Client:

- Running an Audit from the Library
- Running an Audit on a Server from All Managed Servers
- Re-running an Audit from Audit Results

### Running an Audit from the Library

The Library contains all available audits that you can run, organized by operating system, either Windows or UNIX. The list of audits in the Library can be sorted by any of the columns (Name, Last Modified Date, and so on). The search tool (upper right of the window) can also be used to search the audit list by entering a name, ID, person who created the audit, and so on.

To run an audit from the Library, perform the following steps:

- 1** From the Navigation pane, select **Library** ► **By Type** ► **Audit and Remediation**.
- 2** Select Audits, and then select either Windows or Unix.
- 3** Select the audit you want to run, right-click, and select **Run Audit**.
- 4** In the Run Audit window, step one shows you the name of the audit, the source server or snapshot being used in the Audit, the total number of rules defined in the audit, and all targets of the audit (servers and snapshot). Click **View Rule Details** to view the rule definitions.  
  
(If you would like to run the audit immediately, click the **Start Job** button at any point in the process.)
- 5** Click **Next**.
- 6** In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 7** Click **Next**.
- 8** In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
- 9** (Optional) You can specify if you want the email to be sent upon success or failure of the audit job.
- 10** (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opsware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 11** Click **Next**.
- 12** In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

### **Running an Audit on a Server from All Managed Servers**

You can run an audit from this location, if the server is being used as a target for an audit.

To run an audit from the All Managed Servers list, perform the following steps:

- 1** From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**.

- 2 Select a server. From the View drop-down list, select Audit and Remediation. The Details pane area will display below the Content pane.
- 3 From the Details pane Show drop-down list, select Audit - Server is Target.
- 4 Select an audit from the list, right-click, and select **Run Audit**.
- 5 In the Run Audit window, step one shows you the name of the audit, the source server or snapshot being used in the Audit, the total number of rules defined in the audit, and all targets of the audit (servers and snapshot). Click **View Rule Details** to view the rule definitions.

(If you would like to run the audit immediately, click the **Start Job** button at any point in the process.)

- 6 Click **Next**.
- 7 In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 8 Click **Next**.
- 9 In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
- 10 (Optional) You can specify if you want the email to be sent upon success or failure of the audit job.
- 11 (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 12 Click **Next**.
- 13 In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

## Re-running an Audit from Audit Results

You can rerun an audit from an audit results if you would like to run the same audit another time.

Note that when you are viewing the results of an Audit or a Snapshot and re-run the audit from those results, the rules in the original audit may have changed after the results have been capture. Thus it is possible that you will be running the updated audit, and not necessarily the exact audit from which produced these results.

To rerun an audit, perform the following steps:

- 1** From the Navigation pane, select **Library ► By Type ► Audit and Remediation**.
- 2** Select Audits, and then select either Windows or Unix.
- 3** Select an audit, and in the Details pane, select an audit result for the audit. (Each time the audit is run, its results are accumulated in the Details pane.)
- 4** Double-click the audit result to open it.
- 5** From the **Actions** menu, select **Re-Run audit**.
- 6** In the Run Audit window, step one shows you the name of the audit, the source server or snapshot being used in the Audit, the total number of rules defined in the audit, and all targets of the audit (servers and snapshot). Click **View Rule Details** to view the rule definitions.  
  
(If you would like to run the audit immediately, click the **Start Job** button at any point in the process.)
- 7** Click **Next**.
- 8** In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 9** Click **Next**.
- 10** In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
- 11** (Optional) You can specify if you want the email to be sent upon success or failure of the audit job.



- 12** (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 13** Click **Next**.
- 14** In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

## Scheduling an Audit

Scheduling an audit requires specifying when you want an audit to be run (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 or cancel existing scheduled audits. When you delete a scheduled audit, all schedules that you have created associated with that audit will also be deleted.



You must have permissions to create, view, edit, and delete audit schedules. To obtain these permissions, contact your Opware administrator. See the *Opware<sup>®</sup> SAS Policy Setter's Guide* for more information.

## Scheduling a Recurring Audit

After you have created, configured, and saved an audit, you can set up a schedule that specifies when you want the audit to run on a recurring basis. After the schedule is set, you can edit the schedule according to your needs.

To schedule a recurring audit, perform the following steps:

- 1** From the Navigation pane, select **Library > By Type > Audit and Remediation**, and then select Audits.
- 2** Select an OS (Windows or UNIX) and then double-click an audit to open it.
- 3** In the Views pane of the Audit window, select Schedule.
- 4** In the Schedule section, choose to run the audit once, daily, weekly, monthly, or on a custom schedule. Parameters include:
  - **None**: No schedule will be set. To run the audit, select the audit, right-click, and select **Run Audit**.

- **Daily:** Choose this option to run the audit on a daily basis.
- **Weekly:** Choose the day or days of the week to run the audit.
- **Monthly:** Choose the months to run the audit run, and the days of the month.
- **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the audit at 1:00 a.m. every weekday:

```
0 1 * * 1-5
```

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 more information about crontab entry formats, consult the Unix man pages.

- 5** In the Time and Duration section, for each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the audit will keep running indefinitely. To choose a date to end the audit schedule, select End and then choose a date. The Time Zone is set according to the time zone set in your user profile.
- 6** (Optional) Deselect the End option, if you want the audit schedule to run indefinitely.
- 7** To save the audit schedule, from the **File** menu, select **Save**. The audit will now run according to the defined schedule.

## Editing an Audit Schedule

You can edit an audit schedule after you have created (or edited) and saved it.

To edit a scheduled audit, perform the following steps:

- 1** From the Navigation pane, select Jobs and Sessions.
- 2** Select Recurring Schedules.
- 3** From the drop-down list at the top of the Content pane, select Audit Servers.
- 4** Select a scheduled audit job, right-click, and select **Open**.
- 5** In the Audit window, select Schedule in the Views pane to view the audit schedule.
- 6** To edit the audit Schedule, modify the following parameters:
  - **None:** No schedule will be set. To run the audit, select the audit, right-click, and select **Run Audit**.
  - **Daily:** Choose this option to run the audit on a daily basis.

- **Weekly:** Choose the day or days of the week to run the audit.
- **Monthly:** Choose the months to run the audit run, and the days of the month.
- **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the audit at 1:00 a.m. every weekday:

```
0 1 * * 1-5
```

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 more information about crontab entry formats, consult the Unix man pages.

- 7** In the Time and Duration section, for each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the audit will keep running indefinitely. To choose a date to end the audit schedule, select End and then choose a date. The Time Zone is set according to the time zone set in your user profile.
- 8** (Optional) Deselect the End option, if you want the audit schedule to run indefinitely.
- 9** To save the audit schedule, from the **File** menu, select **Save**. The audit will now run according to the defined schedule.

### Viewing a Completed Audit Job

To view information on a completed audit job, perform the following steps:

- 1** From the Navigation pane, select Jobs and Sessions.
- 2** Select Job Logs.
- 3** The Content pane displays all jobs run in this Opsware core. To display only audit jobs, from the drop-down list at the top of the Content pane, select Run Audit Task. If you want to see only your scheduled audits, enter your user ID in the User ID field at the top of the Content pane.
- 4** Open an audit job to view the audit results, and then click **View Results**.

## Viewing and Remediating Audit Results

An audit defines the server object configurations that you want to check on a server (according to the audit's rules); audit results are the end product of running an audit and show any differences between the audit rules and the actual server configuration values for each target server or snapshot.

The type of audit result and remediation that you can perform depends on the types of audit rule. You can view and remediate to types of audit results: server comparison or value-based.



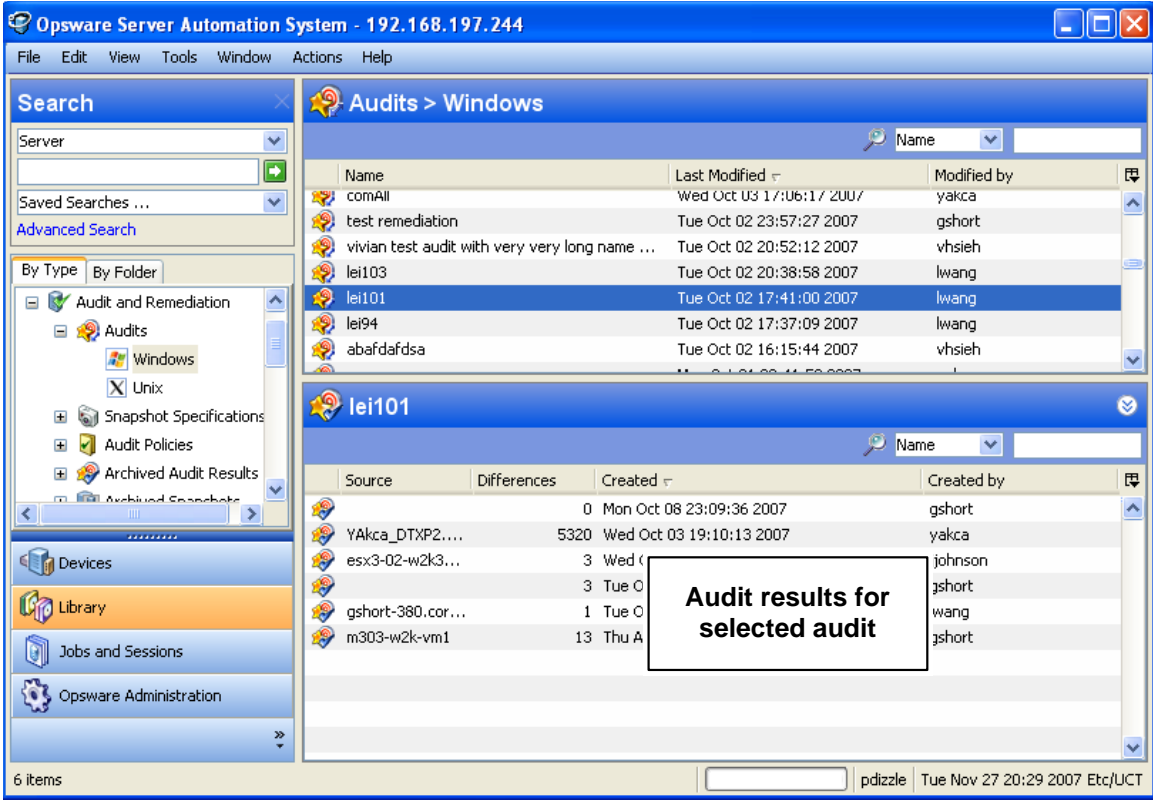
If you have Audit Results with differences from Audits that were created in SAS 5.1, and you have upgraded to SAS 6.x, when you view those Audit Results in the upgraded version of the SAS Client, the Differences column in the Audit Results list will incorrectly display the value of -1 differences. To view the actual number of results, simply open the Audit Results window (double-click it) and you will see all the actual differences in the results.

---

### Accessing Audit Results

In the SAS Client, you can view audit results for any audit, as shown in Figure 2-10.

Figure 2-10: Audit Results



When you select an audit in the library, all results associated with the audit appear in the Details pane below.

### Viewing Comparison-Based Audit Results

Audit results based on a comparison-based audit allow you to view differences between the source server (or snapshot) and target servers or snapshot. If the audit results fails – that is, finds differences between source and target – you can remediate the differences

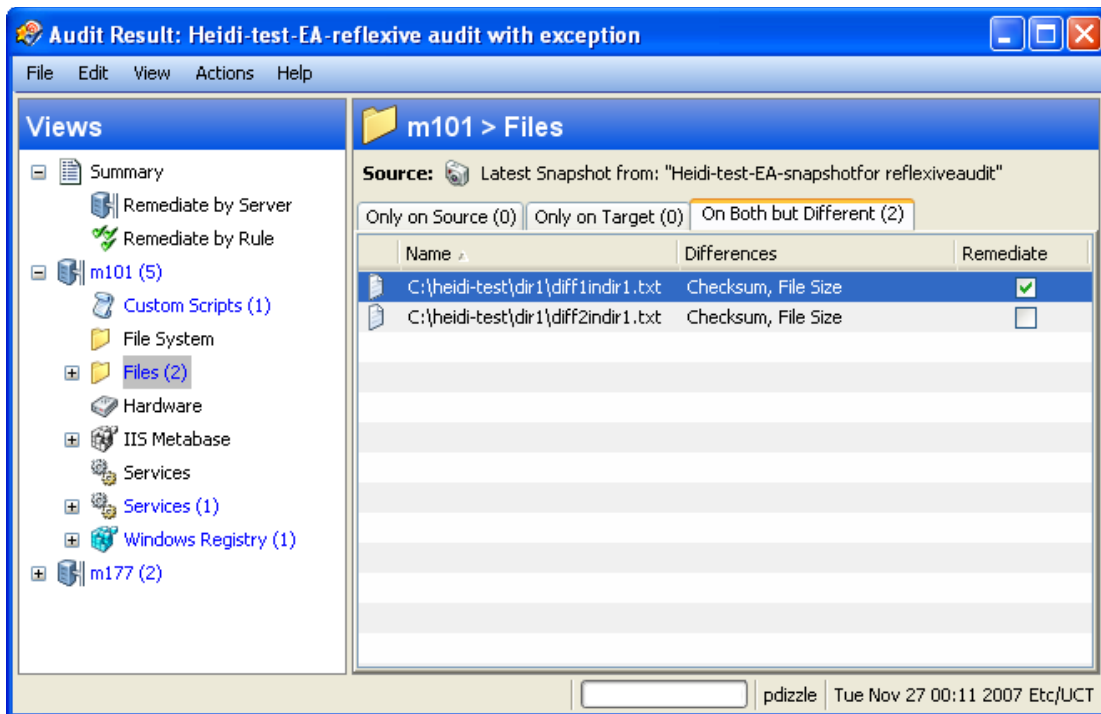
(for most rule types). You can remediate the rule values of the source objects in the audit and overwrite the values on the target (or add values that exist on the source, but do not exist on the target.)

The Audit Results window shows all the objects defined in the audit in the Views pane. It also shows the audit results that failed, the differences found between the audit and the target servers are highlighted in light blue font.

For example, Figure 2-11 shows audit results for a windows file system rule, where the selected file and path exist on both the source (audit rule source server) and the target, but are different, located under the Only Both But Different tab of the Audit Results Window.

From the Audit Results Window, you can select the Files rule, and from the **Actions** menu select **Remediate**.

Figure 2-11: Audit Results For a Comparison-Based Audit Rule



In this example where file difference were found between the source and the target, you can double click the rule to view those differences in a separate window, to make sure you want to perform the remediation. Then, you can select **Remediate** from the **Actions**

menu and remediate the out of compliance rule – or, schedule the audit to run at a later time. When you remediate, the values from the audit (derived from the source) will replace those on the target server.



---

When remediating COM+ objects from snapshot or audit results, the SAS Client does not check the version of the COM+ object, and thus will always remediate the object, whether or not there is any difference between them.

---

### **Viewing Value-Based Audit Results - Audit Rule Remediation**

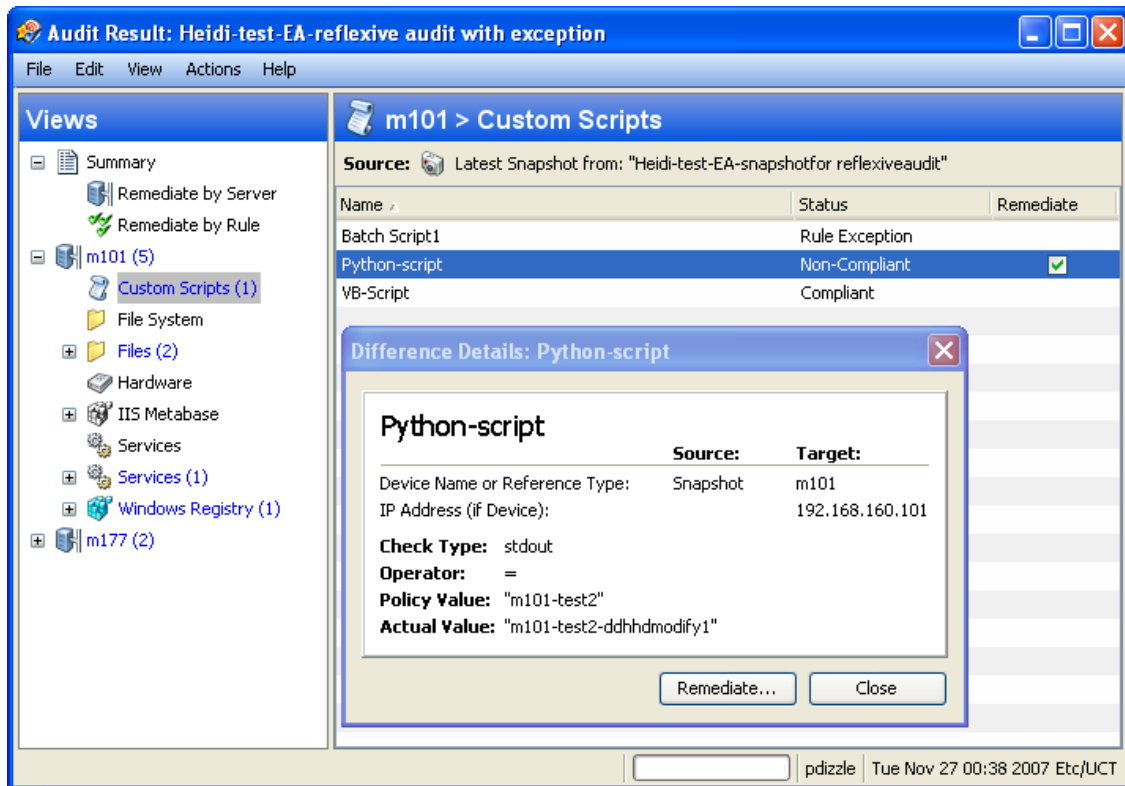
Value-based audit results indicates if the server configuration matches the values defined in the audit rule. You can view the differences between what was defined as the expected value in the rule and the actual value found on the target server. Depending on the rule, you can remediate the difference found on the target server by replacing it with the value specified in the rule.

Some value-based rules are not remediable. For example, Windows/Unix users and groups, the Property value check is not remediable.

Figure 2-12 shows a value-based audit rule in the form of a custom script where the output of the script was different than the results of the same script run on the source server. The Status column for the rule indicates Non-Compliant, which means the output

of the script rule is different between the source and the target. To fix the discrepancy, select the Remediate option and select **Remediate** from the **Actions** menu. Or, double-click the rule and click the **Remediate** button.

Figure 2-12: Audit Results for a Value-Based Audit rule



## Remediation Methods: Rule, Server, or All

When you open audit results, the Summary view of audit results describes general information about the audit, such as the user who created it, when it was created, the source server used in the audit the results are based upon, the number of differences, and so on. You can click **View Rule Details** to see the configured rules of audit the results are based upon (read-only).

From the audit results window, there are three different ways to remediate non-compliant rules in audit results:

- **Remediate All:** Remediate all differences found in the audit results
- **Remediate By Rule:** Remediate specific, individual audit rules



- **Remediate by Server:** Remediate by servers targeted by the audit results

### **Remediate All**

You can select to remediate all the differences found in an audit result for all rules that are remediable. This option remediates all remediable rules on all servers targeted by the audit. Rules that are Compliant are not remediated.

To remediate all differences found in an audit results, perform the following steps:

- 1** From the Navigation pane select **Library > By Type > Audit and Remediation > Audits**.
- 2** Select an audit. From the Details pane below the audit list, you see all audit results associated with the audit.
- 3** Select an audit result, right-click, and select **Open**.
- 4** In the Audit Result window, from the Actions menu select **Remediate All**.
- 5** In the Remediate Audit window, step one shows you the name of the audit, the target of the Audit, the total number of rules defined in the audit.
- 6** Click **Next**.
- 7** In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 8** Click **Next**.
- 9** In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
- 10** (Optional) You can specify if you want the email to be sent upon success or failure of the audit job.
- 11** (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 12** Click **Next**.
- 13** In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

## Remediate By Rule

You can remediate specific differences found in rules in audit results. You can choose to select individual rules that are out of compliance and then run the audit. You can select to remediate by rule for all servers targeted by the audit, or choose only selected servers to have the rule remediated.

To remediate specific differences found in an audit results, perform the following steps:

- 1** From the Navigation pane select **Library > By Type > Audit and Remediation > Audits**.
- 2** Select an audit.
- 3** From the Details pane below the audit list, you see all audit results associated with the audit.
- 4** Select an audit result, right-click, and select **Open**.
- 5** In the Audit Result window, expand the Summary list, and then select Remediate By Rule. You see all differences discovered by rule in the audit results.
- 6** For each rule you want to remediate, select the check mark in the list in the All Servers column, which means that when you remediate the audit results, the rule will be remediated on all servers targeted by the audit that the rule is applied to.  
  
If you want to globally select all rules, right-click and select **Select All**. To deselect all rules, right-click and select **Deselect All**.
- 7** When you have selected the rules you want to remediate, from the **Actions** menu, select **Remediate**.
- 8** In the Remediate Audit window, step one shows you the name of the audit, the target of the Audit, the total number of rules defined in the audit.
- 9** Click **Next**.
- 10** In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 11** Click **Next**.
- 12** In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.

- 13** (Optional) You can specify if you want the email to be sent upon success or failure of the audit job.
- 14** (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 15** Click **Next**.
- 16** In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

### **Remediate by Server**

You can remediate specific differences found in rules in audit results by the server that the audit targets. You can select to remediate all rules on all servers, or, for all rules on selected servers.

To remediate specific differences found in an audit results by server, perform the following steps:

- 1** From the Navigation pane select **Library > By Type > Audit and Remediation > Audits**.
- 2** Select an audit.
- 3** From the Details pane below the audit list, you see all audit results associated with the audit.
- 4** Select an audit result, right-click, and select **Open**.
- 5** In the Audit Result window, expand the Summary list, and then select Servers. You see all differences discovered on all servers targeted by the audit.
- 6** For each server you want to remediate, select the check mark in the list in the All Rules column, which means that when you remediate the audit results, all rules will be remediated on the selected servers.  
  
If you want to globally select all servers in the audit results, right-click and select **Select All**. To deselect all servers, right-click and select **Deselect All**.
- 7** When you have selected the servers you want to remediate, from the **Actions** menu, select **Remediate**.
- 8** In the Remediate Audit window, step one shows you the name of the audit, the target of the Audit, the total number of rules defined in the audit.

- 9** Click **Next**.
- 10** In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select Run Task At, and then choose a day and time.
- 11** Click **Next**.
- 12** In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
- 13** (Optional) You can specify if you want the email to be sent upon success or failure of the audit job.
- 14** (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opsware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 15** Click **Next**.
- 16** In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

### **Viewing and Remediating Audit Results Differences**

For some objects in an audit result, you can view those differences between object that exist on both the target and the source and that have differences between them. You can also see what is different about them and remediate them, if necessary.

For some audit rules, you can view general differences, such as a service's status, the release number for a patch, a registry key's value, and so on. For other server objects, such as files, you can view the differences of the file's contents.

### **Viewing and Remediating File Differences**

For some rules, such as file system, you can view differences between files side by side and line by line. You can see lines that were added, deleted, or modified.

To view and remediate contents of two files that differ in an audit, perform the following steps:

- 1** From the Navigation pane select **Library > By Type > Audit and Remediation > Audits**.
- 2** Select an audit.

- 3** From the Details pane below the audit list, you see all audit results associated with the selected audit.
- 4** Select an audit result, right-click, and select **Open**.
- 5** In the Views pane of the Audit Result window, expand one of the target servers and select a result.
- 6** In the Content pane, expand a target server and select one of the results.
- 7** Next, in the Content pane, select the On Both but Different tab.
- 8** Select a file, right-click, and select View Differences.
- 9** In the Comparison window, select an item from the Encoding drop-down list to specify the character encoding of the data displayed.



---

If the file in question exceeds 2MB in file size, Opsware Audit and Remediation cannot display the file differences.

---

- 10** Click the arrows to find the first, next, previous, or last lines that were added, deleted, or modified. 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.
- 11** Click **Close** to close this window.
- 12** To remediate file differences, from inside the Audit Results window, select either the the Only On Source tab or On Both But Different tab, select a file, right-click and select **Remediate**.
- 13** In the Select Server window, select a server you want to copy the file from the source to, and then click **OK**.

### **Viewing and Remediating Object Differences**

For many server objects, such as Users and Groups, IIS Metabase, Windows Registry, and so on, when there are differences between the source object and the target object, you can view differences in object properties side by side. Each server object will show

different windows, depending on the object and if the audit rule set was comparison-based (comparison between source and target) or value-based (comparison between user-defined audit rule and target).

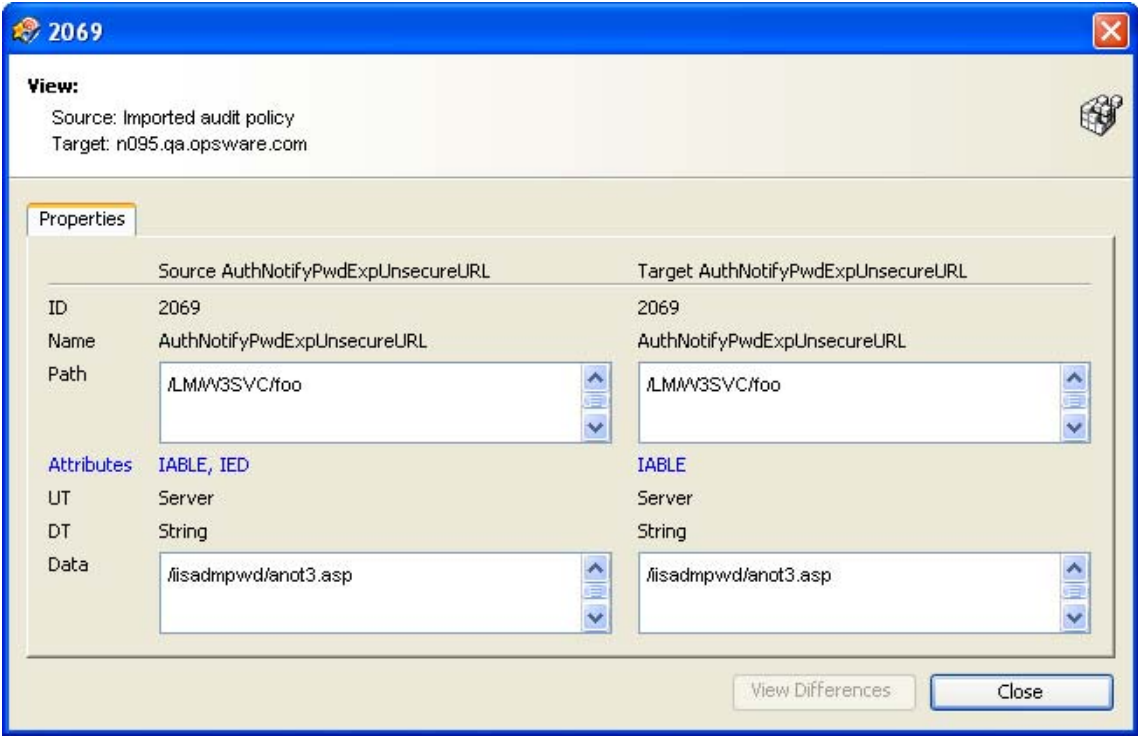
For some value-based audit rules, you can remediate the values on the target server.

To view the contents of two objects that differ, perform the following steps:

- 1** From the Navigation pane select **Library > By Type > Audit and Remediation > Audits**.
- 2** Select an audit.
- 3** From the Details pane below the audit list, you see all audit results associated with the selected audit.
- 4** Select an audit result, right-click, and select **Open**.
- 5** In the Views pane, expand one of the target servers and select a result.
- 6** In the Views pane, select an object.
- 7** In the Content pane, select the On Both but Different tab.
- 8** In the Content pane, select an object, right-click, and select **Open**. You will see a window that shows the differences between the object as defined the audit and the object on the target server.

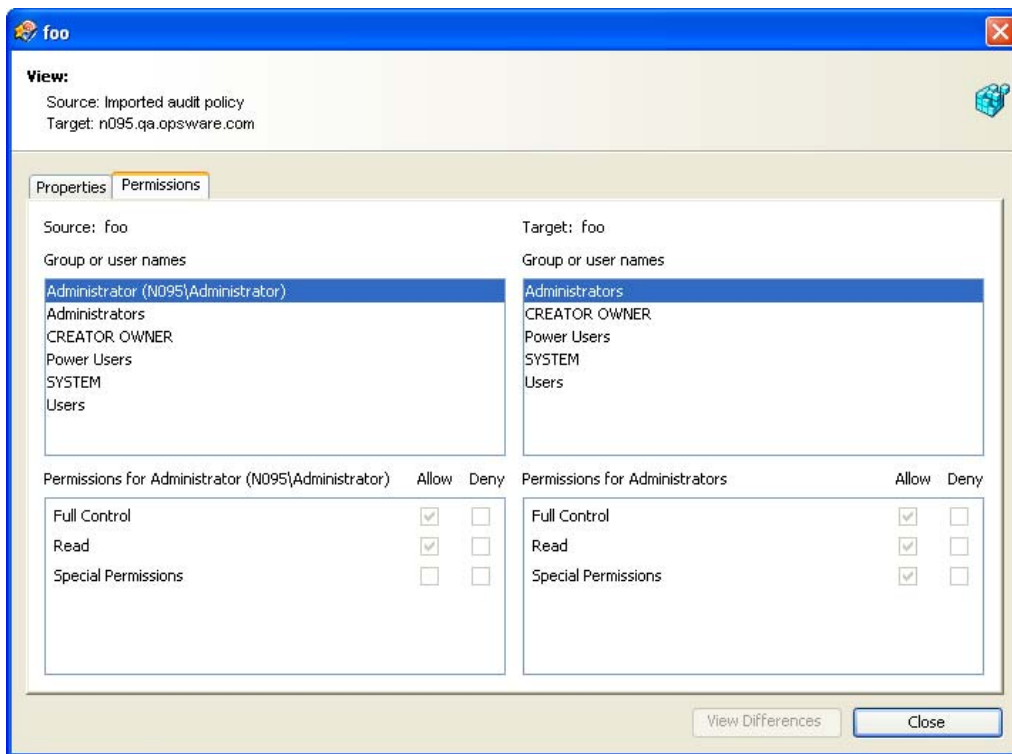
The example in Figure 2-13 displays the audit Result differences for two IIS Metabase objects, showing an attribute of the object that exists on the server but does not exist on the source server, displayed in blue font.

Figure 2-13: Comparison-Based Audit Results Difference: IIS Metabase Objects



For a value-based rule, the difference window will be slightly different and will also include a Remediate option, if remediation is possible. This difference window displays the audit rule, including the policy value and the actual value found on the target server. The example in Figure 2-14 shows the permissions differences for a value-based Windows Registry rule.

Figure 2-14: Rule-Based Audit Results Difference: Windows Registry Permissions Differences




- 9** To remediate the differences, select the Remediate check mark next to each rule.
- 10** From the **Actions** menu select **Remediate**.
- 11** In the Remediate window, follow the steps to run or schedule the remediation. For more information on remediating audit results, see "Viewing and Remediating Audit Results Differences" on page 228.

### Viewing Audit Results with Exceptions

If an audit contains rule exceptions, then the excepted rules are not checked on the target servers when the audit is run. However, your audit results will show which of the rules in the audits are exceptions, including details about the rule exceptions.



The manner in which rule exceptions are displayed in audit results depends on the type of rule that has been excepted:

- Custom script and custom or pluggable check rule exceptions (such as those created by developers or provided by a TON Content Subscription) appear in the Contents pane of the Audit Results window. You can double-click the rule exception for details on the exception.
- All other rule exceptions, such as file system, registry settings, services, IIS Metabase, and COM+ rules, the Audit Results window will display an Exceptions icon  in the Views pane, which you can select and see the details of the exception in the Contents pane.

### Searching for Audits

You can use the SAS Client Search tool to find audits in your facility. You can search for audits by name, by the operating system, and many other criteria.

To search for audits, perform the following steps:

- 1** From inside the SAS Client, ensure that the search pane is activated by selecting View ► Search pane.
- 2** From the top drop-down list, select Audit.
- 3** Click the green arrow button or ENTER to execute the search.
- 4** The results appear in the Content pane.

If you want to extend your search criteria, add new criteria in the search parameters section at the top of the Content pane. You can also save the search by clicking **Save**, or export the Search results to .html or .csv.

### Deleting Audits

To conserve disk space, you can delete audits that you no longer need. You can choose to archive all audit results generated from the audit, if you would like to keep a record of the results.

To delete an audit, perform the following steps:

- 1** From the Navigation pane, select **Library ► By Type ► Audit and Remediation ► Audits**.

- 2** Choosing either Windows or Unix, select one or more audits and then select **Actions** ► **Delete**.
- 3** In the Confirmation Dialog, click **Yes** to delete this audit, or click **No** if you do not want to delete it. You can also select the Archive Audits option, which will archive all audit results generated from the audit. If you do not select the Archive option, all audit results from the selected audit will be deleted.



---

When you delete an audit, all schedules associated with it will be also deleted. See “Scheduling an Audit” on page 217 in this chapter for more information.

---

### Deleting Audit Results

As a best practice, you should delete audit results that you no longer need. If you would like to save audit results, you can choose



---

You must have read permissions for the snapshot to be able to delete it. To obtain these permissions, contact your Opware administrator. See the Opware® SAS Configuration Guide for more information.

---

To delete a snapshot, perform the following steps:

- 1** 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.
- 3** If you want to archive the snapshot instead of delete it, select the snapshot, right-click, and select **Archive**.



---

When you delete a snapshot, you do not delete the snapshot specification that was used to create it. See “Deleting a Snapshot Specification” on page 254 in this chapter for more information.

---

## Archiving Audit Results

Some audits yield numerous results, especially those audits scheduled to run on a recurring basis. You can archive all audits results to keep a record of all audit results run from an audit. When you archive an audit result, it removes its connection to the original audit, but the results and targets of the audit are kept intact.

To archive audit results, perform the following steps:

- 1** From the Navigation pane select **Library ► By Type ► Audit and Remediation ► Audits**.
- 2** Select an audit.
- 3** From the Details pane below the audit list, you see all audit results associated with the selected audit.
- 4** To archive an audit result, select it, right-click, and select **Archive**.
- 5** You are asked to confirm if you want to archive the audit result, since doing so will remove the link between the result and the audit. Click **Yes** to archive the audit result.
- 6** To view all archived audit results, From the Navigation pane select **Library ► By Type ► Audit and Remediation ► Archived Audit Results**.

## Snapshots

A snapshot captures the configuration of a managed server at a particular point in time, and provides a means of capturing the current state of a known working (or, not working) server. A snapshot is useful for capturing a server configuration that you know represents a desired state of configuration. You can also compare the snapshot with other servers in your facility by using the snapshot in an audit.

A snapshot is also a useful 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.

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, application configurations, COM+ objects, hardware information, installed patches, and more. You can even create custom scripts that gather data from the target managed servers.

## **Snapshot Specification and Snapshot**

Snapshots are configured in similar way as you configure an audit. First you create a *snapshot specification*, which is like a template that defines exactly what you want to capture of a server's configuration. Then, you configure the snapshot specification's rules, and then run it. The results are a snapshot – a picture of a server's configuration. The main difference between a snapshot and an audit is that a snapshot takes a picture of a server's configuration, whereas an audit compares a server configuration with the rule values that you define.

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.

## **Snapshot Used in an Audit**

You can use a snapshot in an audit to compare managed servers, groups of servers, and snapshots. By using a snapshot in an audit, you can compare a problematic server (target of the audit) with a known working server (snapshot as source for the audit). To further extend the audit definition, you can also define rules for server objects.

When a snapshot is used as the source for an audit, all server configuration values captured in the snapshot results are available to use as rules for the audit. For more information about using a snapshot in an audit, see "Configuring an Audit" on page 152.

## **Snapshot Specification Used in an Audit**

You can use a snapshot specification as the source of an audit if you want to keep track of a server's configuration over time and monitor any changes that occur. For example, you might want to keep track of a specific application to make sure that its configuration remains correct over a period of time. If this application runs on several servers, you can create a snapshot specification that defines a desired state of server configuration, and then run the snapshot.

Next, you can create an audit and use the original snapshot specification as the source for your audit. Each server that was targeted by the snapshot are now also included as targets of the audit. Next, when you run the audit (either on-demand or on a scheduled basis), each server's current configuration will be compared with the state originally captured when you took the initial snapshot. Any changes are displayed in the audit results window.

For more information, "Configuring an Audit" on page 152.

## Audit Policies and Snapshot Specification

An audit policy is collection of rules that defines a desired state of a server's configuration. An audit policy can be used inside a snapshot specification, either through linking or importing. An audit policy is useful because it allows a policy setter to define server configuration compliance values. These can then be used by others in the context of a snapshot specification.

An audit policy can be linked to an audit or snapshot specification, so whenever a change is made to the policy, the audit or snapshot specification using the policy will also reflect the latest changes. Or, an audit policy can be imported into a snapshot specification, without keeping the link to the source audit policy. When you import an audit policy into a snapshot specification, you can choose to replace any current values in the audit or merge values from the audit policy with those in the snapshot specification.

For more information on importing or linking an audit policy to a snapshot specification, see "Linking and Importing Audit Policies" on page 211.

## Snapshot Specification Elements

An snapshot specification consists of the following elements:

- **Properties:** The name and description of the snapshot specification. If you want to create an inventory of some snapshot specification rules, you can select the Perform Inventory and the snapshot result will collect all information about the specific rules from the target servers. This option is applies to the following rules: Discovered Software, Internet Information Server, Local Security Settings, Registered Software, Runtime State, Windows and Unix Users and Groups.
- **Targets:** The servers that you want to take a snapshot of – that is, capture the specific server configuration as defined in the snapshot specification's rules. You can choose as many servers and groups of servers as you want.
- **Source:** The source of a snapshot specification. If you choose a server then you can select server objects from that server as the basis of your snapshot. The source of a snapshot specification can be a server, or no source at all. (Some rules require a source server. Other rules can be defined by your own custom values without a source.)

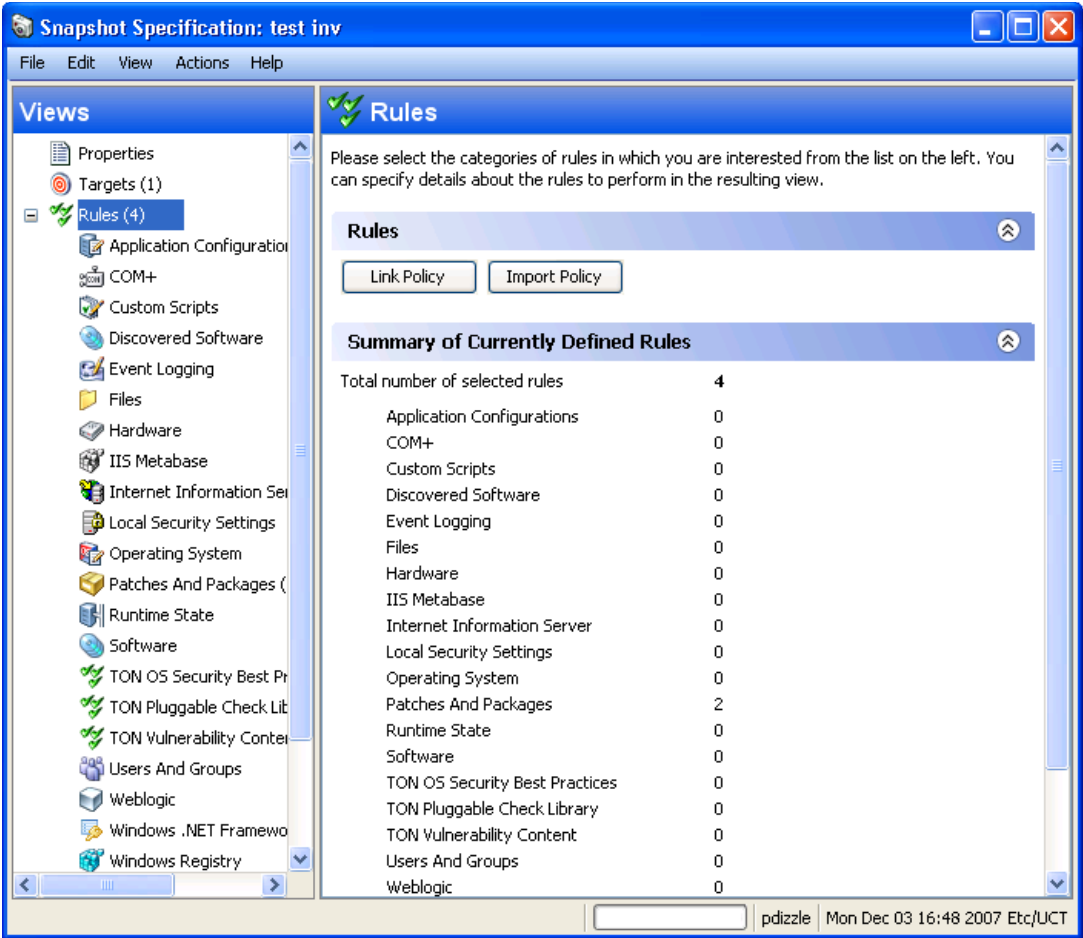
Note that the value of a source parameter is not used when taking a snapshot. It only has meaning when defining a snapshot specification.

- **Rules:** A check on a particular server object with a desired value and an optional remediation value. For example, you might check if a server contains a specific Windows Service, and if found, determine if the service is turned off. For a description of server objects that you can define rules for in a snapshot specification, see “Audit and Remediation Rules” on page 163.
- **Schedule:** The time the snapshot will run. You can run the snapshot specification as a job on a onetime basis, or on a recurring schedule.
- **Notifications:** The email notification send after the snapshot has run. You can base the notification on success, failure, or simply the completion of the snapshot specification job.

When you set up a snapshot specification, you select the objects to check for on the target server. You can also apply rules to these objects that define their desired configuration state. For some rules, you can define remediation values, in the event that the resulting snapshot is used as the source for an audit.

Figure 2-15 shows a snapshot specification that has three rules that will capture configuration information about the target server for event logging, operating system, and windows services.

Figure 2-15: Snapshot Specification Elements



## The Snapshot Process

Taking a snapshot of a server configuration requires the two following basic steps:


- Creating a snapshot specification, which is a template that defines the configuration parameters captured on a target server.
- Running the snapshot specification job that results in a snapshot.

Figure 2-16 illustrates an example of the snapshot process.


Figure 2-16: Snapshot Process

### SNAPSHOT PROCESS - Windows Server Snapshot

#### Part A: Create Snapshot Specification to define “Golden” Server Configuration



#### Part B: Run Snapshot Specification Job and View Results in the Snapshot





## Creating a Snapshot Specification

You can create a snapshot specification from two different locations inside the SAS Client, depending on your purpose. You can create a snapshot specification from the following locations inside the SAS Client:

- Creating a Snapshot Specification from a Server
- Creating a Snapshot Specification from the Library



You must have a set of permissions to create and modify snapshot specifications. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Policy Setter's Guide* for more information.

### Creating a Snapshot Specification from a Server

When you create a new snapshot specification from a managed server, the snapshot specification will use the selected server as its source. You can choose several different server sources for the snapshot specification as you define the rules, or choose no source at all and define your own custom rules. Some rules, however, require a source.



To take a snapshot of a managed server, the server must be reachable and you must have access to the server.

To create a snapshot specification from a server, perform the following steps:

- 1** From the Navigation pane, select **Devices > Servers > All Managed Servers**.
- 2** Select a server, then select **Actions > Create Snapshot Specification**.

### Creating a Snapshot Specification from the Library

If you want to create a new snapshot specification and set all your own rules, create the audit from the SAS Client Library by performing the following steps:

- 1** From the Navigation pane, select **Library > By Type > Audit and Remediation**.
- 2** In the Navigation pane, select snapshot specifications, then Windows or Unix.

## Configuring a Snapshot Specification

To configure a snapshot specification, performing the following tasks:

- Name and describe the snapshot specification, and decide if you want to perform an inventory.
- Choose target servers you want to take a snapshot of. You can choose to snapshot multiple servers or groups of servers.
- Configure your own custom rules, or choose settings from a source server to serve as the basis for the snapshot specification rules.
- Schedule the snapshot specification job to run once or on a recurring schedule.
- Set up email notifications to notify users when the snapshot specification job finishes successfully, if the job fails, or on both conditions.
- Save the snapshot specification.



---

If you take a snapshot of COM+ objects from a 32 bit Windows server, and you attempt to remediate the results using copy to onto a Windows 64 bit server, it may not work

---

## Configuring a Snapshot Specification

To configure a snapshot specification, perform the following steps:

To create a snapshot specification from a server, perform the following steps:

- 1** From the Navigation pane, select **Library** ► **By Type** ► **Audit and Remediation**.
- 2** In the Navigation pane, select Snapshot Specifications, then either Windows or Unix.
- 3** From the **Actions** menu, select **New**.
- 4** In the Snapshot Specification window, enter the following information:
  - **Properties:** Enter a name and description for the snapshot specification. Also, for certain snapshot specification rules (Discovered Software, Internet Information Server, Local Security Settings, Packages and Patches, Runtime State, Windows and Unix Users and Groups), you can select the Perform Inventory option, which will capture all resources associated with the rule.
  - **Source:** Select a source for the snapshot specification. By default, the source server for the snapshot specification will be the managed server that you chose as the source for the snapshot specification. Browse the source server for values to

populate the snapshot specification's rules. You can also choose a different source server as the basis of the snapshot specification for each rule category, or no source at all. If you choose no source, you must define your own rules, or choose to link to an audit policy in the rules section.

- **Rules:** Choose a rule category from the list to begin configuring your snapshot specification's rules. Since each rule is unique and requires its own instructions, to configure specific rules, see "Audit and Remediation Rules" on page 163.

If you want to use an audit policy to define the rules of your snapshot specification, click either **Link Policy** or **Import Policy**. When you link an audit policy, the snapshot specification maintains a direct connection with the audit policy, so if any changes are made to the policy, the snapshot specification will update it with the new changes. If you import an audit policy, the snapshot specification will use all the rules defined in the policy but will not maintain a link to the audit policy. For information on how to import or link to a snapshot specification, see "Linking and Importing Audit Policies" on page 211.

- **Targets:** Choose the Targets of the snapshot specification. These are servers or groups of servers that you want the configured snapshot specification rules to capture. To add a server or group of servers, click **Add**. To choose a source server to use to create the snapshot specification rules, click **Select**.
- **Schedule:** Choose to run the snapshot specification immediately, or on a recurring schedule. Choose whether you want to run it once, daily, weekly, monthly, or on a custom schedule. Parameters include:
  - **None:** No schedule will be set. To run the snapshot specification, select the snapshot specification, right-click, and select **Run snapshot specification**.
  - **Daily:** Choose this option to run the snapshot specification on a daily basis.
  - **Weekly:** Choose a day of the week to run the snapshot specification.
  - **Monthly:** Choose the months to run the snapshot specification.
  - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

```
0 0 * * 1-5
```

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 more information about crontab entry formats, consult the Unix man pages.

- **Time and Duration:** For each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the snapshot specification will keep running indefinitely. To choose an end date to end the snapshot specification schedule, select End, and from the calendar selector, choose a date. The Time Zone is set according to the time zone set in your user profile.
- **Notifications:** Enter the email addresses (separated by a comma or a space) of those you want to receive an email when the snapshot specification Job finishes running. You can choose to send the email notification on both success and the failure of the snapshot specification job (not the success of the audit rules). To add an email address, click **Add Notification Rule**.

- 5 When you have finished configuring the snapshot specification, from the **File** menu, select **Save**.



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 one gigabyte (GB). If you specify that you want to collect the full contents of files in the selection criteria, the data collected might exceed the maximum size that can be successfully recorded in a snapshot.

---

### Configuring Snapshot Specification Rules

For information on how to configure specific snapshot specification rules, see “Audit and Remediation Rules” on page 163.

## Saving a Snapshot Specification as an Audit Policy

You can save selection criteria used in a snapshot and save it as an audit policy. This can be useful if you would like to use the rules configured in a snapshot specification for other snapshot specifications or audits.



In order to save a snapshot specification or an audit as a policy, the policy must be saved to a SAS Client Library folder, and your user must have write permissions to the folder you want to save to. For more information on permissions, see the *Opware® SAS Administration Guide*.

To save your snapshot specification as an audit policy, perform the following steps:



- 1** Launch the SAS Client. From the Navigation pane, select **Library > By Type > Audit and Remediation**.
- 2** Select Snapshots Specification, and then double click a snapshot specification you want to save as an audit policy.
- 3** In the Snapshot Specification window, select **File > Save As**.
- 4** In the Save As window, enter a name and (Optional) description.
- 5** From the Type drop-down list, select Audit Policy.
- 6** Click **Save**. The selected snapshot specification has been saved as an audit policy. To view the audit policy, from the Navigation pane, select **Library > By Type > Audit and Remediation > Audit Policies**. For more information about using audit polices, see “Audit Policies” on page 208.

## Running a Snapshot Specification

When you run a snapshot specification, it captures from the target servers all configuration parameters configured in the rules. After you run a snapshot specification, the results of the snapshot job become a snapshot and can be viewed inside the snapshot.

To run a snapshot specification, perform the following steps:

- 1** From the Navigation pane, select **Library > By Type > Audit and Remediation**.
- 2** In the Navigation pane, select Snapshot Specifications then either Windows or Unix.

- 3** Select a snapshot specification, right-click, and select **Run**.
- 4** In the Run Snapshot Specification window, step one shows you the name of the snapshot, the total number of rules defined, and all targets). Click **View Rules Details** to view the rule definitions.
- 5** Click **Next**.
- 6** In the Scheduling page, choose if you want the audit to run immediately, or some later time and date. To run the audit at a later time, select the second option and choose a day and time.
- 7** Click **Next**.
- 8** In the Notifications page, by default your user will have a notification email sent when the Audit finishes, whether or not the audit job is successful. To add an email notifier, click **Add Notifier** and enter an email address.
- 9** (Optional) You can specify if you want the email to be sent on success of the audit job (  ) or failure of the audit job (  ).
- 10** (Optional) You can specify a Ticket Tracking ID in the Ticket ID field. The ticket ID field is only used when Opware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 11** Click **Next**.
- 12** In the Job Status page, click **Start Job** to run the audit. When the audit has run, click **View Results** to view the results of the audit.

## Scheduling Snapshot Jobs

A snapshot specification job enables you to specify when you want the SAS Client to create a snapshot (either once or on a recurring basis) and who you want to receive email notification about the status of the job. You can also view, edit, and delete existing snapshot specification schedules. When you delete a snapshot specification, all schedules associated with that snapshot specification will be deleted.

This section discusses the following topics:

- Scheduling a Recurring Snapshot Job
- Viewing and Editing a Snapshot Job Schedule

- Viewing and Editing a Snapshot Job Schedule
- Deleting a Snapshot Job Schedule

### Scheduling a Recurring Snapshot Job

After you have created, configured, and saved an snapshot specification, you can schedule snapshot specification a recurring snapshot job. After the schedule is set, you can edit the schedule according to your needs.

To schedule a recurring snapshot specification, perform the following steps:

- 1** From the Navigation pane, select **Library > By Type > Audit and Remediation > Snapshot Specifications**.
- 2** Select either Windows or Unix, and then double-click a snapshot specification to open it.
- 3** From the Snapshot Specification window Views pane, select Schedule.
- 4** In the Schedule section, choose to run the snapshot job immediately or on a recurring schedule. Choose to run it once, daily, weekly, monthly, or on a custom schedule:
  - **None:** No schedule will be set. To run the snapshot job, select the snapshot specification, right-click, and select **Run Audit**.
  - **Daily:** Choose to run the snapshot job on a daily basis.
  - **Weekly:** Choose a day of the week to run the snapshot specification job.
  - **Monthly:** Choose the months to run the snapshot specification job.
  - **Custom:** In the Custom Crontab string field, enter a string the indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:
 

```
0 0 * * 1-5
```

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 more information about crontab entry formats, consult the Unix man pages.
  - In the Time and Duration section, for each type of schedule, specify the hour and minute you want the daily schedule to start. Unless you specify an end time, the snapshot specification job will keep running indefinitely. To choose an end date to end the audit schedule, select End, and then choose an end date. The Time Zone

is set according to the time zone set in your user profile.

- (Optional) Deselect the End option if you want the snapshot specification job to run indefinitely.

- 5** To save the snapshot specification job schedule, from the **File** menu select **Save**. The snapshot specification will now run according to the defined schedule.

## Viewing and Editing a Snapshot Job Schedule

You can edit a snapshot specification schedule after you have created (or edited) and saved it.

To edit a scheduled snapshot specification, perform the following steps:

- 1** From the Navigation pane, select Jobs and Sessions.
- 2** Select Recurring Schedules.
- 3** From the drop-down list at the top of the Contents pane, select Create Snapshot. The list shows all scheduled snapshot specification jobs.
- 4** To view a scheduled snapshot specification, double-click one.
- 5** Select the Schedule object in the Views pane.
- 6** To edit the snapshot specification job schedule, modify the following parameters:
  - **Schedule:** Choose to run the snapshot specification immediately, or on a recurring schedule. Choose to run it once, daily, weekly, monthly, or on a custom schedule. Parameters include:
    - **None:** No schedule will be set. To run the snapshot specification, select the snapshot specification, right-click, and select **Run snapshot specification**.
    - **Daily:** Choose to run the snapshot job on a daily basis.
    - **Weekly:** Choose the day of the week you want the snapshot job to run.
    - **Monthly:** Choose the months to run snapshot specification job.
    - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values. For example, the following crontab string will create the snapshot at midnight every weekday:

```
0 0 * * 1-5
```



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 more information about crontab entry formats, consult the Unix man pages.

- **Time and Duration:** For each type of schedule, specify the hour and minute, the day of the week (and month) you want the daily schedule to start. Unless you specify an end time, the snapshot specification job will keep running indefinitely. To choose a date to end the snapshot specification job schedule, select End and then choose a date. The Time Zone is set according to the time zone set in your user profile.
- (Optional) Deselect the End option if you want the snapshot specification schedule to run indefinitely.

- 7** To save the snapshot specification schedule, from the **File** menu select **Save**. The snapshot job will now run according to the defined schedule.

### Deleting a Snapshot Job Schedule

To delete a snapshot job schedule, perform the following steps:

- 1** From the Navigation pane, select Jobs and Sessions.
- 2** Select Recurring Schedules.
- 3** From the drop-down list at the top of the Contents pane, select Create Snapshot.
- 4** The Content pane displays all snapshot specification jobs that have been run on this Opsware core. To display only snapshot specification jobs, from the drop-down list at the top of the Content pane, select Run Snapshot Task. If you want to see only those snapshot specifications that you have scheduled or run, enter your user ID in the User ID field at the top of the Content pane.
- 5** To delete the schedule, select it, right-click, and select **Delete Schedule**.

## Locating Snapshots

After you have created a snapshot, you can find it in several locations inside the SAS Client.

### *Locating Snapshots In the Library*

- 1** From the Navigation pane, select **Library** ► **By Type** ► **Audit and Remediation** ► **Snapshot Specifications**.
- 2** Select either Windows or Unix.
- 3** From the list, select a snapshot specification. The Details pane at the bottom of the application window displays all snapshots run from the selected snapshot specification.

### *Locating Snapshots in the Device Explorer*

To locate snapshots associated with a specific server, you can view them in server's Device Explorer by performing the following steps,

- 1** From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**.
- 2** Select a server from the list, right-click, and select **Open**.
- 3** In the Device Explorer window, select **Inventory** ► **Snapshot Specification**.
- 4** In the Content pane, select a snapshot specification and all associated snapshots appear in the Details pane at the bottom of the window.
- 5** To view a snapshot, double-click it to open.

## Searching for Snapshots

You can use the SAS Client Search tool to find snapshots in your facility. You can search for snapshots by name, by the operating system, and many other criteria.

To search for snapshots, perform the following steps:

- 1** From inside the SAS Client, ensure that the search pane is activated by selecting **View** ► **Search Pane**.
- 2** From the top drop down list, select Snapshot.
- 3** Click the green arrow button or ENTER to execute the search. The results appear in the Content pane. If you want to extend your search criteria, you can add new criteria in the search parameters section at the top of the Content pane. You can also save the search by clicking **Save**, or export the Search results to .html or .csv.

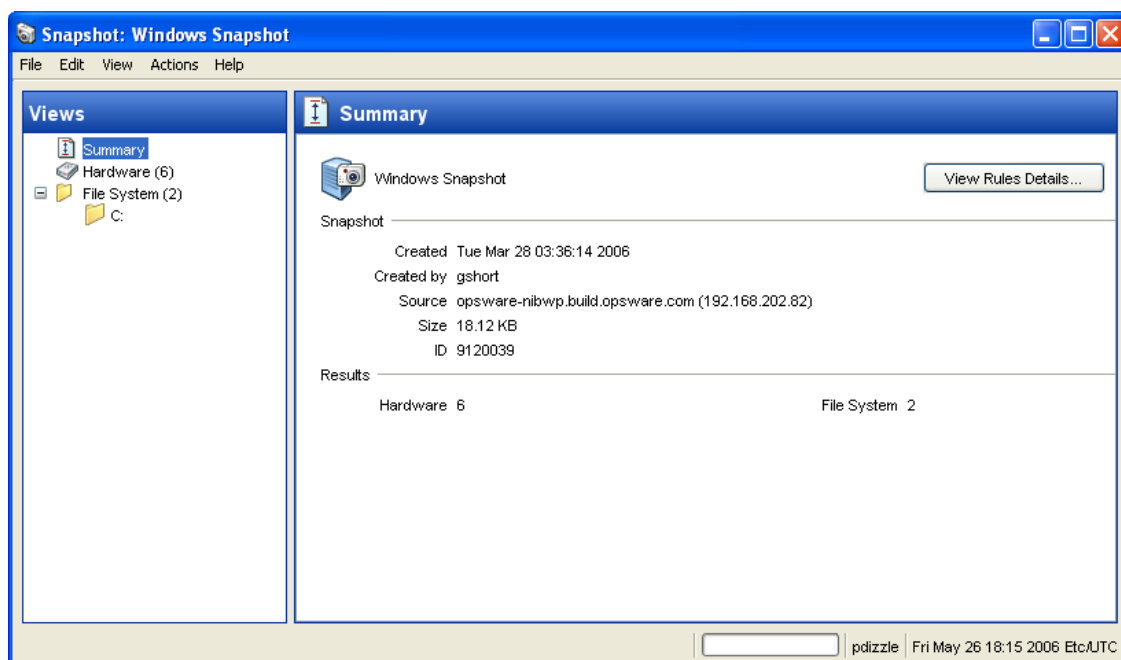
## Viewing Snapshot Contents

You can view the contents of a snapshot and view detailed information about the server configurations that were recorded.

To view the contents of a snapshot, perform the following steps:

- 1 From one of the starting points described in “Locating Snapshots” on page 250, open a snapshot.

Figure 2-17: Sample Snapshot Window of a Windows Server



- 2 In the snapshot window, you can select:
  - **Summary:** Displays 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. You can also click **View Rules Details** to see the snapshot specification which this snapshot is based on.
  - **Installed Hardware:** 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:** Displays information about the installed patches that were recorded in the snapshot, such as the patch type.

- **Installed Packages:** Displays information about the installed packages that were recorded in the snapshot, such as package type, package version, and release number.
- **Event Logging:** Displays security, application, and system log files recorded in the snapshot.
- **File System:** Displays the directories, file properties, attributes, and contents of the files recorded in the snapshot.



---

If a file in the snapshot exceeds 2MB in file size, Opware Audit and Remediation cannot display the file contents.

---

- **Windows Services:** Displays information about the running services recorded in a snapshot, such as the name, description, startup state, startup type, and log on account.
- **Windows Registry:** Displays 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. Opware Audit and Remediation supports the following Windows Registry keys: HKEY\_CLASSES\_ROOT, HKEY\_CURRENT\_CONFIG, HKEY\_LOCAL\_MACHINE, and HKEY\_USERS.
- **COM+:** Displays 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.

Opware SAS provides warning messages that explain how Windows COM folders were processed. The following scenarios apply:

- When you create a snapshot and select a Windows COM folder that does not contain any objects, the snapshot window displays a summary. Opware 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 specification and select a Windows COM+ object that does not exist on a target, Opware SAS displays a warning that the folder is invalid.

- When you create a snapshot and select a Windows COM+ folder that does not contain any objects, Opsware SAS displays a warning that the folder is empty.
- **Metabase:** Displays information about IIS Metabase objects in the snapshot, such as the ID, name, path, attributes, and data of the object.
- **Custom Scripts:** Displays information about the custom script rule recorded in the snapshot.
- **Users and Groups:** Displays information about users and groups on servers, such as user name for last login, whether or not CTRL + ALT + DELETE is enabled, and so on.

**3** Click **Close** to close the object browser.

### Archiving Snapshots

Some snapshot specification yield numerous snapshots, especially those scheduled to run on a recurring basis. You can archive all snapshots to keep a record of all snapshots run for a server or group of servers.

When you archive a snapshot, it detaches the snapshot from the server and removes its connection to the original snapshot specification.

To archive audit results, perform the following steps:

- 1** From the Navigation pane select **Library > By Type > Audit and Remediation > Audits**.
- 2** Select an audit.
- 3** From the Details pane below the audit list, you see all audit results associated with the selected audit.
- 4** To archive an audit result, select it, right-click, and select **Archive**.
- 5** You are asked to confirm if you want to archive the audit result, since doing so will remove the link between the result and the audit. Click **Yes** to archive the audit result.
- 6** To view all archived audit results, From the Navigation pane select **Library > By Type > Audit and Remediation > Archived Snapshots**.

## Deleting a Snapshot Specification

To conserve disk space, you can delete snapshot specifications that you no longer need. You can choose to archive all snapshots generated from the snapshot specification, if you would like to keep a record of the results. Or, you can choose to delete the snapshot specification and all snapshots associated with it.

To delete an snapshot specification, perform the following steps:

- 1** From the Navigation pane, select **Library > By Type > Audit and Remediation > Snapshot Specifications**.
- 2** Choosing either Windows or Unix, select one or more Snapshot Specification and then select **Actions > Delete**.
- 3** In the Confirmation Dialog, click **Yes** to delete this snapshot specification, or click **No** if you do not want to delete it. You can also select the Archive Snapshots option, which will archive all snapshots generated from the snapshot. If you do not select the Archive option, all snapshots generated from the selected snapshot specification will be deleted.



When you delete a snapshot specification, all schedules associated with it will be also deleted. See “Scheduling Snapshot Jobs” on page 246 in this chapter for more 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 Opware administrator. See the Opware® SAS Configuration Guide for more information.

---

To delete a snapshot, perform the following steps:

- 1** 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.

- 3 If you want to archive the snapshot instead of delete it, select the snapshot, right-click, and select **Archive**.



---

When you delete a snapshot, you do not delete the snapshot specification that was used to create it. See “Deleting a Snapshot Specification” on page 254 in this chapter for more information.

---

## Copying Objects from a Snapshot to a Server

After viewing snapshot contents, you can copy certain objects to a target server. Opsware Audit and Remediation allows you to copy directories, files, windows services (state only), IIS Metabase objects, COM+ objects and categories, and Windows Registry keys to a managed server.



---

In order to copy COM+ rule snapshot results from a snapshot to a server, you must have selected the Archive all associated files option when you configured the COM+ rule. Also the COM+ object being copied must not be in use by any application in order for the copy to remediation to work. For more information, see “Configuring COM+ Rule” on page 174.

---

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, Opsware Audit and Remediation 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, Opsware Audit and Remediation 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, Opsware Audit and Remediation 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.
- ACLs are not copied along with COM+ objects or Microsoft IIS objects to the target server.
- When remediating COM+ objects from snapshot results using copy to, the SAS Client does not check the version of the COM+ object, and thus will always copy the object, whether or not there is any difference between them.



---

You must have write permission on the destination server to be able to copy an object to it. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Policy Setter's Guide* for more information.

---

### Copying Objects to a Server from a Snapshot

To copy an object from a snapshot to a managed server, perform the following tasks:

- 1** From one of the starting points described in “Locating Snapshots” on page 250, open a snapshot.
- 2** In the Views pane, select a file system, Windows Services, or Windows Registry object.
- 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 107 in Chapter 18 for more information.

---



# Chapter 3: Server Compliance

## IN THIS CHAPTER

This section discusses the following topics:

- Overview of SAS Server Compliance
- Compliance Terms and Concepts
- Viewing Compliance in the SAS Client
- Adding and Removing Compliance View Columns
- Filtering By Compliance Status
- Refreshing For Latest Compliance Information
- Scanning for Compliance
- Exporting Compliance View Information
- Compliance Remediation
- Audit Compliance
- Software Compliance
- Patch Compliance
- Application Configuration Compliance

## Overview of SAS Server Compliance

The SAS Client Server Compliance feature allows you to view overall compliance levels for all servers and groups of servers in your facility and enables you to remediate servers that are out of compliance. You can view compliance for an individual servers, multiple servers, groups of servers, or for all servers under Opware management.

The Compliance View displays the results of all compliance statuses on servers (or groups of servers) for audits, software policies, patch policies, and application configurations. A server's compliance status is based upon a compliance *policy*, which defines unique server configuration settings or values to ensure that your IT environment

is configured as it should be.

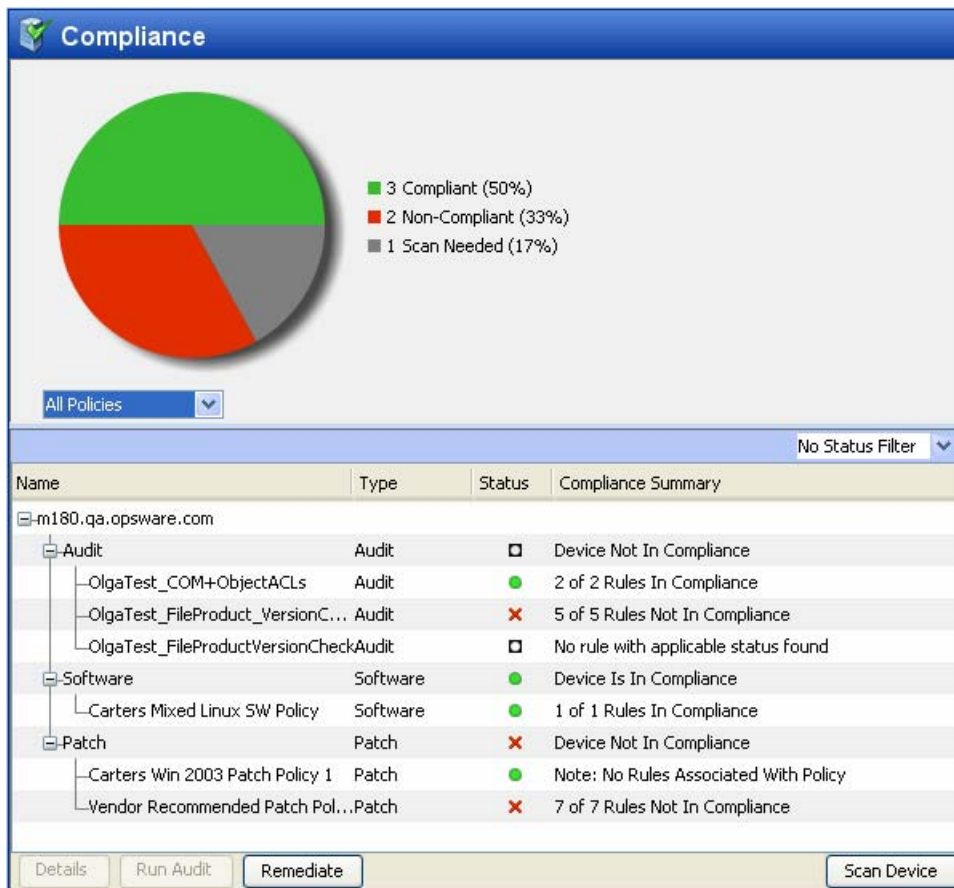
Compliance policies are created and defined by a user often known as a policy setter, though sometimes an ad-hoc policy might be created by a systems administrator. The policy setter creates compliance policies and then attaches them to servers in order to ensure that servers are compliant with the organization's standards and policies.

For example, a policy setter can create a software policy that defines a standard set of patches and packages that should be installed on a server, or the policy setter could define the manner in which certain application files should be configured on a server. A server or group of servers is considered *compliant* if its configuration matches the rules defined by the policy setter in the compliance policy.

The Compliance View in the Device Explorer allows you to determine if the server's actual installed software, packages, patches, and configuration files settings match the configuration defined in the software policy. The Compliance View in the Device Group Explorer allows you to view compliance for groups of servers, showing a compliance

status rollup for all members and (sub-group members) of a group. From the Compliance View, you can discover servers and groups of servers that are out of compliance and remediate any problems.

Figure 3-1: Compliance View in the SAS Client



### Compliance Usage: Proactive and Reactive

You can proactively use the Compliance View by viewing it on a regular basis to assess server compliance levels, and to take the necessary action to fix problems.

For example, you can use the Compliance View to determine the status of an individually scheduled audit that makes sure a Web application's configuration (such as Apache's http.conf file) meets the standards set by your group. You want to ensure that no one has changed the application's configuration. To verify that no unwanted changes have been

made, you should regularly check the Compliance View on this server's Device Explorer to see if this scheduled audit's compliance status has changed to red (non-compliant), and if so, view the audit results and remediate the problem.

In other situations, you can reactively use the Compliance View to answer a specific question or diagnose a specific problem. For example, you can create a scheduled audit that defines security standards for a group of servers in your facility. The audit will ensure that all Windows 2003 servers contain a specific patch. When Microsoft releases a new security patch, you want identify the Windows 2003 servers that contain the new patch and those that do not. You can update the audit to contain the new patch and then browse the Windows 2003 servers in the Device Group's Compliance View. When you rerun the audit, you can discover the servers that need the patch and remediate them by installing the new patch.

## Compliance Terms and Concepts

- **Compliance:** The degree to which a server's actual configuration conforms to the rules defined in a compliance policy.
- **Compliance Category:** The Compliance View displays compliance statuses for four different compliance categories, including Audit, Software, Patch, and App Config (Application Configuration).
- **Compliance Policy:** The user-defined configuration that expresses the desired state for a server or device configuration or setting. For example, a patch policy defines the specific patches that should be installed on a computer. An audit policy might define that a certain Windows service should be disabled at all times. An application configuration policy defines the way in which a configuration file should be configured.
- **Compliance Rule:** The content or setting inside of a policy that defines an ideal configuration for a server, such as a patch or package, a file configuration, software installation order, user and group membership and privileges, and so on.
- **Compliance Statuses:** Indicates the compliance status for a compliance category, reporting the difference between what should be (compliance policy) and what actually is (server configuration). For example, software compliance category in the Compliance View displays a status of Compliant if all configurations defined in the policy match the server configuration. Compliance calculation for groups is slightly different than individual servers. For more information on compliance statuses, see "Compliance Statuses" on page 262.

- **Compliance Scan Results:** The results of a compliance scan. These results report the compliance status, details, and can also include remediate options.
- **Compliance Scan:** The mechanism that checks servers targeted by a compliance policy (audit, software, patch, and application configuration) and returns the results to the SAS Client. A compliance scan could check to see what patches are installed on a computer targeted by a patch policy or software policy and return the results, or, it can check a configuration file's contents and determine if it matches the rules defined in an application configuration. In the Compliance View, you can perform a compliance scan for the Software, Patch, and App Config compliance categories. Audits do not have a scan feature, but running an audit achieves the same results. Running an audit checks the servers targeted by the audit to determine if they are in compliance with an audit's rule definitions.
- **Compliance View:** Displays overall and individual compliance levels for all managed servers or groups of servers in your facility.

For more information on compliance scans, see "Scanning for Compliance" on page 283.

### Server Compliance Categories

The Compliance View for servers and groups of servers displays compliance for the following categories:

- **Audit:** Audit compliance represents an aggregate of all audits that run on a recurring schedule and indicates whether or not the rules defined in a scheduled audit match what is installed and configured on a the target server or servers.

For more information, see "Audit Compliance" on page 289.

For more information about creating and running audits, see "Audit and Remediation" on page 137.

- **Software:** Software compliance is determined by whether or not a software policy definition matches what is installed on a server. A software policy defines patches, packages, and application configurations, scripts as well as a host of other server objects such as services, Windows registry, COM+, IIS Metabase, and so on. A software policy can also contain other software policies.

For more information, see "Software Compliance" on page 293.

For more information on creating software policies, see “Software Management Setup” on page 9.

- **Patch:** Patch compliance is determined by whether or not the patch policy definition matches the patches are installed on a server or group of servers. The Compliance View displays compliance information for Windows patches only.

For more information, see “Patch Compliance” on page 298.

For more information on creating Windows patch policies, see “Patch Management for Windows” on page 323.


- **Application Configuration:** An application configuration’s compliance is determined by whether or not the application configuration definition matches the application configurations on a server or group of servers. An application configuration defines the configuration settings and values for application configuration files.


For more information, see “Application Configuration Compliance” on page 302.

For more information on creating, configuring, and using application configurations, see “Application Configuration Management” on page 495.

For information on how to remediate servers and groups of servers are out of compliance for each of these compliance categories, see “Compliance Remediation” on page 285.

## Compliance Statuses

In general, a server or group of servers can be *Compliant* or *Non-Compliant*. A server is considered Compliant if the rules defined in the policy match the actual configuration on the server that the policy is attached to. When a server is in compliance with the policy attached to it, the Compliance View displays a Compliant icon .

If the server’s actual configuration does not match the rules configured in a policy, then the Compliance View will display a status of Non-Compliant .

For example, you can configure an audit to make sure that a Windows 2003 server has the Windows CIS recommended minimum password length of at least eight characters. When the audit runs and checks the server’s user password and discovers a user password that is only four characters, then the Device Explorer’s Compliance View in the show the server’s audit policy as Non-Compliant.

If the server has more than one audit attached to it, then the Compliance View shows an aggregate, or roll up, compliance status for all audits attached to the server. If at least one of the audits targeting the server is Non-Compliant, then the overall Audit compliance status for the server is Non-Compliant.

If this server belongs to a device group of multiple servers, you can access the Compliance View for the group to see compliance status levels for all audits that run on all servers in the group, as well all servers in any sub-groups. The method used for determining compliance statuses for groups is based upon a default calculation. The group of servers is considered Compliant if at least ninety five percent of the servers that belong to the group have a status of Compliant. If less than ninety five percent of the servers have a status of Compliant, then the status of the group will be Non-Compliant.

The default compliance status threshold for groups of servers can be customized to fit your needs. For more information, see “Changing Device Group Compliance Settings” on page 267.

### Compliance Status Definitions

Table 3-1 lists default compliance statuses for a policies, servers, and device groups.

Table 3-1: Compliance Compliance Status Statuses


ICON	COMPLIANCE STATUS DESCRIPTION
	<p><b>Compliant</b></p> <ul style="list-style-type: none"> <li>• <b>Policy:</b> All rules or items defined in the policy match the actual server configuration.</li> <li>• <b>Servers:</b> Compliance scan ran successfully and the server configuration matches <i>all</i> of the rules defined in <i>all</i> of the policies attached to the server.</li> <li>• <b>Device Groups:</b> Compliance scan ran successfully and at least ninety five percent of all servers in the groups have a status of Compliant.</li> </ul>

Table 3-1: Compliance Compliance Status Statuses







ICON	COMPLIANCE STATUS DESCRIPTION
	<p><b>Partial</b></p> <ul style="list-style-type: none"> <li>• <b>Policy:</b> One or more rules or items defined in the policy does not match the actual server configuration, due to an exception applied to one of the rules. (Windows Patch policies only.)</li> <li>• <b>Servers:</b> Compliance scan ran successfully and the server configuration did not match at least one of the rules defined in any of the policies attached to the server, due to an exception applied to one of the rules. (Windows Patch policies only.)</li> <li>• <b>Device Groups:</b> Compliance scan ran successfully, and more than two percent of all servers in the group have a status of Non-Compliant, and at least ninety percent the servers in the group have a status of Compliant.</li> </ul>
	<p><b>Non-Compliant</b></p> <ul style="list-style-type: none"> <li>• <b>Policy:</b> One or more rules or items defined in the policy does not match the actual server configuration.</li> <li>• <b>Servers:</b> Compliance scan ran and the actual server configuration does not match at least one or more of the rules defined in the policy.</li> <li>• <b>Device Groups:</b> Compliance scan ran and more than five percent of servers (and all servers in all sub-groups) have a status of Non-Compliant.</li> </ul>
	<p><b>Scan Failure</b></p> <p>Compliance scan was unable to run.</p>
	<p><b>Scan Needed</b></p> <p>Results unavailable, perhaps because a compliance scan was never run (for example, on a new installation), or the configuration on the server (or servers in the device group) changed since the last time information was reported to the SAS Client.</p>



Table 3-1: Compliance Compliance Status Statuses

ICON	COMPLIANCE STATUS DESCRIPTION
	<b>Scanning:</b> Compliance scan currently running.
	<b>No Tests Defined</b> No compliance policies of this type are attached to the server or all servers in the device groups (including all servers in any sub-groups).



It is possible that actual server configurations as well as policy information might have changed from the last time you viewed compliance for a server or group in the Compliance View. To get the latest compliance data from the Opsware SAS core, select **Refresh** from the **View** menu. (Or, press F5.) Or, you can run a compliance scan on the server or group to determine compliance status. For more information, see “Scanning for Compliance” on page 283.

## **Compliance Status Thresholds – Policy, Server, and Group**

Compliance status for a policy – an audit, a software policy, a patch policy, an application configuration – is based upon all the rules in the policy. All it takes is one of the rules in a policy to be Non-Compliant (does not match the actual configuration on the server) and the entire policy is also considered Non-Compliant for a server.

Compliance status for a server is based upon all the policies attached to or that target the server. If any one of the compliance categories has a compliance status of Non-Compliant, then the server's overall compliance status is also considered Non-Compliant. Stated another way, all of the policies in all of the compliance categories must be Compliant for the server's overall compliance status to be Compliant.

For information on how compliance status for a server is displayed in the Device Explorer, see “Device Explorer Compliance Summary Pie Chart and Details” on page 270.

## **Device Groups Compliance Status Thresholds**

Whether or not a server is considered Compliant or Non-Compliant is important when viewing device group compliance in the Compliance View, which is based upon a default threshold calculation – and which you can configure and customize).

In the Device Group Compliance View, in order for a compliance category (Audit, Software, Patch, or App Config) to display a status of Compliant, ninety five percent of all servers in a group must have a status of Compliant for each of the categories. If more than five percent of the servers in a group are not in compliance in a given compliance category, the status for the group is displayed as Non-Compliant in each category.

The Partial status default threshold is set by default at ninety percent, which means that ninety percent of the servers are Compliant but at least five percent of the servers are Non-Compliant, the compliance status displays as Partial (partially compliant) for a given compliance category.

Or, another way to state Partial compliance for a group is to say that if greater than two percent of the servers in the group have a status of Non-Compliant, but at least 90% of the servers in the group have a compliance status of Compliant, the category will show a status of Partial.

Group status is calculated recursively based on all policies (in all compliance categories) attached to all servers that belong to the group. This includes servers in all sub-groups that are children to the selected group.

You can change the default thresholds used to calculate compliance status. For example, you could configure that group compliance status be calculated non-recursively, which would exclude all sub-group server members from the compliance calculation.

For more information on how to change default compliance settings for device groups, see “Changing Device Group Compliance Settings” on page 267.

## Changing Device Group Compliance Settings

By default, the SAS Client allows you to configure the manner in which compliance for a device group is determined.

To change the settings for device group compliance, perform the following steps:

**1** From the Navigation panel ► **Opware Administration** ► **Compliance Settings** ► **Device Group Compliance**.

**2** Click **Edit Settings**.

**3** In the Device Group Settings window, you can configure the following settings:

- **Display Device Group Rollup Compliance:** This option allows you to show or hide the icon that indicates compliance status of the parent group shown at the top of each compliance category column. This icon indicates a compliance status rollup for all members of a selected group.

For example, if this option is selected, when you select a group and from View drop-down list select Compliance, the top column heading for each compliance category column (Audit, Software, Patch, App Config) shows an icon that indicates the compliance status for all servers in the selected group. You can mouse-over this column heading to view compliance status counts all devices in this category.

- **Member Calculations:** This option allows you to choose whether or not you want to include servers that belong to sub-groups when calculating overall group compliance level for a compliance category. For example:
  - **Server and group members are considered:** This means that the compliance status for a device group will check compliance for all servers in a group, and all servers in all sub-groups that belong to the selected device group.
  - **Only server members are considered:** This means that the compliance status for the selected device group will only check compliance for servers at the top level of the group, and will exclude any servers that belong to any sub-group members.

- **Thresholds:** Allows you to change the compliance threshold calculation used to determine device group compliance status for all compliance categories. By default, a group will display a status of Non-Compliant if less than five percent of its members are Non-Compliant. A group will display a status of Partial if less than two percent of its members are Non-Compliant. You can set your own default compliance status thresholds here.

## Viewing Compliance in the SAS Client

In the SAS Client, you can view compliance for individual servers, servers and groups together, and for groups of servers:

- Viewing Individual Server Compliance
- Viewing Compliance for Multiple Servers
- Viewing Group Compliance in the Device Group Explorer



When viewing compliance status for groups, it is possible that there are servers in the group that your user does not have permission to see. In addition, your user account might not have permissions to view some of the policies (audit, software, patch) used to calculate the compliance status for a group of servers.

In these cases, even though you cannot see some servers and some policies, you will still be able to see overall compliance status for groups your user has access to view, and you will still be able to see compliance category roll ups, even though some of the policies may be hidden from your view.

---

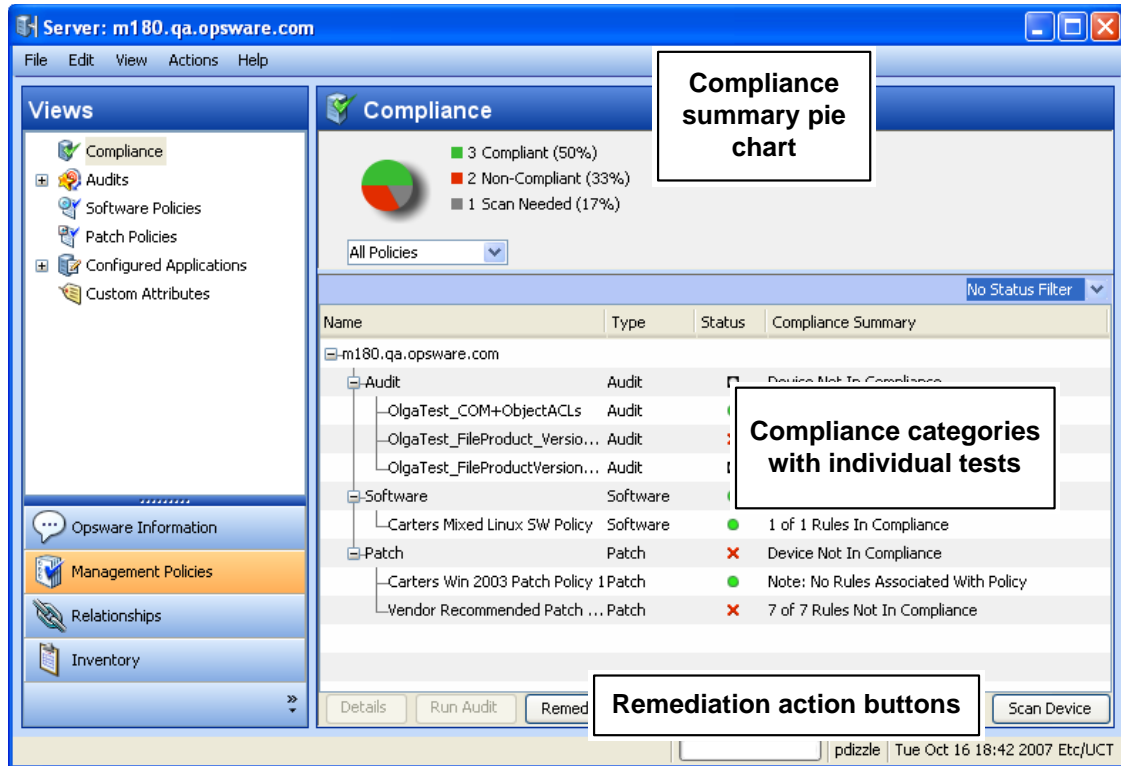
## Viewing Individual Server Compliance

To view compliance information for an individual server, perform the following steps:

- 1** From the Navigation pane, select **Devices** ► **All Managed Servers** (or **Virtual Servers**).
- 2** From the Content pane, select a server, right-click, and select **Open**. (Or, you can double-click the server).
- 3** In the Device Explorer, from the Views pane, select **Management Policies**.

- 4 Select Compliance from the Views pane. On the right side, the Content pane displays a compliance summary pie chart of compliance statuses for each compliance category, as well as detailed status information for individual policies, as shown in Figure 3-1.

Figure 3-1: Compliance View for an Individual Server



- 5 To perform an action on one of the compliance categories, or an individual policy in the categories, make a selection in the Details pane and select either Details, Run Audit (for audits only), Remediate, or Scan Device. For more information on types of remediation you can perform for each policy category, see “Compliance Remediation” on page 285.



The ability to both view policies and perform remediation operations on them is determined by your user’s permissions. If you are not able to view a policy or perform an action on one, consult your Opware administrator.

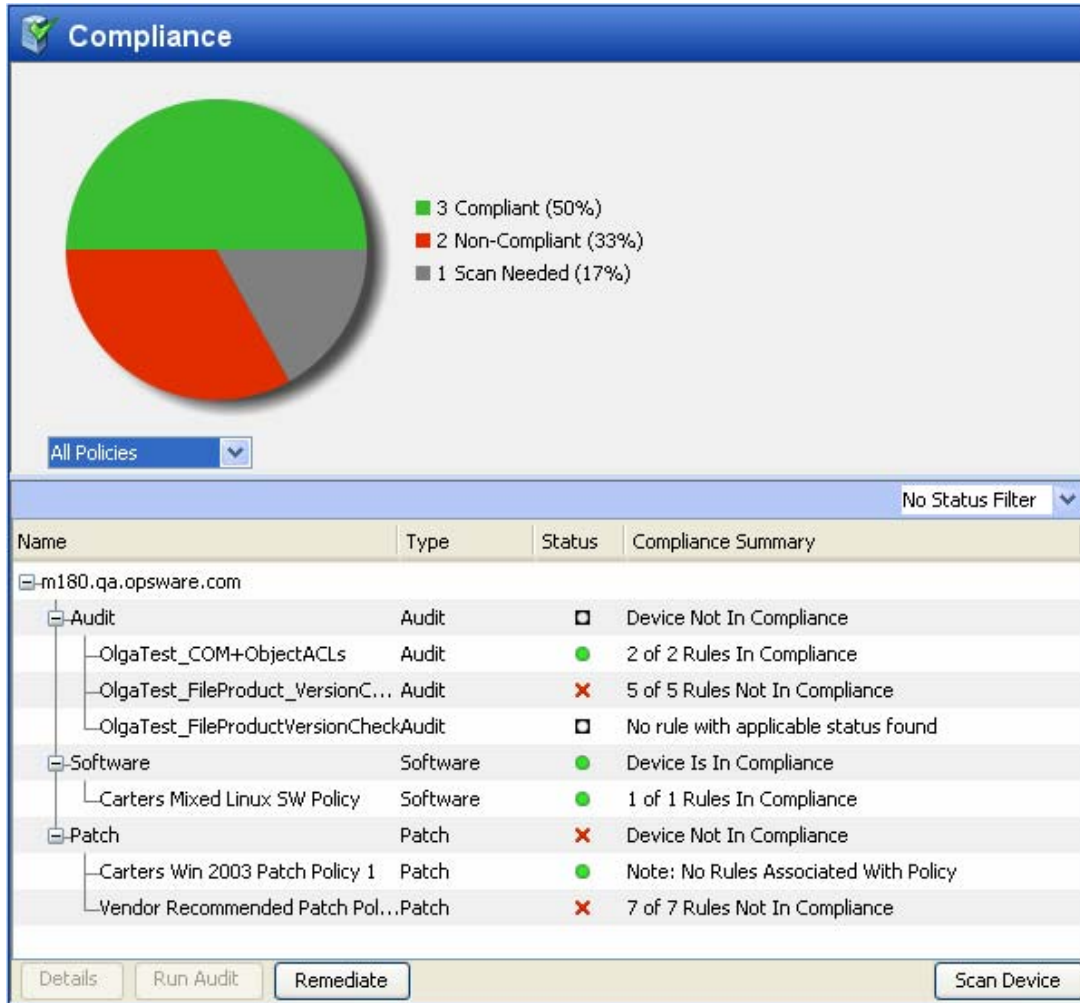
### **Device Explorer Compliance Summary Pie Chart and Details**

The Device Explorer's Compliance View contains two main sections: the compliance summary pie chart and the compliance summary details list.

- Compliance summary pie chart (upper pane), which provides a graphical display of the overall compliance status for all policies attached to the selected server, and breaks down the percentage of each status level by category, such as Audit, Software, and AppConfig. This pie chart can also be filtered to show status only for a specific compliance category.
- Compliance summary details (lower pane) chart, which allows you to drill down in each category to see overall compliance status for each category, the individual policies contained in each category and the compliance status for each policy (Compliant, Non-Compliant, and so on), and a summary description for each. Depending upon your selection, you can launch actions to remediate Non-Compliant policies, such as

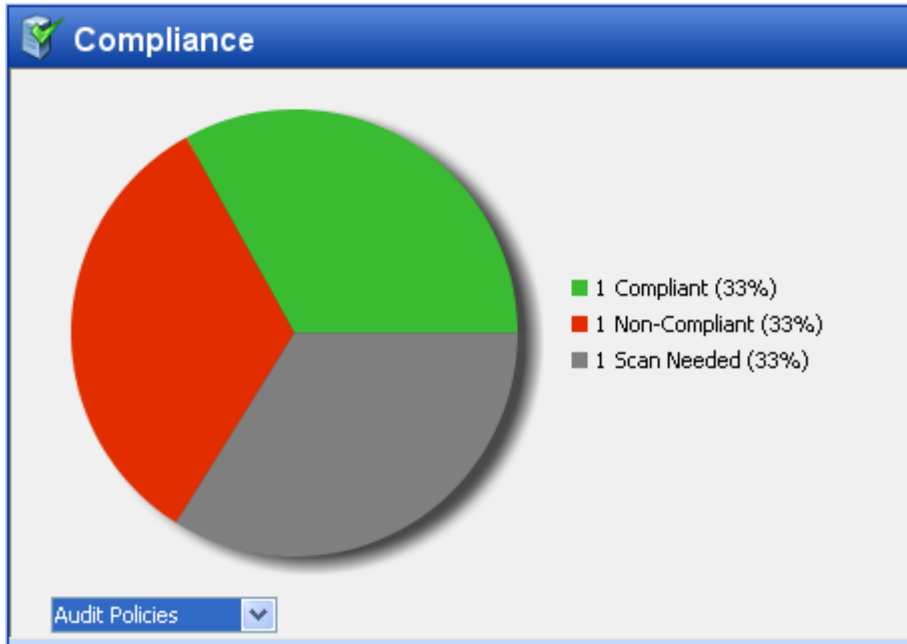
scanning the device for compliance, running an audit, viewing details of a policy, as shown in Figure 3-1.

Figure 3-2: Compliance Summary Pie Chart for an Individual Server



You can select the drop-down list beneath the pie chart to view the pie chart filtered by each compliance test category, such as selecting Audits Policies, as shown in Figure 3-3.

Figure 3-3: Compliance Summary Pie Chart Showing Compliance Levels for Audits Only



You can also choose to filter the compliance policy breakdowns in the details pane below the pie chart to see all compliance policies that contain a certain compliance status. For example, in Figure 3-4, the compliance view has been filtered to show only all compliance policies that are non-compliant.

Figure 3-4: Server Compliance Filtered to Show Only Non-Compliant Policies

Non-Compli...			
Name	Type	Status	Compliance Summary
m180.qa.opsware.com			
[-] Audit	Audit	<input type="checkbox"/>	Device Not In Compliance
└OlgaTest_FileProduct_VersionCheck	Audit	✘	5 of 5 Rules Not In Compliance
[-] Patch	Patch	✘	Device Not In Compliance
└Vendor Recommended Patch Policy f...	Patch	✘	7 of 7 Rules Not In Compliance

At the bottom of the table, there are four buttons: 'Details', 'Run Audit', 'Remediate', and 'Scan Device'.



In the above example, the Compliance View details pane shows all Non-Compliant policies attached to the server. A policy is considered Non-Compliant if at least one of the rules configured in the policy does not match the configuration on the server.

For a list of the actions you can take for a compliance test, such as remediate or scan device, see “Compliance Remediation” on page 285

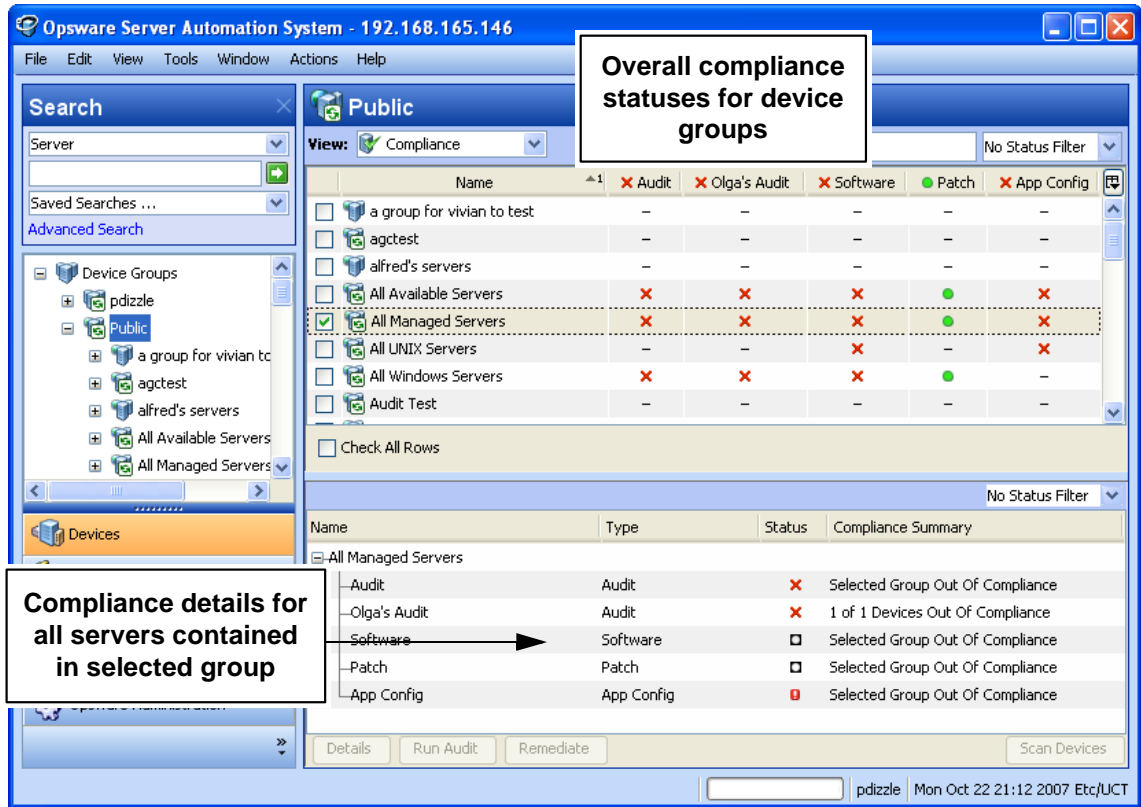
### **Viewing Compliance for Multiple Servers**

To view compliance information for groups of servers, perform the following steps:

- 1** From the Navigation pane, select **Devices ► Device Groups**.
- 2** In the Device Groups tree, select Public. Or, select your own user’s group list. The Content pane on the right side displays the contents of all the groups in the list, either all public groups or all groups your user created.
- 3** From the Views drop-down list above the Content pane, select Compliance.
- 4** For one or more of the device groups (or, any servers), select the check mark next to it to include it in the Compliance View’s Details pane.

- 5 The Contents pane displays compliance summary and detail information for the selected group, as shown in Table 3-1.

Figure 3-1: Compliance View for Device Groups



- 6 You can use the Status Filter drop-down list to filter the view by status; for example, you could choose to view only those groups that show a compliance status of Non-Compliant.

In the Details pane, you can select one of the categories, and depending upon on the category (and your user's permissions), click one of the action buttons at the bottom of the pane for more details, launch an audit, remediate a software or patch policy, or run a compliance scan on all members of the group.

For a list of the actions you can take for a compliance test, such as remediate or scan device, see "Compliance Remediation" on page 285

### Device Group Compliance – Content Pane

The Device Group Content pane displays a summary of compliance status roll ups for all the group members (and contents of groups) that you have selected from the Navigation pane ► **Devices** ► **Device Groups**.

Compliance status (Compliant, Non-Compliant, Partial, and so on) icons in the column heading at the top of the list indicate the rollup status for all groups in the list. If you move your mouse pointer over the top of the column for a category, pop-up text displays the overall compliance for the compliance category for all the visible groups.

In each row of the list, this view displays compliance status for each group in all four compliance categories for each group in the list, which include Audit, Software, Patch, and App Config, as well as any individually scheduled audits that you choose to display in this view, as shown in Figure 3-2.

Figure 3-2: Compliance Roll Ups for Groups

	Name	Audit	Olga's Audit	Software	Patch	App Config
<input type="checkbox"/>	a group for vivian to test	-	-	-	-	-
<input type="checkbox"/>	agctest	-	-	-	-	-
<input type="checkbox"/>	alfred's servers	-	-	-	-	-
<input type="checkbox"/>	All Available Servers					
<input checked="" type="checkbox"/>	All Managed Servers					
<input type="checkbox"/>	All UNIX Servers	-	-		-	
<input type="checkbox"/>	All Windows Servers					-
<input type="checkbox"/>	Audit Test	-	-	-	-	-

Check All Rows

In Figure 3-2, each compliance category (Audit, Software, and so on) displays a compliance status for all policies of each type that are attached to servers in the group. The group named All Managed Servers, for example, displays all categories as Non-

Compliant ✘ except for the Patch category. This means that other than Patch, more than five percent of the servers in the group have a status of Non-Compliant in for Audits, Software, and App Config (as well as the custom column named Olga's Audits).

The Patch category, however, shows a Compliant ● status, which means that at least 95 percent of patch policies attached to servers in this group have a Compliant status. (For information on how to change the compliance status thresholds for device groups, see "Changing Device Group Compliance Settings" on page 267.)

In addition, the scheduled audit named Olga's Audits" has been added to the list, which shows the status of all the servers targeted by that specific audit. For information on how to add or remove compliance categories, see "Adding and Removing Compliance View Columns" on page 280.

### Device Group Compliance – Details Pane

When you select one or more groups from the Content pane (or all of them), the Details pane displays device compliance status rollup in each compliance category for all members of the group, as displayed in Figure 3-3.

Figure 3-3: Device Group Members Compliance Status Rollup in the Details Pane

Name	Type	Status	Compliance Summary
[-] All Available Servers			
- Audit	Audit	✘	Selected Group Out Of Compliance
- Olga's Audit	Audit	✘	1 of 1 Devices Out Of Compliance
- Software	Software	☐	Selected Group Out Of Compliance
- Patch	Patch	●	Selected Group Is In Compliance
- App Config	App Config	⚠	Selected Group Out Of Compliance
<input type="button" value="Details"/> <input type="button" value="Run Audit"/> <input type="button" value="Remediate"/> <input type="button" value="Scan Devices"/>			

You can use the Status Filter drop-down list to filter the view by status; for example, you could choose to view only those groups that show a compliance status of Non-Compliant.

You can also select one of the compliance categories, and depending upon on the category (and your user's permissions), click one of the action buttons at the bottom of the pane for more details, launch an audit, remediate a software or patch policy, or run a

compliance scan on all members of the group. For a list of the actions you can take for a compliance test, such as remediate or scan device, see “Compliance Remediation” on page 285

### **Viewing Group Compliance in the Device Group Explorer**

To view detailed compliance information about a group of servers, open a group’s Device Group Explorer and select its Compliance View, which shows a rollup of compliance statuses for all members of the group as a whole, as opposed to compliance status for individual servers. This gives you a sense of whether or not the group is compliant for each compliance category, and for all servers in the group (and any sub-groups).

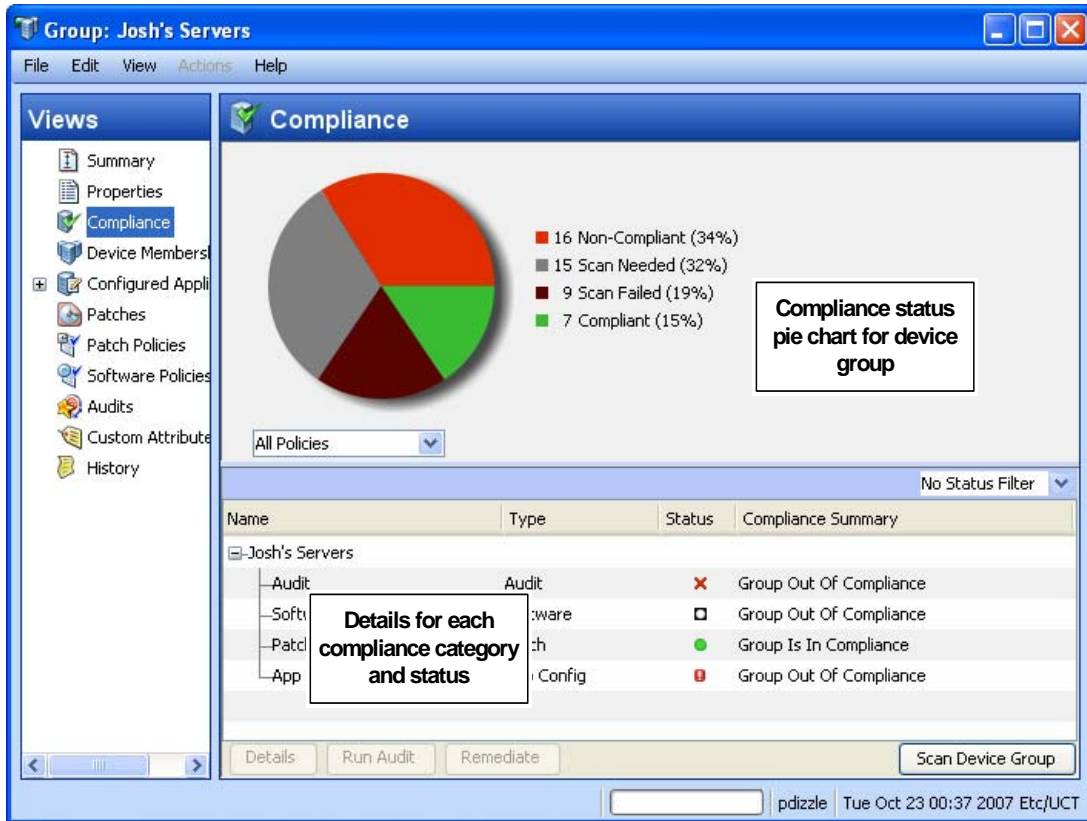
You can use the Status Filter drop-down list to filter the view by status; for example, you could choose to view only those groups that show a compliance status of Non-Compliant. You can also select one of the categories, and depending upon on the type of category (and permissions to view), click one of the action buttons at the bottom of the pane for more details, launch an audit, remediate a software or patch policy, or run a compliance scan on all members of the group.

To view a group of servers in the Device Group Explorer, perform the following steps:

- 1** From the Navigation pane, select **Devices ► Device Groups**.
- 2** In the Device Groups tree, navigate and select Public (or, select your own user’s group list), and then select a group specific group.
- 3** From the **Actions** menu, select **Open**.

- 4 From the View pane of the Device Group Explorer, select Compliance. The Compliance View displays summary and rollup compliance status information about all servers in the group. shows the Compliance View for a device group in the Device Group Explorer.

Figure 3-4: Device Group Compliance View



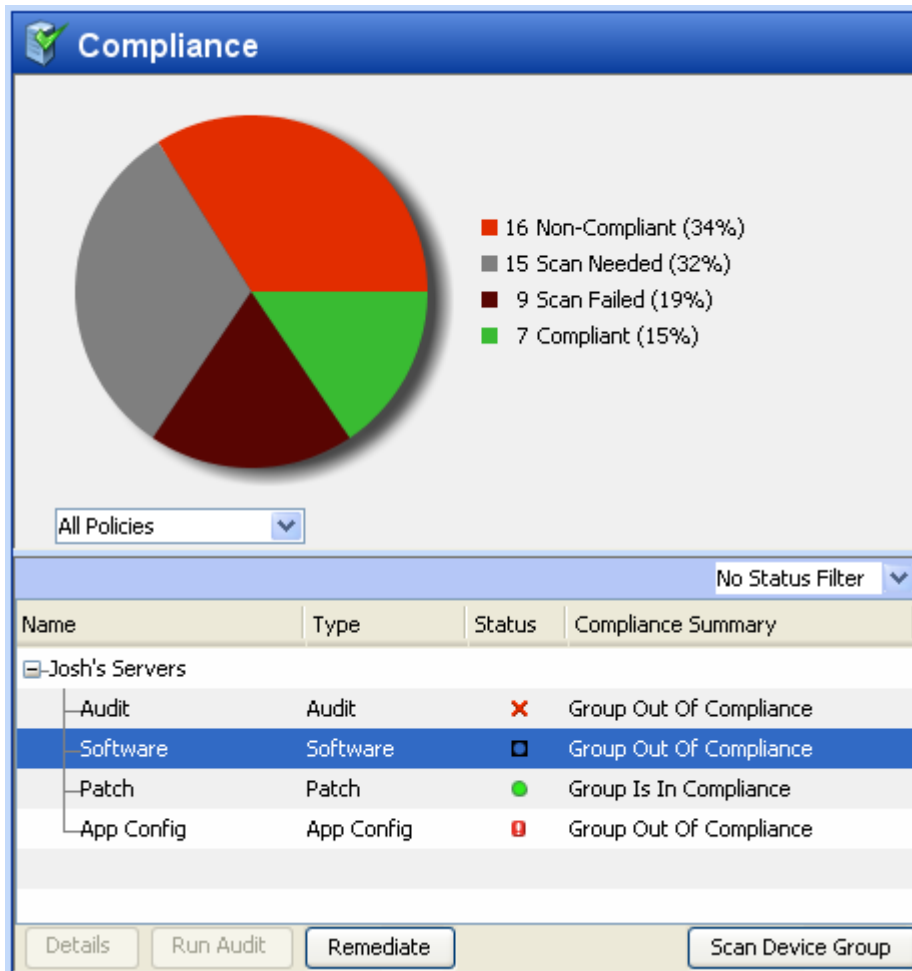
### Device Group Compliance Summary Pie Chart and Details Pane

The Device Group Explorer's Compliance View contains two main sections:

- Compliance summary pie chart (upper pane), which provides a graphical display of the overall compliance status for all policies attached to all the servers in the group, and breaks down the percentage of each status level by category, such as Audit, Software, and AppConfig. This pie chart can also be filtered to show status only for a specific compliance category.
- Compliance summary details (lower pane) chart, which allows you see device group compliance status for each category (Compliant, Non-Compliant, and so on) and view

a summary for each. Depending upon your selection (and your user's permissions), you can launch actions to remediate non-compliant policies, scan device for compliance, run an audit, view details of a policy, as shown in Figure 3-5.

Figure 3-5: Device Group Browser Compliance Summary Pie Chart and Details Pane



By default, a device group is Non-Compliant if more than five percent of the servers in the group have a status of Non-Compliant. You can filter the pie chart to show only those servers that have a specific compliance, such as, show all servers in the group that have a status of Compliant.


Figure 3-5, shows that sixteen of the policies (thirty four percent) attached to servers in the selected device group meet have a status of Non-Compliant. You can select a compliance category from the Details pane below the pie chart (for example, Software),

and perform a remediation action on the group, depending upon your user's permissions. For example, you can select Software and then click Remediate to remediate the Software Policy on to the servers in the group.

## Adding and Removing Compliance View Columns

When you view compliance for device groups in the Compliance View, by default, all four compliance categories are displayed as columns in the Content pane – Audit, Software, Patch, and App Config. You can, however, add or remove any of these categories, as well as add or remove any individual policy in each category.

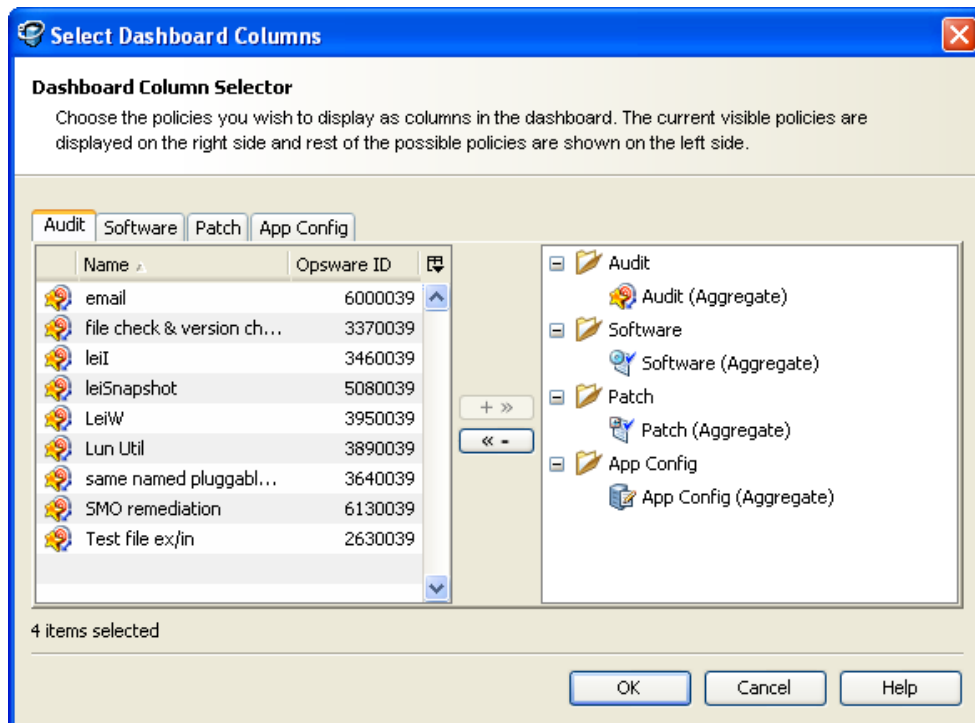
To add or remove device group compliance categories in the Compliance View, perform the following steps:

- 1** From the Navigation pane, select **Devices ► Device Groups**.
- 2** In the Device Groups tree, navigate and select Public (or, select your own user's group list), and then select a specific group. The Content pane shows a list of all four compliance categories and statuses for each member of the group.
- 3** To add or remove a category, click the Column Selector  button on the upper right corner of the Content pane.
- 4** In the Select Dashboard Columns window, the left side of the window displays four tabs, one for each compliance category and all compliance policies in those categories your user has permissions to see. The right side of the window displays



the currently visible policies in each category in the Compliance View. By default, the Compliance View displays the Aggregate (rollup) of all policies in the category, as shown in Figure 3-6.

Figure 3-6: Select Compliance View Columns



- 5 To add an individual policy as a column in the Compliance View, from the left side, select a compliance category tab and then a policy and click the right arrow button.
- 6 To remove an individual policy or an aggregate column from the Compliance View, select one from the right-side of the window and then click the left arrow button.
- 7 When you are finished, click **OK**. You can now view your changes in the Compliance View.

## Filtering By Compliance Status


When you view compliance for groups of servers (and individual servers) in the Compliance View, you can filter the view to show only groups and servers that have at least one server that matches a specific compliance status for any of the displayed compliance categories.

For example, when you select a group and then select Compliance View, you can use the status filter to only show members of the selected group (individual servers and those in any sub-groups) that have a Non-Compliant status for each of the compliance categories, such as Audit, Software, and so on.

To filter the Compliance View by compliance status, perform the following steps:

- 1** From the Navigation pane, select **Devices ► Device Groups**.
- 2** In the Device Groups tree, navigate and select Public (or, select your own user's group list), and then select a group specific group. The Content pane shows the Compliance View statuses for all members of the selected group.
- 3** To filter this view by compliance status, select one from the compliance status drop-down list, as shown in Figure 3-7.

Figure 3-7: Compliance Status Drop-Down List



- 4** The Compliance View displays only those members of the group (individual servers and those in any sub-groups) have a status of Non-Compliant.
- 5** You can select any of the servers or sub-groups in the group listed, and Details pane below will show the compliance status information for those servers. You can further filter the Details pain by using the status filter on the upper right corner of the Details pane.

## Refreshing For Latest Compliance Information

When you first select the Compliance View, the information displayed shows the latest information reported from the Opsware core for each compliance category. It is possible, however, that a server's configurations has changed since you last looked at the Compliance View. It is also possible that a policy has changed since you last viewed server and groups in the Compliance View.

If this is the case, you might want to scan for compliance, or rerun an audit in order to generate new data for the Compliance View to display.

As a best practice, it's useful to refresh the Compliance View to ensure that you are looking at the latest compliance information in your core. To get the latest compliance information from the core, from the **View** menu, select **Refresh**, click **Refresh**, or press F5.

## Scanning for Compliance

When you scan for compliance, you are scanning the servers targeted by a compliance policy – Software, Patch, or App Config – to determine if the target server configurations match the policy's rule definitions. For an audit, when you run an audit it checks the target server configuration to determine the extent to which it matches the audit's rule definitions.

For example, a compliance scan can check to see what patches are installed on a computer, compare that with a patch or software policy, and then return the results to the Compliance View. Or, a compliance scan can check the contents of a configuration file on a server in order to determine if it matches the rules defined in an application configuration.

In the Compliance View, you can perform a compliance scan for the Software, Patch, and App Config compliance categories. Audits do not have a scan feature, but running an audit achieves the same results.

Specifically, each different feature category performs the following actions when scanning for compliance:

- **Software Compliance Scan:** Compares configuration files on a server to determine if they match the values stored in the software policies attached to the server or group of servers. The results of this scan show you the servers that are in compliance (have all required software policy items installed) and the servers that are out of compliance (do

not have all required software policy items installed).For more information on scanning software policies, see “Checking Software Compliance Scan” on page 468.

- **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 show you the servers that are in compliance (have all required patches installed) and the servers that are out of compliance (do not have all required patches installed). Scanning for compliance relates only to Windows patching; Unix patching is encompassed within software policies.

For more information on patch compliance see “Patch Management for Windows” on page 323.

- **App Config Compliance Scan:** Compares configuration files on a server with the template definitions defined application configurations that are attached to that server. The results of this scan show you the servers that are in compliance (configuration file definitions match the configuration templates) and the servers that are out of compliance (configuration file definitions do not match the configuration templates).

For more information on App Config compliance, see “Application Configuration Compliance” on page 302.

## Exporting Compliance View Information

If you want to view all the information displayed in the Compliance View to a file, you can export the view to either .html or .csv.

To export Compliance View information to a file, perform the following steps:

- 1** To view the Compliance, from the Navigation pane, select **Devices** ► **Device Groups**.
- 2** Select a group that you want to view compliance for, and from the **View** menu, select **Compliance**.
- 3** Right-click inside the Contents pane and select **Export**.
- 4** In the Export Compliance View window, enter a name for the file, and choose if you want to export to .html or .csv. You can also change the encoding if you want the saved file to use a specific encoding scheme.
- 5** Click **Export**.

## Compliance Remediation

In addition to providing compliance status information for servers and groups, the Compliance View enables you to remediate server configurations that are not in compliance with your organization's standards, as defined by your audit, software, patch, and application configuration compliance policies.

Generally speaking, the act of remediating a server or group of servers means finding how and where a server or group is out of compliance (Non-Compliant), and then making sure that a server's actual configuration conforms to your compliance policies.

From the Compliance View for a server or group of servers, you can perform the following actions:

- Remediating a software or patch policy
- Running, viewing, and remediation audit results
- Pushing an application configuration on to a server
- Running a compliance scan for patches, software, or application configurations to get the latest compliance information for your servers.

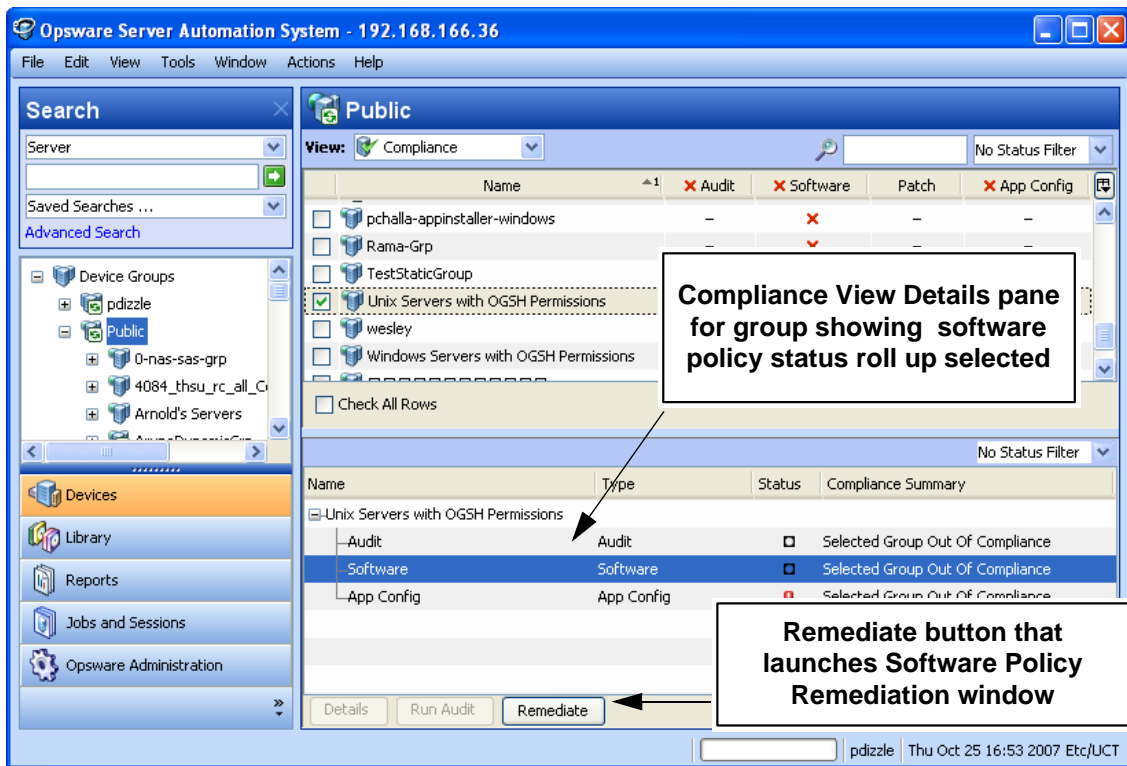
When you select a server or group of servers from in the Compliance View, or view them in the Device or Device Group Explorer, the Details pane provides action buttons for actions that help you discover and remediate out of compliance policies. The type of actions available depending upon which varies depending upon the type of policy,

whether you select a single server or group of servers, and whether or not you select an individual policy, multiple policies, or the roll up of a compliance category (such as Audit, Software, Patch, or App Config) in the Details pane.

### Group Compliance Remediation

Figure 3-8 displays how the Compliance View enables a group's compliance remediation options as located in the Details pane.

Figure 3-8: Details Pane Displaying Remediation Actions Available for a Group of Servers



In Figure 3-8, the Details pane for the selected group shows a summary of all policies attached to all servers in the group (and all servers in any sub-groups) arranged by compliance category – Audit, Software, Patch, and App Config. When you select a group, you can only remediate an entire category of policies, such as all software or patch policies attached to all servers in the group that are out of compliance. If you select the

Software category in the Details pane, the **Remediate** button activates and when clicked launches the Software Policy Remediation window, allowing you to remediate any out of compliance policy configurations for all servers in the group.

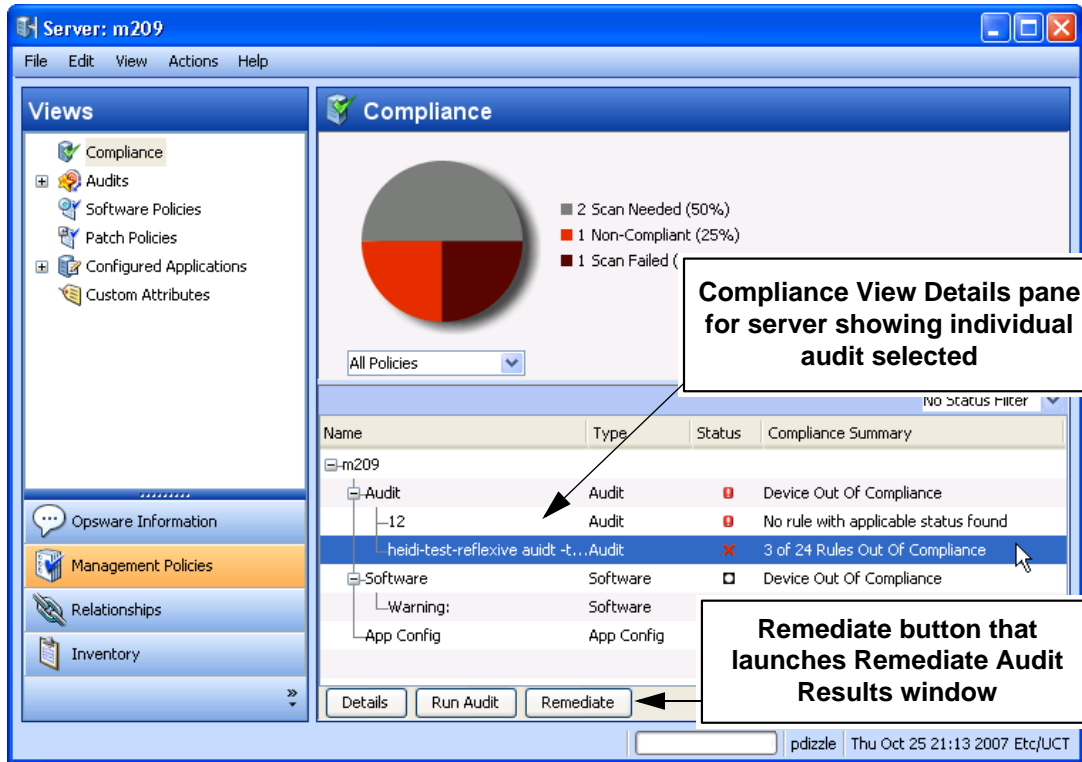
You can view this same information and access these option by selecting the group, and from the Actions menu select **Open**. Doing this launches the Device Group Explorer and displays the same Details pane for the group, along with the action buttons at the bottom of the pane.

### **Compliance Remediation for Servers**

With groups, your remediation options always apply to all members of the group. For an individual server, however, you can launch the server's Device Explorer and you can remediate either all or specific policies that are attached to the server. For example, you can launch a server, and from the server's Device Explorer select **Management Policies** ► **Compliance**, and view all the compliance policies attached to the server.

From the Details pane, you can select an audit or a software policy, and view the audit, run the audit, remediate the software policy, scan the device for compliance, and so on, as shown in Figure 3-9.

Figure 3-9: Device Explorer Showing Compliance and Remediation Options





## Audit Compliance

Opsware's Audit and Remediation feature allows you to define server configuration policies in an *audit*, which helps ensure that the servers in your facilities meet your audit policy standards. An audit consists of a collection of rules that you can define to model those standards.

For example, an audit might consist of Windows COM+ configurations, registry settings, services, file system settings, hardware configuration, user and group password settings, software installation, packages, storage settings, and so on, that define an ideal server configuration. Or, the audit might represent a negative server configuration that enables you to determine the way a server should *not* be configured.

*Audit compliance* determines if the rules defined in a recurring audit match the actual server configuration for all servers targeted by the audit. The Compliance View allows you to see both the aggregate and individual compliance status of all audits that run on a recurring schedule on a server or group of servers. If any of the audits are Non-Compliant, you can remediate any differences found between the audit and the audit's target server or servers.

The Compliance View derives audit compliance servers and groups of servers from regularly scheduled audits. For more information on audit remediation, and how to create and schedule recurring audits, see "Audit Compliance Remediation" on page 290.

## Audit Compliance Status

Audit compliance status is determined by the following criteria:

- **Audit Compliance – Single Server:** If a single rule inside an audit does not match the target server's configuration, then the server's audit compliance status is Non-Compliant. The Details pane of a server's Device Explorer window shows the Audit category as Non-Compliant, and the summary column indicates how many rules are Non-Compliant out of the total number of rules.

For example, if an audit has ten rules, and four of the rules are Non-Compliant, then the audit's status is listed as Non-Compliant and the summary description reads: "4 of 10 Rules Out of Compliance."

If more than one audit targets the server, and if at least one of those audits is Non-Compliant, then the aggregate compliance status for audits is displayed as Non-Compliant as well. You can expand the Audit category of the Details pane to see which

of the audits are not in compliance, as well as a breakdown of how many rules in each audit are in compliance or out of compliance.

- **Audit Compliance – Device Groups:** An audit that targets a group of servers (and all the servers in all sub-groups) is considered Compliant if at least 95 percent of the servers in the group that are targeted by the audit have a compliance status of Compliant.

If more than five percent of the servers in the group targeted by an audit have an audit compliance status of Non-Compliant, then the aggregate compliance for audits will display as Non-Compliant.

However, if more than two percent of all servers in the group have a status of Non-Compliant, and at least ninety percent the servers in the group have a status of Compliant, then the Audit compliance level for the group will have a status of Partial.

The Details pane for a group of servers in the Compliance View shows whether or not all of the audits are compliant or not, but does not expand to show a breakdown of individual servers and audits.

You can modify the thresholds used to determine compliance for groups of servers. For more information, see “Changing Device Group Compliance Settings” on page 267.

## **Audit Compliance Remediation**

The Compliance View allows you to view all audits that target the a server or group of servers and to remediate those results that are out of compliance, to ensure that a server's configuration complies with the rules defined in an audit.

For each audit rule that is out of compliance on the target server (the server's configuration either mismatched the rule definition or simply did not exist), remediation copies the rule object onto the target server so it matches the rule. Or, in the case of a value-based audit rule, changes the target server's configuration to match the rule.

For example, you have an audit that consists of several Windows registry settings, COM+ definitions, and a few Opsware TON predefined Users and Group rules. After the audit runs against an actual Windows server, it is possible that several of the rules are out of compliance. When you remediate, it copies the correct objects from the rules and

changes the out of compliance settings discovered in the audit results. For a group of servers, remediation has the same results, only the remediation operation applies to all servers in the group, including all the server contained in any sub-groups.

### **Remediating Audits for One or More Servers**

You can remediate an audit that is attached to a single server or one attached to several servers that you select in the Device Groups list.



You cannot run an audit on a group of servers from the Compliance View. However, you can create an audit that runs against a group of servers and remediate those audit results for a group of servers from the Audit Results window. For more information, see “Configuring an Audit” on page 152 and “Viewing and Remediating Audit Results” on page 220.

To remediate an individual audit on one or more servers, perform the following steps

- 1** To remediate an individual audit on a single server in the Device Explorer, from the Navigation pane, select **Devices > Servers > All Managed Servers**.
- 2** From from the **Actions** menu, select **Open**.
- 3** From inside the Device Explorer’s View pane, select **Management Policies > Compliance**.
- 4** In the Details pane of the Compliance View, expand the Audit category and select an individual policy.

Or

- 5** To remediate an individual audit for several servers, from the Navigation pane, select **Devices > Device Groups** and select a group.
- 6** Select multiple servers by selecting the check box next to the server.
- 7** From the View drop-down list, select Compliance.
- 8** In the Details pane of the Compliance View, expand the Audit category and select an individual audit that is targeting all of the selected servers.
- 9** The following list describes the types of remediation you can perform for an individual audit on a single server or one that targets multiple servers:

- **Details:** Launches the Launch Audit Results window, which allows you to view all differences found between the audit and the target, and to remediate the differences by rule or by server. For more information, see “Viewing and Remediating Audit Results” on page 220.
- **Run Audit:** Launches the Run Audit task window and allows you to run the audit immediately or schedule to run the audit at a later time. The audit will run against all servers targeted by the audit.
- **Remediate:** Launches the audit Remediate Audit Results window, which allows you to remediate target server configurations that are out of compliance with the audit rules. You can remediate differences by rule or by server. For more information, see “Viewing and Remediating Audit Results” on page 220.
- **Scan Device:** Launches the Scan Compliance window, which enables to you scan the selected server for all Software, Patch, and App Config policies attached to the server. This does not have any effect on the audits that target this server. For more information, see “Scanning for Compliance” on page 283.

## Software Compliance

Opsware's Software Management feature allows you to create *software policies* that enable you to install software and configure applications simultaneously. A software policy can contain several different kinds of items, such as packages, RPM packages, patches, application configurations, and other software policies. After creating a software policy, you can attach it to servers or groups of servers.

Software compliance indicates whether or not the items in a software policy are compliant with the actual server configuration. If the actual server configuration does not match the software policy definitions, then the server's software policies are Non-Compliant.

The Compliance View derives software compliance information for software policies when you scan a server or group for software compliance. For more information on how to scan for software compliance, see "Scanning for Compliance" on page 283.

For more information software management and how to create and manage software policies, see "Software Management Setup" on page 9.

## Software Compliance Status

Software compliance status is determined by the following criteria:

- **Software Compliance – Single Server:** If at least one item in a software policy does not match what is discovered (or does not exist) on the server the policy is attached to, the server's software compliance status is Non-Compliant. The Details pane of a server's Device Explorer window shows the Software category as Non-Compliant and the summary column indicates how many rules (software policy items) are Non-Compliant out of the total number of rules.

For example, if a software policy contains ten items, and six of the items are Non-Compliant, then the software policy's status is listed as Non-Compliant and the summary description reads: "6 of 10 Rules Out of Compliance."

If more than one software policy targets a single server, and if at least one of those policies is Non-Compliant, then the aggregate compliance status for Software is displayed as Non-Compliant as well. You can expand the Software category of the Details pane to see which of the policies are not in compliance, as well as a breakdown of how many rules in each policy are either in or out of compliance.

- **Software Compliance – Device Groups:** A software policy that is attached a group of servers is considered Compliant if at least 95 percent of the servers in the group (and all the servers in all sub-groups) have a software compliance status of Compliant. If

more than five percent of the servers in the group have a software compliance status of Non-Compliant, then the aggregate compliance status will display Non-Compliant.

However, if more than two percent of all servers in the group have a status of Non-Compliant, and at least ninety percent the servers in the group have a status of Compliant, then the Software compliance level for the group will have a status of Partial.

The Details pane for a group of servers in the Compliance View shows whether or not all of the software policies are compliant or not, but does not expand to show a breakdown of individual servers and policies.

You can modify the thresholds used to determine compliance for groups of servers. For more information, see “Changing Device Group Compliance Settings” on page 267.

## **Software Compliance Remediation**

The Compliance View allows you to view all software policies attached to a server or groups of servers and to remediate those servers that are out of compliance, and in the process ensure that a server's software configuration complies with the software policy definition.

For each software policy item – such as software, packages, patches, scripts, application configurations –software remediation installs (or for a script, executes) those items on the target server. If the items do not exist on the server, then they get installed. If the items existed but did not match the policy, they get updated with the correct version.

For example, you have a software policy that consists of several packages, patches, a few scripts, and an application configuration all organized in the order in which they are to be installed and executed. First, you remediate the software the policy onto a servers to make sure the server is in compliance with your company's software installation

standards. Over time, some of the items in the software policy get updated – such as a new set of packages gets added– and for whatever reason, a software item on the server was uninstalled on the target server.

When you perform a software compliance scan, the scan determines the server's compliance status by comparing the software policy contents with the actual software installed on the server. Even if only one software item attached to one of the server is not in compliance with the policy, the server will have a software compliance status of Non-Compliant.

When you remediate a server or group of servers, the patches, packages, and application configurations specified in the policy are installed and applied in the order specified in the policy. For a group of servers, remediation has the same results, only the remediation operation applies to all servers in the group, including all the server contained in any sub-groups. (For information on how to change this setting, see "Changing Device Group Compliance Settings" on page 267.)

For more information on software policy remediation, see "Software Management" on page 437.

### **Remediating Software Compliance – Single or Multiple Servers**

When you remediate software compliance for a single server or multiple servers, you can choose to remediate all of the policies attached to the servers or select to remediate individual policies.

You can remediate software policies for a single server by viewing the server's Device Explorer or you can remediate software policies for multiple servers by selecting them in the Device Groups list.

To remediate software policies on single or multiple servers, perform the following steps:

- 1** To remediate software policies for a single server in the Device Explorer, from the Navigation pane, select **Devices > Servers > All Managed Servers**.
- 2** From from the **Actions** menu, select **Open**.
- 3** From inside the Device Explorer's View pane, select **Management Policies > Compliance**.

- 4** In the Details pane of the Compliance View, expand the Software category and select an individual software policy or the top level Software category, which will enable you to remediate all of the policies attached to the server.  
  
Or
- 5** To remediate software policies for multiple servers, from the Navigation pane, select **Devices ► Device Groups** and select a group.
- 6** In the Content pane that shows a list of servers that belong to the group, select multiple servers by selecting the check box next to the server.
- 7** From the View drop-down list, select Compliance.
- 8** In the Details pane of the Compliance View, expand the Software category and select a software policy that is attached to the selected servers. Or, select the top level Software category if you want to remediate all of the software policies attached to the selected servers.
- 9** From the bottom of the Details pane, you have the following options:
  - **Remediate:** Remediate the selected software policy or policies against the selected server or servers.
  - **Scan Device:** Launches the Scan Compliance window, which enables to you scan the selected server for all Software, Patch, and App Config policies attached to the server. For more information, see “Scanning for Compliance” on page 283.

### **Remediating Software Compliance – Groups**

When you remediate software policies for a group or multiple groups of servers, you can remediate all the policies attached to all servers in the group or multiple groups. However, when you select a group or multiple groups, you can only remediate *all* of the software policies attaches to all the servers in the group and any sub-groups.



By default, remediate actions apply to all members of a group, including servers the belong to any sub-groups. If you would like to change this setting, see “Changing Device Group Compliance Settings” on page 267.

---



To remediate software policies for groups or multiple groups of servers, perform the following steps:

- 1** To remediate software policies for a single server in the Device Explorer, from the Navigation pane, select **Devices > Servers > All Managed Servers**.
- 2** From from the **Actions** menu, select **Open**.
- 3** From inside the Device Explorer's View pane, select **Management Policies > Compliance**.
- 4** In the Details pane of the Compliance View, expand the Software category and select an individual software policy or the top level Software category, which will enable you to remediate all of the policies attached to the server.

Or

- 5** To remediate software policies for multiple servers, from the Navigation pane, select **Devices > Device Groups** and select a group.
- 6** In the Content pane that shows a list of servers that belong to the group, select multiple servers by selecting the check box next to the server.
- 7** From the View drop-down list, select Compliance.
- 8** In the Details pane of the Compliance View, expand the Software category and select a software policy that is attached to the selected servers. Or, select the top level Software category if you want to remediate all of the software policies attached to the selected servers.
- 9** From the bottom of the Details pane, you have the following options:
  - **Remediate**: Remediate the selected software policy or policies against the selected server or servers.
  - **Scan Device**: Launches the Scan Compliance window, which enables to you scan the selected server for all Software, Patch, and App Config policies attached to the server. For more information, see "Scanning for Compliance" on page 283.

## Patch Compliance

Opware's Patch Management (Windows and Unix) feature enables you to identify, install, and remove patches on managed servers and groups of servers. With Windows Patch Management you can identify and install patches for the Windows 2000, Windows 2003, and Windows NT 4.0 operating systems, include Service Packs, Update Rollups, and hotfixes.

In the Compliance View, you can view compliance status for Windows patch policies (Unix patch compliance is not currently supported in this release) in order to see whether or not your Windows servers have the correct patches installed on them. During a Patch compliance scan, Patch Management checks managed servers and public device groups to determine whether all patches in a policy and a policy exception were installed successfully. If the patches installed (or not installed) on the server does not match the patch policy definitions, then the Compliance View displays the server's patch policies as Non-Compliant.

Compliance scans can be run on a one time basis, or can be scheduled on a recurring basis. You can remediate a patch policy to a server in order to ensure a server's or groups patch compliance.

For more information on Windows patch compliance, see:

- "Patch Compliance" on page 368
- "Verifying Patch Policy Compliance" on page 362
- "Scanning for Compliance" on page 283
- "Scheduling a Patch Compliance Scan" on page 375

## Patch Compliance Status

Patch compliance status is determined by the following criteria:

- **Patch Compliance – Single Server:** If at least one item in a patch policy does not match what is discovered (or does not exist) on the server the policy is attached to, the server's patch compliance status is Non-Compliant. The Details pane of a server's Device Explorer window shows the Patch category as Non-Compliant and the summary column indicates how many rules (patch policy items) are Non-Compliant out of the total number of rules.

For example, if a patch policy contains ten items, and six of the items are Non-Compliant, then the patch policy's status is Non-Compliant and the summary description reads: "6 of 10 Rules Out of Compliance."

If more than one patch policy targets a single server, and if at least one of those policies is Non-Compliant, then the aggregate compliance status for Patch is displayed as Non-Compliant as well. You can expand the Patch category of the Details pane to see which of the policies are not in compliance, as well as a breakdown of how many rules in each policy are either in or out of compliance.

- **Patch Policy – Rule Exception:** If a rule exception is applied to one of the patch policy items, then the server's Patch compliance will display a compliance status of Partial. Patch is the only compliance category that allows rule exceptions at the policy level, and thus us the only category that can have a Partial compliance status.
- **Patch Compliance – Device Groups:** A patch policy that is attached a group of servers is considered Compliant if at least 95 percent of the servers in the group (and all the servers in all sub-groups) have a patch compliance status of Compliant. If more than five percent of the servers in the group have a patch compliance status of Non-Compliant, then the aggregate patch compliance status will display Non-Compliant.

However, if more than two percent of all servers in the group have a status of Non-Compliant, and at least ninety percent the servers in the group have a status of Compliant, then the Patch compliance level for the group will have a status of Partial.

The Details pane for a group of servers in the Compliance View shows whether or not all of the patch policies are compliant or not, but does not expand to show a breakdown of individual servers and policies.

You can modify the thresholds used to determine compliance for groups of servers. For more information, see “Changing Device Group Compliance Settings” on page 267.

For more information on creating and using patches and patch policies, see “Patch Management for Windows” on page 323 or “Patch Management for Unix” on page 401

### **Remediating Patch Compliance – Single or Multiple Servers**

When you remediate patch compliance for a single server or multiple servers, you can choose to remediate either all of the policies attached to the servers or only remediate individual policies.

You can remediate patch policies for a single server by viewing the server's Device Explorer, or you can remediate patch policies for multiple servers by selecting the policies in the Device Groups list.

To remediate patch policies on single or multiple servers, perform the following steps:

- 1** To remediate patch policies for a single server in the Device Explorer, from the Navigation pane, select **Devices > Servers > All Managed Servers**.
- 2** From from the **Actions** menu, select **Open**.
- 3** From inside the Device Explorer's View pane, select **Management Policies > Compliance**.
- 4** In the Details pane of the Compliance View, expand the Patch category and select an individual policy or the top level Patch category, which will enable you to remediate all of the patch policies attached to the server.

Or

- 5** To remediate patch policies for multiple servers, from the Navigation pane, select **Devices > Device Groups** and select a group.
- 6** In the Content pane that shows a list of servers that belong to the group, select multiple servers by selecting the check box next to the server.
- 7** From the View drop-down list, select Compliance.
- 8** In the Details pane of the Compliance View, expand the Patch category and select a software policy that is attached to the selected servers. Or, select the top level Patch category if you want to remediate all of the policies attached to the selected servers.
- 9** From the bottom of the Details pane, you have the following options:

- **Remediate:** Remediate the selected patch policy or policies against the selected server or servers.
- **Scan Device:** Launches the Scan Compliance window, which enables to you scan the selected server for all Software, Patch, and App Config policies attached to the server. For more information, see “Scanning for Compliance” on page 283.

### **Remediating Patch Compliance – Groups**

When you remediate patch policies for a group or multiple groups of servers, you can remediate all the policies attached to all servers in the group or multiple groups. However, when you select a group or multiple groups, you can only remediate all of the patch policies attaches to all the servers in the group and any sub-groups.



By default, remediate actions apply to all members of a group, including servers the belong to any sub-groups. If you would like to change this setting, see “Changing Device Group Compliance Settings” on page 267.

To remediate patch policies for groups or multiple groups of servers, perform the following steps:

- 1** To remediate patch policies for a single server in the Device Explorer, from the Navigation pane, select **Devices > Servers > All Managed Servers**.
  - 2** From from the **Actions** menu, select **Open**.
  - 3** From inside the Device Explorer’s View pane, select **Management Policies > Compliance**.
  - 4** In the Details pane of the Compliance View, expand the Patch category and select an individual patch policy or the top level Patch category, which will enable you to remediate all of the patch policies attached to the server.
- Or
- 5** To remediate patch policies for multiple servers, from the Navigation pane, select **Devices > Device Groups** and select a group.
  - 6** In the Content pane that shows a list of servers that belong to the group, select multiple servers by selecting the check box next to the server.
  - 7** From the View drop-down list, select Compliance.

- 8** In the Details pane of the Compliance View, expand the Patch category and select a policy that is attached to the selected servers. Or, select the top level Patch category if you want to remediate all of the policies attached to the selected servers.
- 9** From the bottom of the Details pane, you have the following options:
  - **Remediate:** Remediate the selected patch policy or policies against the selected server or servers.
  - **Scan Device:** Launches the Scan Compliance window, which enables to you scan the selected server for all Software, Patch, and App Config policies attached to the server. For more information, see “Scanning for Compliance” on page 283.

## Application Configuration Compliance

An application configuration (App Config) manages configuration files on a managed server. An application configuration can manage one or several configuration files for a server or group of servers. Each application configuration is made up of one or more templates which model an ideal configuration state and are targeted to manage configuration values (key-value pairs) for specific files on a server.

For example, you can create an application configuration that manages the hosts file for servers in your data center. You can define the IP address-hostname key-value pairs for a standard Unix hosts file, and then attach the application configuration to a several servers or a group of servers. The application configuration serves as the policy that enables you to ensure that the hosts files on the target servers have the correct IP address-hostname definitions.

Application configuration compliance indicates whether or not all of the Application Configurations attached to a server are compliant with the actual application configuration files on the server. In the case of the hosts file example, if the information inside the hosts file on actual server configuration does not match the values defined in the application configuration, then the server's App Config is Non-Compliant. If more than one application

configuration is attached to a server, and any one of the actual configuration files targeted by the application configuration is different, then the entire server is Non-Compliant in the Compliance View.

Conversely, if there are no differences found between the application configuration and the files on server, then the App Config compliance status is Compliant. All application configurations must be 100 percent compliant for the server's App compliance status to be considered Compliant in the Compliance View.

To check the latest state of a configuration file targeted by an application configuration, you can perform an application configuration compliance scan to determine if there are any differences between the application configuration and the actual configuration files on the server.

For more information on running a compliance scan, see "Scanning for Compliance" on page 283.

For more information on creating and using Application Configurations, see "Application Configuration Management" on page 495.

### **Application Configuration (App Config) Compliance Status**

Application Configuration (App Config) compliance status is determined by the following criteria:

- **App Config Compliance – Single Server:** If any differences are discovered between the application configuration and the actual configuration file on the target server, the server's App Config compliance status is Non-Compliant. The Details pane of a server's Device Explorer window shows the App Config category as Non-Compliant. If the server has several application configurations attached to it, and any one of the actual configuration files targeted by the application configuration is different than the application configuration, then the entire server is considered Non-Compliant in the Compliance View.
- **App Config Compliance – Device Groups:** An application configuration that is attached to a group of servers is considered Compliant if at least 95 percent of the servers in the group (and all the servers in all sub-groups) have an App Config compliance status of Compliant. If more than five percent of the servers in the group have an App Config compliance status of Non-Compliant, then the aggregate compliance status displays a status of Non-Compliant.

However, if more than two percent of all servers in the group have a status of Non-Compliant, and at least ninety percent the servers in the group have a status of Compliant, then the App Config compliance level for the group will have a status of Partial.

The Details pane for a group of servers in the Compliance View shows whether or not all of the application configurations are compliant or not, but does not expand to show a breakdown of individual servers and application configurations.

You can modify the thresholds used to determine compliance for groups of servers. For more information, see “Changing Device Group Compliance Settings” on page 267.

### **Remediating App Config Compliance – Servers and Groups**

Remediation for an application configuration is slightly different than the other compliance category types. Rather than remediating a policy onto a server, as you can with Software or Patch or Audits, to remediate an application configuration you select an application configuration from inside either the Device Explorer or Device Group Explorer and use the Push function to push the values defined in the application to onto the actual configuration files on the server (or group of servers).

When you push an application configuration, all values defined in the application configuration templates add or replace those on the target configuration files.



---

The manner in which some value in an application configuration get pushed – for example, sequences (of lists and scalars) – depends upon how those values have been set in the application configuration inheritance hierarchy and what sequence merge modes have been configured in the configuration template. For more information about sequence merging, see “Sequence Merging and Inheritance” on page 508.

---

To remediate application configurations server or group of servers, perform the following steps:

- 1** To remediate application configuration for a single server in the Device Explorer, from the Navigation pane, select **Devices > Servers > All Managed Servers**.

Or

- 2** To remediate patch policies for a group of servers, from the Navigation pane, select **Devices > Device Groups** and select a group.



- 3** From from the **Actions** menu, select **Open**.
- 4** From inside the Device Explorer's View pane, select **Management Policies** ► **Compliance**.
- 5** In the Details pane of the Compliance View, expand the Configured Applications category and select an individual application configuration that you want to remediate (push). You have several options for viewing and remediating a application configuration:
  - **Run Script:** Executes the application configuration's Data-manipulation script. If the application configuration contains no data-manipulation script, **Run Script**.
  - **Preview:** Launches a file differencing window that allows you to compare the contents between the application configuration and the actual configuration file on the managed server.
  - **Push:** Applies any changes made to the application configuration on to the target configuration file (listed in the Filename field above the value set editor). This also saves all changes made to the application configuration.
  - **Schedule:** Launches the Schedule Job window so you can schedule the Application Configuration push to run at a later date and time or on a recurring basis.

For more information, see "Pushing Changes to a Server or Group" on page 532.



# Chapter 4: SAS Client Reports

## IN THIS CHAPTER

This section discusses the following topics:

- Overview of SAS Client Reports
- Reports Features
- Opsware SAS Client Reports
- User Permissions
- Launching the Reports Feature
- Reports Display
- Running a Report
- Report Results

## Overview of SAS Client Reports

The Opsware SAS Client Reports feature provides comprehensive, real-time information about managed servers, network devices, software, patches, customers, facilities, operating systems, compliance policies, and users and security in your environment. These parameterized reports are presented in graphical and tabular format, and are actionable—which means that you can perform appropriate actions on objects, such as a policy or an audit, within the report. These reports are also exportable to your local file system (as .html, .pdf, or .xls files) to facilitate use within your organization.

This section contains information about the types of SAS Client reports, how to modify report parameters, how to run the reports, and how to perform actions in the report results.

## Reports Features

SAS Client Reports enable you to perform enterprise health assessments by providing the following features:

- Actionable reports that enable you to take the appropriate action on objects within the reports. For example, in the list view of a compliance report, you can select a server and open a Remote Terminal or Server Explorer to browse it, perform an audit, create a snapshot, create a package, and so on.
- A single entry point in the SAS Client Dashboard for all reports.
- Reports that are data-secured—controlled by the user's permissions. You can view all objects that you have read permissions for. You can perform actions on objects that you have write permissions for.
- Reports that are exportable to .html, .pdf, and .xls formats. You can export reports to your local file system for use within your organization.

## Opware SAS Client Reports

Table 4-1 lists the SAS Client Reports by report folders.

Table 4-1: SAS Client Reports

REPORT FOLDER	REPORT TITLE
Server Reports	Servers by Customer
	Servers by Facility
	Servers by Manufacturer
	Servers by Model
	Servers by Operating System
	Servers by Use

Table 4-1: SAS Client Reports (continued)

REPORT FOLDER	REPORT TITLE
Virtualization Reports	Virtualization by Virtual Technology All Virtual Servers Solaris 10 Virtual Servers by Hypervisors (zones only) Resource Allocation by Hypervisors (zones only) VMware ESX 3 Virtual Servers by Hypervisors (VMs only) Resource Allocation by Hypervisors (VMS only)

Table 4-1: SAS Client Reports (continued)

REPORT FOLDER	REPORT TITLE
Compliance Reports	Audit Policy Compliance (All Servers)
	Audit Policy Compliance by Customer
	Audit Policy Compliance by Facility
	Patch Policy Compliance (All Servers)
	Patch Policy Compliance by Customer
	Patch Policy Compliance by Facility
	Software Policy Compliance (All Servers)
	Software Policy Compliance by Customer
	Software Policy Compliance by Facility
	Audit Compliance : Devices by Audit
	Patch Compliance : Servers by Policy
	Software Compliance : Servers by Policy
	Server Audit Compliance
	Server Patch Policy Compliance
	Server Software Policy Compliance
	Server Audit Compliance Detail
	Server Patch Compliance Detail
	Server Software Policy Detail
Non-Compliant Patches By Server	
Non-Compliant Patches By Patch Policy	

Table 4-1: SAS Client Reports (continued)

REPORT FOLDER	REPORT TITLE
SOX	SOX Compliance Summary Audit Results With Failures Audit Results Without Failures Current Users Defined Patch Policies Defined Server Audit Policies Defined Server Audits Defined Software Policies Deleted Users Jobs With an Associated Ticket ID Jobs Without an Associated Ticket ID Servers Audited Without Failures Servers in Compliance With Their Patch Policies Servers in Compliance With Their Software Policies Servers Not Audited or Audited With Failures Servers Not in Compliance With Their Patch Policies Servers Not in Compliance With Their Software Policies Servers With Associated Audits Servers With Attached Patch Policies Servers With Attached Software Policies Servers Without Associated Audits Servers Without Attached Patch Policies Servers Without Software Patch Policies
	Users Created in the Last 30 Days User Groups Membership

Table 4-1: SAS Client Reports (continued)

REPORT FOLDER	REPORT TITLE
User and Security Reports	Client and Feature Permissions Customer/Facility Permissions and Device Group Permission Overrides User Groups Memberships
Network Reports	Connections by Network Device Connections by Server Duplex Compliance (All Servers) Duplex Compliance by Customer Duplex Compliance by Facility

See the following documentation for more information about the SAS Client features that support information in these reports:

- “Software Management” on page 437
- “Audit and Remediation” on page 137
- “Exploring Servers and Groups in the SAS Client” on page 65
- “Patch Management for Windows” on page 323
- “Patch Management for Unix” on page 401
- “NAS Integration” in the *Opware® SAS User's Guide: Server Automation*
- “Server Management in SAS Web Client” in the *Opware® SAS User's Guide: Server Automation*



## User Permissions

Reports are controlled by the user's permissions. You can view all objects that you have read permissions for, and you can perform actions on objects that you have write permissions for.

To view or run a network report, NAS Integration must be installed. See "NAS Integration" in the *Opware® SAS User's Guide: Server Automation*.

To view or run a user and security report, system administrator permissions are required.

## Launching the Reports Feature

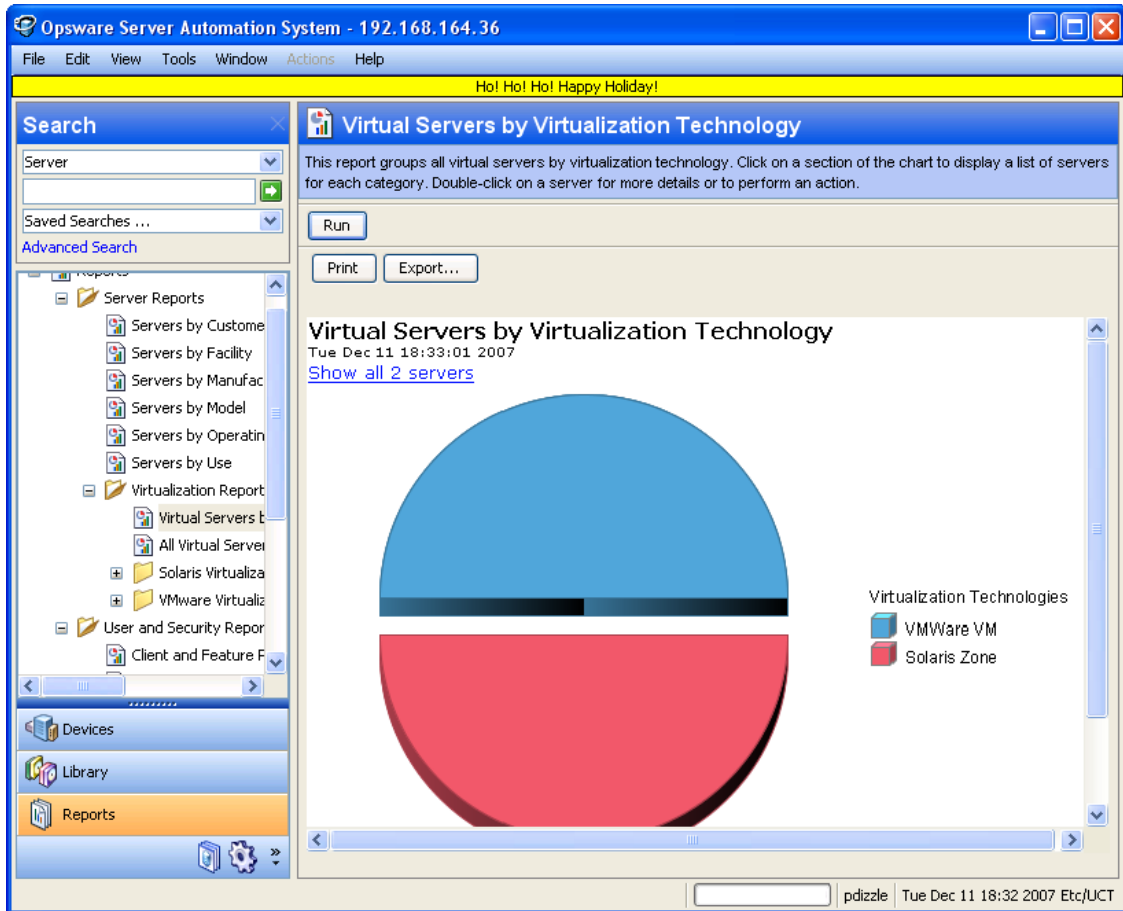
To launch the Reports feature, perform one of the following steps:

- From the **View** menu, select **Reports** ► **Dashboard**.
- From the **View** menu, select **Reports** ► **Reports**.
- From the Navigation pane, select Reports.

## Reports Display

The Reports feature display consists of a Search pane, report parameters, report folders, and other filtering tools.

Figure 4-1: The Reports Feature Display



### Search Pane

In the Reports feature, you can use the SAS Client Search feature to find reports by defining specific filter criteria. See "SAS Client Search" in the *Opware® SAS User's Guide: Server Automation*.

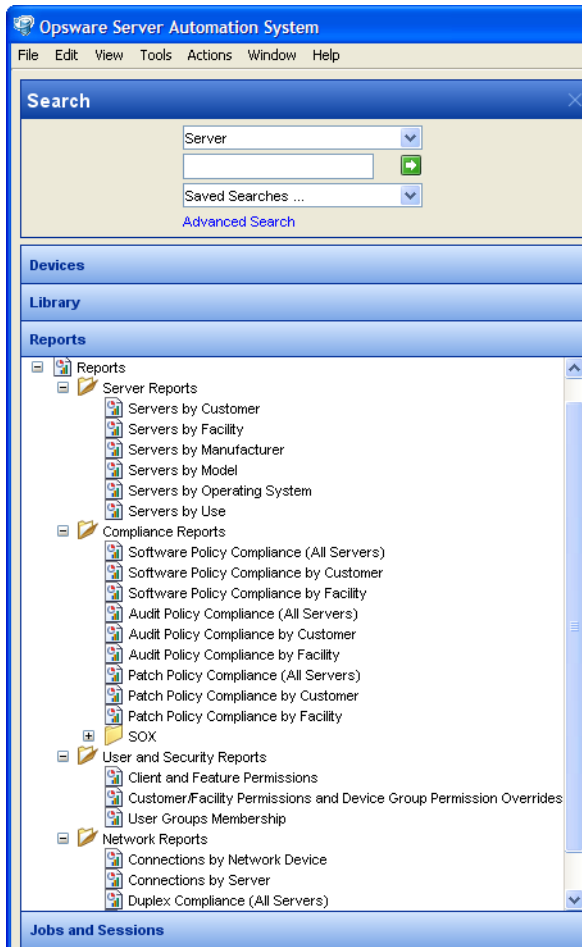
### **Report Folders**

Reports are organized in the following folders according to regulatory or IT best practice standards:

- **Server Reports:** This folder contains reports about servers by customer, facility, manufacturer, model, operating system, and server usage.
- **Compliance Reports:** This folder contains reports about compliance for software policies, audit policies, and patch policies by servers, customer, and facility.
- **SOX Reports:** This folder contains reports about compliance standards based on Sarbanes-Oxley, including the COSO process model and the CobiT control model.
- **Network Reports:** This folder contains reports about connections and duplex compliance for network devices and servers. You must have Opsware NAS installed to see this folder in the Navigation pane.
- **User and Security Reports:** This folder contains reports about client and feature permissions; customer, facility, and device group permissions; and user group memberships. You must have system administrator permissions to see this folder in the Navigation pane.
- **Custom Reports:** This folder contains any custom reports you have created. For more information on creating custom reports, check the Opsware Knowledge Base.

- Figure 4-2 illustrates the Report folders in the Navigation pane, including the reports you will find in each folder.

Figure 4-2: Report Folders



## Report Parameters

Many reports require input parameters in order to be run. For reports that require parameters, you can run the report with its default parameter values or modify the parameter values. If you want to run a report that includes or excludes certain servers, customers, or hardware models, you need to specify this criteria in the report parameters. See “Running a Report” on page 317.

## Running a Report

To run a report, perform the following steps:

- 1** From the Navigation pane, select Reports.
- 2** Expand the Reports folder and then expand the Server Reports and Virtualization Reports.
- 3** Select one of the virtualization report listed in the folder.
- 4** If there are no report parameters in the Content pane, click **Run**.
- 5** If there are report parameters in the Content pane, you can either use the default parameters or change them:
  - To use the default report parameters, click **Run** to run the report.
  - To change the report parameters, see “Modifying Report Parameters” on page 317.

### **Modifying Report Parameters**

To modify the default parameters and run a report that includes certain servers, customers, hardware models, and so on, perform the following steps:

- 1** In the drop-down list for (the Server, Customer, Model, and so on), select Contains, Equals, Begins With, or Ends With.
- 2** (Optional) Select the ellipsis button to open the Select Values window.
- 3** In the Select Values window, select a value in the Available or Selected pane and then use the directional buttons to include it in or exclude it from your search criteria.
- 4** Click **OK** to save your changes.
- 5** Click **Run** to run the report.



---

If data cannot be found to run the report, a “No records to display!” error displays.

---

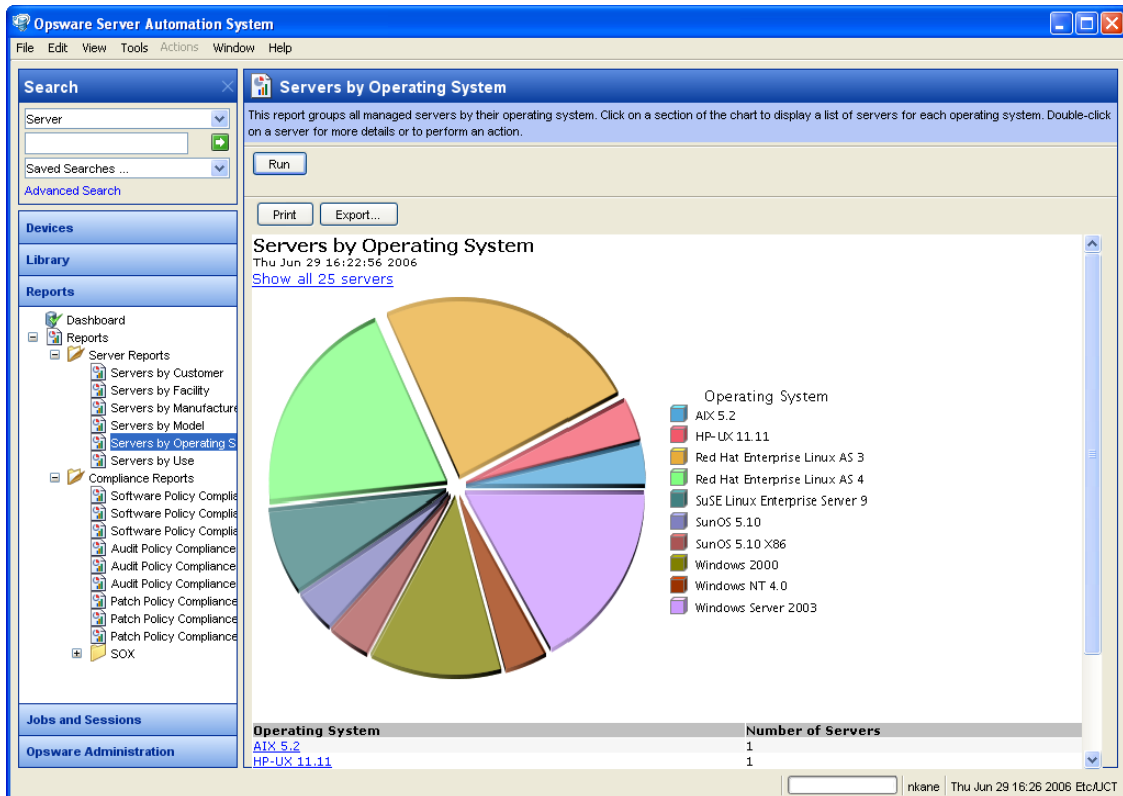
## Report Results

Report results initially appear in a graphical or list view. The graphical report is an overview of available data for this report displayed in a pie chart or in a bar graph. You can drill down for more detail in the chart or graph by clicking on any of the sections or bars. For example, you can drill down to individual servers that appear in a report and get detailed information about them.

## Graphical Report

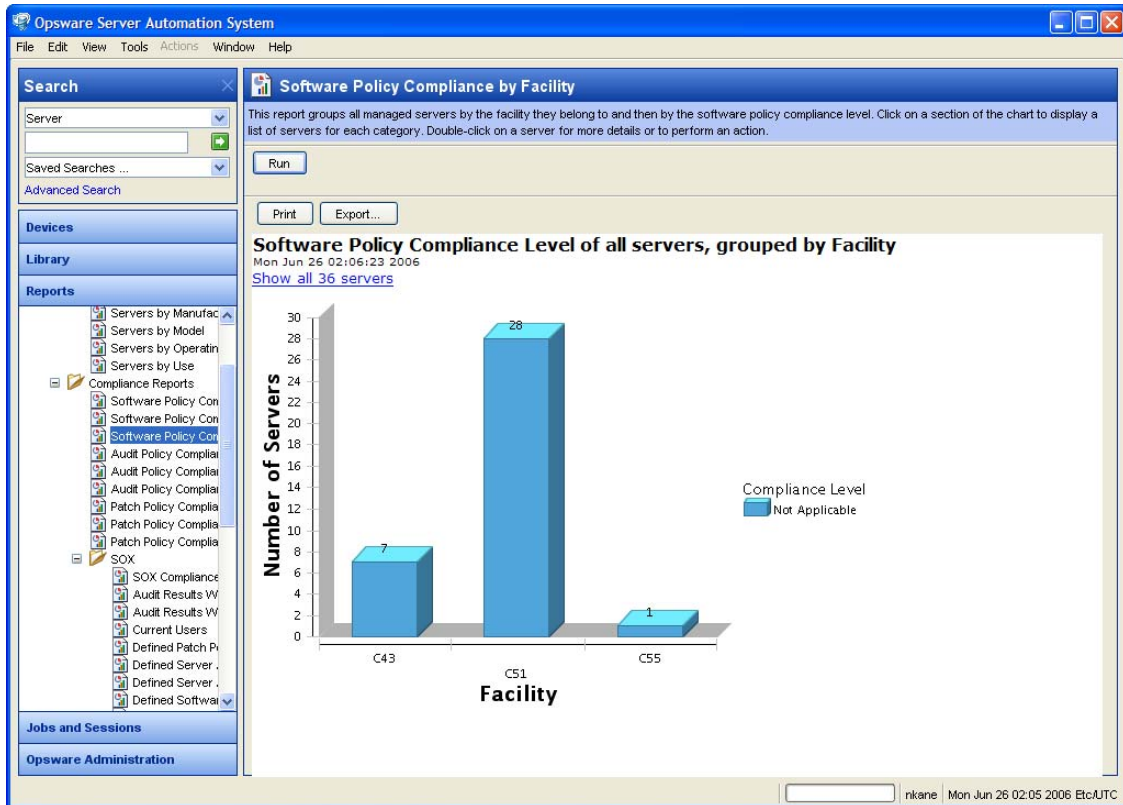
A graphical report is a pie chart or a bar graph. Click on a section of the chart or graph to drill down for more details or to perform an action. You can also click on the “Show all <number> servers” link to display a list of servers. See Figure 4-3 and Figure 4-4 for examples.

Figure 4-3: Pie Chart



To display the corresponding list view of a bar graph, click on the front part of the bar. Do not click on the top or shaded part of the bar.

Figure 4-4: Bar Graph



## List Report

A list report is a tabular display of information. Double-click on a row in the list, such as a server, audit, or policy, for more detail or to perform an action. See Figure 4-5 for an example.

Figure 4-5: List Report

The screenshot shows the Opware Server Automation System interface. The main window is titled "Software Policy Compliance by Customer". It features a search bar on the left and a navigation tree. The central area displays a report titled "Server Listing" with the following data:

Hostname	IP Address	Operating System	Customer	Compliance Level	Date Scanned
m208.dev.opsware.com	192.168.165.130	Red Hat Enterprise Linux AS 3	Not Assigned	Not Applicable	
erendirra.snv1.corp.opsware.com	192.168.9.224	Red Hat Linux 7.3	Not Assigned	Not Applicable	
BDay_LTXP.snv1.corp.opsware.com	192.168.8.186	Windows XP	Not Assigned	Not Applicable	
m282-rhel4-vm1.dev.opsware.com	192.168.202.68	Red Hat Enterprise Linux AS 4	Not Assigned	Not Applicable	
andy	192.168.165.132	Red Hat Enterprise Linux AS 3	Not Assigned	Not Applicable	
m279-rhel2-vm1.dev.opsware.com	192.168.202.36	Red Hat Enterprise Linux AS 2.1	Not Assigned	Not Applicable	
m282-nt-vm1	192.168.202.63	Windows NT 4.0	Not Assigned	Not Applicable	
loki.corp.opsware.com	192.168.9.128	Red Hat Linux 8.0	Not Assigned	Not Applicable	
m279-solaris10-vm1.dev.opsware.com	192.168.202.41	SunOS 5.10 X86	Not Assigned	Not Applicable	
fake_Satya0001.opsware.com	157.23.29.86	AIX 4.3	Not Assigned	Not Applicable	
m049.dev.opsware.com	192.168.194.58	AIX 5.2	Not Assigned	Not Applicable	
fake_Satya0001.opsware.com	158.190.233.168	Windows 2000	Not Assigned	Not Applicable	
fakeSatya0001.opsware.com	214.146.204.139	Windows 2000	Not Assigned	Not Applicable	
m047	192.168.194.55	HP-UX 11.11	Not Assigned	Not Applicable	



## Exporting a Report

You can export a report for use in other applications in your environment and attach a report for email distribution. Depending on the report format, you can export a report to your local file system in either .html, .pdf, or.xls file formats. You can export a graphical report to .html or .pdf only. You can export a list report to either .html, .pdf, or.xls file formats.



---

When you export a report in the SAS Client, the time that you will see marked on the exported report will be the time when the report was exported, not the time when the report was generated.

---

To export a report, perform the following steps:

- 1** From the report, click **Export** to open the Save window.
- 2** In the Save in field, enter a location that identifies where you want to save the file to, or select from the drop-down list.
- 3** Enter a file name.
- 4** Select the file type.
- 5** Click **Save**.

## Printing a Report

To print a report, perform the following steps:

- 1** From the report, click **Print** to open the Print window.
- 2** Use the default print options or modify them, and then click **OK**.



# Chapter 5: Patch Management for Windows

## IN THIS CHAPTER

This section discusses the following topics:

- Overview of Patch Management for Windows
- Roles for Windows Patch Management
- Patch Management Process
- Patch Properties
- Policy Management
- Patch Compliance
- Patch Administration for Windows
- Locales for Windows Patching
- 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 Opware SAS 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. This feature also supports patching on 64 bit for Windows 2003 operating systems and for 32 bit for Windows XP operating systems.

This section contains information about how to install Windows patches using patch policies and how to uninstall patches using a sequence of tasks. It also contains information about running patch compliance scans and generating patch policy compliance reports.

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 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.

## **Patch Management for Windows Features**

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 that identify servers that require patching
- Auditing abilities for tracking the deployment of important patches

These features enable you to browse patches by a certain operating system, schedule patch downloads and installations, set up email notifications, preview a patch installation, use policies and remediation to install patches, and export patch information to a reusable file format.

### **Types of Patch Browsing**

The SAS 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.

### **Scheduling and Notifications**

In Patch Management, you can separately schedule when you want patches imported from Microsoft (either automatically or on demand) into Opware 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 reboot preferences to adopt, override, postpone, or suppress the vendor's reboot options.

### **Patch Policies and Exceptions**

To provide flexibility in how you identify and distribute patches on managed servers or groups of servers, Patch Management allows you to create patch policies that define groups of patches that you need to install. By creating a patch policy and attaching it to a server or a group of servers, you can effectively manage which patches get installed where in your organization. If you want 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.

### **Patch Installation Preview**

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 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 that 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 installation, 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 for 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 supersede other patches or are superseded by other patches.

### **Patch Policy Remediation**

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 installation, you can also preview a patch uninstallation.

### **Exporting Patch Data**

To help you track the patch state of servers or groups of servers, 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.

## **Library**

The SAS Client Library provides flexibility in searching for and displaying Microsoft patches by operating system, severity level, release date, bulletin ID, and so on. See Figure 5-1. 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 Library. Use the column selector to control the columns of patch metadata data that you want to display.

Since the Library is integrated with Microsoft patch metadata, you can review vendor information (in real-time) in the Preview pane.

Figure 5-1: Windows Patches in the Opware SAS Client Library

The screenshot displays the Opware Server Automation System interface. The main window has a menu bar (File, Edit, View, Tools, Actions, Window, Help) and a search bar. The left sidebar shows a tree view of the library, with the 'Patches' folder expanded to show various operating systems. The 'Windows Server 2003' folder is selected and highlighted in blue. The right pane shows the details for the selected patch, Q890923, including its name, type, OS, size, and availability.

**Search**

Server [v]  
  
 Saved Searches ... [v]  
[Advanced Search](#)

**Windows Server 2003**

View: Properties [v]

Name	Type	Severity
Q893066	Windows H...	Critical
Q890923	Windows H...	Critical

**Q890923**

**General**

Name: Q890923  
 Type: Windows Hotfix  
 OS: Windows 2003  
 Size: 3.85 MB  
 Opware ID: 7800040  
 Availability: Available [v]  
 Title: Cumulative Security Upd  
 KB #: 890923  
 Bulletin: MS05-020

## Patch Management for Windows Prerequisites

The managed servers that will be patched have the following requirements:

- Either Microsoft Core XML Services (MSXML) 3.0 (or later) or Internet Explorer (IE) 6.0 (or later) must be installed on the managed servers. These versions of MSXML and IE support the Microsoft XML parser and related DLL files that are required for the native Microsoft Baseline Security Analyzer (MBSA) tool (mbsacli.exe). Opware SAS uses version 2.0 of the MBSA tool for patch management. (From Opware SAS 5.5 through 6.0, version 1.2.1 of MBSA was also used.) Vendor-recommended patches that are installed during the patch remediation process are based on MBSA 2.0.

- Windows Installer 3.1 must be installed on the managed servers. This installer is available at the following URL:

<http://support.microsoft.com/kb/893803/>

- On the managed servers, the Automatic Update service must be set to either Automatic or Manual. To set a Windows service, from the Windows Control Panel select Administration Tools ► Services. This service setting is required because Patch Management relies on the MBSA 2.0 scanning engine (mbsacli) to detect installed and recommended patches. If the Automatic Update service is disabled, mbsacli will not work properly and the patching process will not continue after reboot. In this situation, the Agent is unable to report a complete set of installed and recommended patches.
- For Windows 2000 managed servers, SP4 must be installed. Servers with earlier service packs are not supported by Patch Management.
- For Windows XP managed servers, SP2 must be installed.
- To use Patch Management on managed servers with Opware Agent versions earlier than 6.1, the language (locale) of the managed server must be either English, Japanese, or Korean. To set the language, on the managed server, open the Control Panel, open the Regional and Language Options window, select the Regional Options tab, and select an item from the drop-down list at the top.
- Specific versions of the Opware Agent are required to support the functions of Patch Management, as listed in Table 5-1.

Table 5-1: Opware Agent Requirements

PATCH FUNCTIONALITY	OPSWARE AGENT VERSION
Install Patch	4.5 or later



Table 5-1: Opsware Agent Requirements

PATCH FUNCTIONALITY	OPSWARE AGENT VERSION
Uninstall Patch	4.5 or later
Remediate	5.5 or later

### 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 the patches that 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 Opsware SAS Client or with a script.

Once every 24 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 then stores the data 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 SAS 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 SAS 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 SAS 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 SAS 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 Windows 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, reboot instructions, 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 5-2: 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 Update Rollup
Windows 2003	Windows Hotfix Windows OS Service Pack Update Rollup
Windows XP	Windows Hotfix Windows OS Service Pack Update Rollup

## Supporting Technologies for Patch Management

Patch Management uses patching utilities and technologies for each supported Windows operating system. Opware SAS uses these tools behind the scenes. This 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. It also detects the application files that 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**: Enables access to the Microsoft framework for patch updates.

See "Importing Windows Patch Utilities" on page 376.

## Windows Hotfixes

After a Microsoft Windows hotfix is imported into Opware 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 Opware 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 require that the server is rebooted, the reboot can sometimes be postponed until all hotfixes have been installed. The user performing the installation must first run qchain.exe before performing the reboot. This ensures 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 for Q815062.

If a Windows Service Pack or Security Rollup Package is being installed in the same hotfix chaining process, a reboot is required. This reboot 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 can be safely postponed until the end of the chained installation.

Opsware SAS will *not* change the installation 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 installation operation.

Opsware SAS runs qchain.exe on the managed server after the installation of each Windows hotfix and before any associated reboot. This guards 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.

### **Searching for Patches and Policies**

In the SAS Client, you can search for information about your operational environment by using the SAS Client Search feature. The Search feature enables you to search for patches, patch policies, servers, and so on. See “SAS Client Search” in the *Opsware® SAS User’s Guide: Server Automation*.

## Roles for Windows Patch Management

Opsware SAS provides support for rigorous change management by assigning the functions of patch management to several types of users in an organization. These users include 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 and identifies the vendor patches that 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 is able to assess the necessity of applying patches issued by vendors. A policy setter is also 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 to 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 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 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<sup>®</sup> SAS Administration Guide*.

---

## Patch Management Process

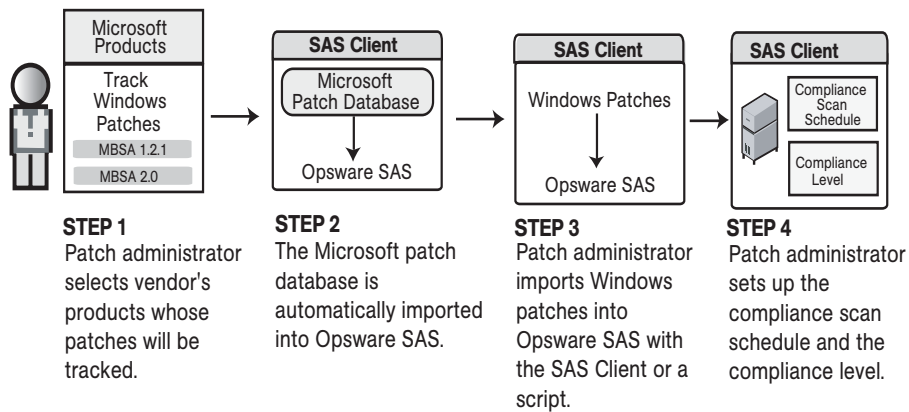
The Windows patching process consists of several key phases: setup, policy management, patch compliance, and deployment. Setup steps include getting the Microsoft database (patches and metadata) into Opsware SAS, identifying products you want to track patches for, and configuring patch compliance. Policy 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 groups of servers. Patch compliance steps include running compliance scans to determine whether a server is out of compliance, remediating policies, setting up installation options, and installing applicable patches. To deploy patches on demand, you

can import the required patches, test them, update policies, create new policies, mark them as available to use, specify install options, and install the required patches. Figure 5-2 and Figure 5-3 illustrate these phases and steps.

Figure 5-2: Windows Patching Process: Part A and Part B

## WINDOWS PATCHING PROCESS

### Part A: Set Up Patch Management



### Part B: Create and Attach Patch Policies to Servers


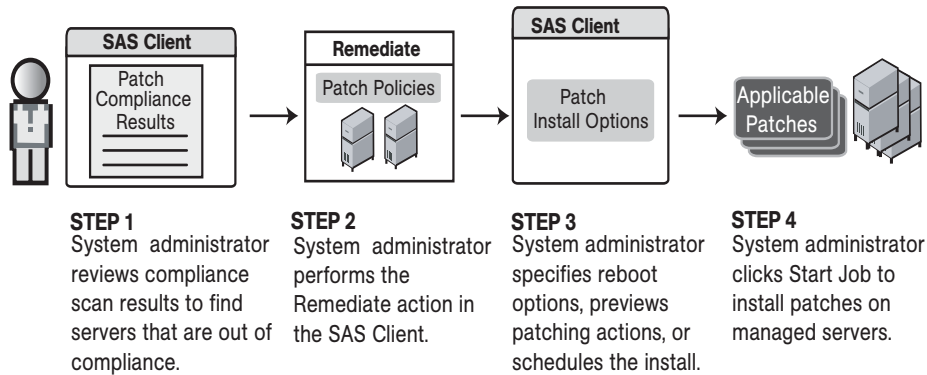




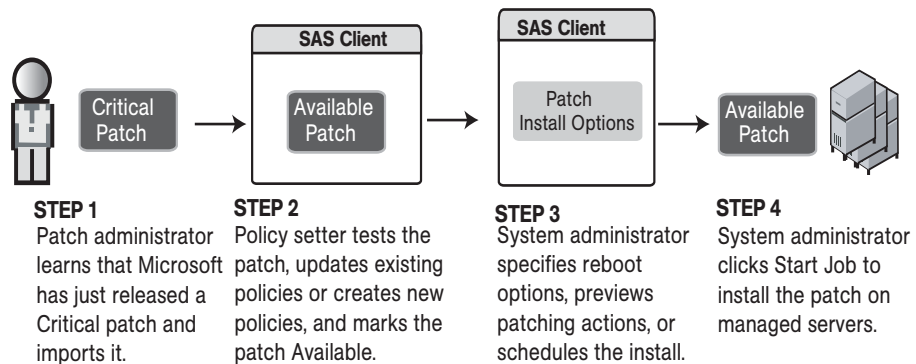
Figure 5-3: Windows Patching Process: Part C and Part D

## WINDOWS PATCHING PROCESS

### Part C: Install Patches By Remediating Policies



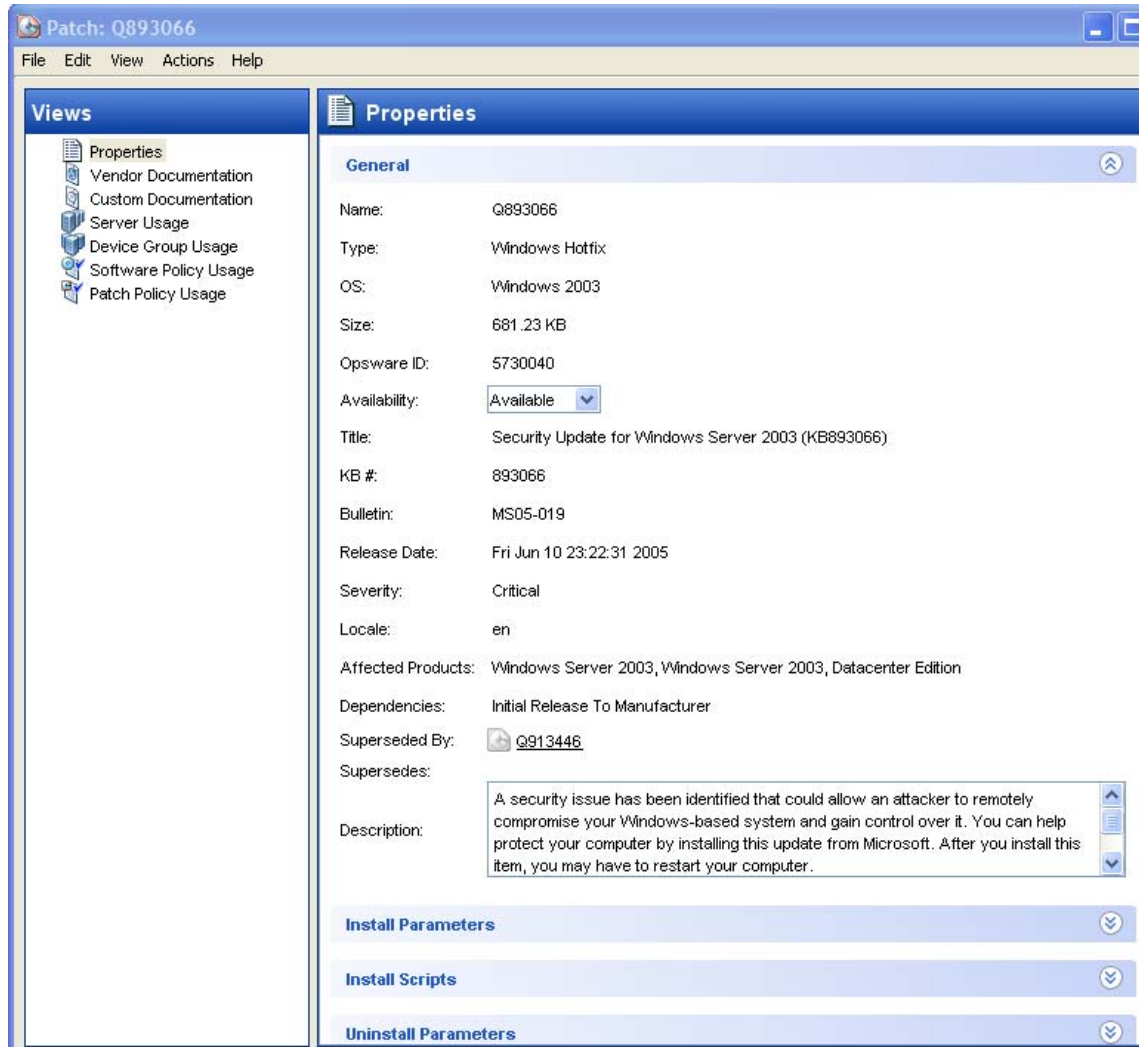
### Part D: Install Patches on Demand



## Patch Properties

Patch Management displays detailed information (properties) about a patch.

Figure 5-4: Windows Patch Properties



Patch properties include the following information:

- **Name:** The Microsoft name of the patch, such as QNumber, Windows 2000 Service Pack 4, and so on.
- **Type:** The type of patch, such as Windows Hotfix or Windows Update Rollup.
- **OS:** The Windows operating systems that are known to be affected by this patch.

- **Size:** The size of the patch file, in kilobytes (KB) or in megabytes (MB).
- **Opware ID:** The Opware SAS unique ID for the patch.
- **Availability:** The status of a patch within Opware 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 Opware SAS.
  - **Limited:** The patch has been imported into Opware SAS but cannot be installed. This is the default patch availability.
  - **Available:** The patch has been imported into Opware 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.
- **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 for this patch.
- **Release Date:** The date that Microsoft released this patch.
- **Severity** (Optional): The Microsoft severity rating for this patch, which can be one of the following:
  - **Critical:** A patch that if exploited could allow the propagation of an internet worm, without user action.
  - **Important:** A patch that if exploited could result in a compromise of the confidentiality, integrity, or availability of user data, or of the integrity or availability of processing resources.
  - **Moderate:** A patch that if exploited could result in minimal impact. Exploitability is mitigated to a significant degree by certain factors, such as default configuration, auditing, or difficulty of exploitation.
  - **Low:** A patch that is difficult to exploit or if exploited, could result in minimal impact.
- **Locale:** The locale this patch applies to.

- **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** (Optional): A list of patches that this patch is superseded by. This relationship does not apply to MBSA 1.2.1 patches.
- **Supersedes** (Optional): A list of patches that this patch supersedes. 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 supersede or are 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 replaced 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 remediating 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. Patch Management does not analyze this information for MBSA 1.2.1 patches.



Patch Management does not detect whether two patches are mutually exclusive, which is when 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 SAS Client displays information about Microsoft Windows patches that have been imported into Opware SAS.

To view information about a patch, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► 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 Windows 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 Opware 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** From the Navigation pane, select Library ► By Type ► 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 From the Navigation pane, select Library ► By Type ► 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 Views pane, select Custom Documentation.
- 5 From the **Actions** menu, select **Import Custom Documentation** or click **Import**.
- 6 In the Import Custom Documentation window, locate a text file and specify encoding.
- 7 Click **Import**.

### Deleting 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 delete custom documentation for a patch, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► 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 Views pane, select Custom Documentation.
- 5 From the **Actions** menu, select **Delete Custom Documentation**.
- 6 In the Delete Custom Documentation window, click **Delete**.

### Finding Vendor-Recommended Windows Patches

To find the patches that Microsoft recommends for a particular server (based on MBSA 2.0), perform the following steps:

- 1** From the Navigation pane, select Devices ► Servers ► All Managed Servers.
- 2** From the View drop-down list, select Patches.
- 3** From the Content pane, select a server that is running Opware 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 from the drop-down list. This displays the types of patches for the selected server.

### Finding Servers That Have a Windows Patch Installed

To find the servers that have a particular patch installed, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► 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 Server Usage.
- 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 Windows Patch Installed

To find the servers that do not have a particular patch installed, perform the following steps:

- 1** From the Navigation pane, select Library ► 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 Server Usage.
- 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 Opware 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 SAS Client or with a script. For information about the script, see "Automatically Importing Windows Patches" on page 344.

To import a patch with the SAS Client, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Package Repository.
- 2 Expand the Package Repository 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**.

### Automatically Importing Windows Patches

The `populate-opware-update-library` shell script downloads the Microsoft Patch Database and patches from the Microsoft site. The script also imports the database and patches into Opware SAS. (To be imported, a patch must be in the Microsoft Patch



Database that has been imported into the Software Repository.) Optionally, the script sets the initial status (Available or Limited) of newly imported patches. The script can also filter the patches imported according to operating system (such as Windows 2003). The functionality of the script is also available in the SAS Client, as described in “Importing the Microsoft Patch Database” on page 373.

To run the `populate-opsware-update-library` script, you need to log onto the Software Repository server as `root`. Typically, you schedule the script to run periodically as a `cron` job on the Software Repository server. To end users of the SAS Client, the patches imported with the script appear to have been automatically imported. Do not run concurrent instances of the script.

The `populate-opsware-update-library` script is in the following directory:

```
/opt/opsware/mm_wordbot/util/
```

Table 5-3 describes the script's options.

Table 5-3: Options of `populate-opsware-update-library`

OPTION	DESCRIPTION
<code>--spin hostname-or-IP</code>	Hostname or IP address of Data Access Engine (spin) host. Default value: spin
<code>--theword hostname-or-IP</code>	Hostname or IP address of Software Repository (theword) host. Default value: theword
<code>--cert_path file-path</code>	File specification of cert file to be used for Spin connection. Default value: <code>/var/opt/opsware/crypto/wordbot/wordbot.srv</code>
<code>--ca_path file-path</code>	File specification of CA file to be used for Spin connection. Default value: <code>/var/opt/opsware/crypto/wordbot/opsware-ca.crt</code>
<code>--verbose</code>	Display copious output, including patches skipped during the upload.

Table 5-3: Options of `populate-opsware-update-library` (continued)

OPTION	DESCRIPTION
<code>--no_nt4</code>	Do not process NT4 patches.
<code>--no_w2k</code>	Do not process W2K patches.
<code>--no_w2k3</code>	Do not process W2K3 patches.
<code>--no_w2k3x64</code>	Do not process Windows 2003 (64 bit) patches.
<code>--no_xp</code>	Do not process Windows XP (32 bit) patches.
<code>--use_proxy_url url</code>	When downloading binaries, connect via this proxy URL.
<code>--proxy_userid userid</code>	Basic-auth userid to provide to proxy server.
<code>--proxy_passwd passwd</code>	Basic-auth passwd to provide to proxy server.
<code>--set_available</code>	Set availability status to Available when uploading patches. The <code>--set_available</code> and <code>--set_limited</code> options cannot be specified at the same time.
<code>--set_limited</code>	Set availability status to Limited when uploading patches.
<code>--no_hotfixes</code>	Do not upload hotfixes.
<code>--no_servicepacks</code>	Do not upload servicepacks.
<code>--no_updaterollups</code>	Do not upload updaterollups.
<code>--no_wsusscan_upload</code>	Do not upload the MBSA 2.0 patch database.
<code>--wsusscan_url_override url</code>	Download the MBSA 2.0 patch database from this URL.
<code>--update_all</code>	Refresh the patches already uploaded into Opware SAS.

Table 5-3: Options of `populate-opsware-update-library` (continued)

OPTION	DESCRIPTION
<code>--download_only path</code>	Download files from the vendor's web site to the specified path (directory), but do not upload them into Opsware SAS. The files are downloaded into the <code>platform_ver/locale</code> subdirectory beneath the specified path.
<code>--upload_from_update_root path</code>	Upload files from the specified path (directory), not from the vendor's web site. The script looks for patches in the <code>platform_ver/locale</code> subdirectory beneath the specified path. If it cannot find the patch in the that subdirectory, the script looks for the patch in the specified path. If a patch is not found, the script skips the patch and does not upload it. This option is ignored if <code>--download_only</code> is also specified.
<code>--help</code>	Display the syntax of this script.

### Exporting a Windows Patch

To export a patch from Opsware SAS to the local file system, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► 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 Windows Patch Information

You can export information about patches installed on a server and patches recommended by the vendor. You can also export information from patches recommended by the vendor along with model information on the selected server (such as patch policies or patch policy exceptions). The following information is exported into a .csv file:

- **Server Name:** The name of the managed server.
- **OS:** The operating system of the server.
- **Service Pack:** The service pack level of the server being reported, such as Service Pack 0, Service Pack 1, and so on.
- **KB#:** The Microsoft Knowledge Base Article number for the patch.
- **Bulletin:** 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:** A brief description of the purpose of the patch.
- **Time Queried:** The last software registration by the Agent.
- **Time Installed:** The time that the patch was installed.
- **Type:** The patch type.
- **Compliance Level:** An integer that represents the compliance level.
- **Compliance:** Text that displays when you place your cursor over the Compliance column in the Patch Preview pane.
- **Exception Type:** The type of exception, such as Always Install or Never Install.
- **Exception Reason:** A description that explains the purpose of the exception.



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** From the Navigation pane, select **Devices** ► **All Managed Servers**.
- 2** From the Content pane, select one or more managed servers.
- 3** From the Show drop-down list, select an option.
- 4** From the **Actions** menu, select **Export Patch Info to CSV**.
- 5** In the Export to CSV window, navigate to a folder and enter the file name.
- 6** 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.
- 7** 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

When you delete a patch, it is removed from Opsware SAS, but it is not uninstalled 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.

---

To delete a patch, perform the following steps:

- 1** From the Navigation pane, select **Library** ► **By Type** ► **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**.

- 5 In the Delete Patches windows, click **Delete**.

## Policy Management

In Patch Management, patch policies and patch policy exceptions enable you to customize patch distribution in your environment. Policies and exceptions define the Windows patches that 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 deviate from the patch policies and exceptions and perform ad hoc patch installs, then you need to remediate. The remediation process ensures that the applicable patches get installed on servers.

### Patch Policy

A patch policy is a group of patches that you want to install on Opware SAS managed servers. All patches in a patch policy must apply to the same Windows operating system.

A patch policy provides broad flexibility for distributing 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) and installs 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 groups of servers. 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 Remediate window allows you to select an individual patch policy to remediate. You do not have to remediate 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. The patch compliance process compares patch policies with corresponding patch policy exceptions.

Patch Management supports the following types of patch policies:

- **User-defined patch policy:** This allows an Opware SAS user to specify the patches that 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). This enables the policy setter to apply only those patches that their environment needs.

- **Vendor-recommended patch policy:** Membership of patches is defined by MBSA recommendations 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 is associated with it. See the *Opware<sup>®</sup> SAS User's Guide: Server Automation*.
- A patch policy is always public.

- A patch policy can be attached to zero or more servers or public device groups.
- More than one patch policy can be attached to a server or public device group.
- 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 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 a group of servers). You can do this by deselecting or adding individual patches to 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. The patch compliance results 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 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 device 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 group of servers), you control the patches that should be installed or not installed on a server. A precedence hierarchy in Patch Management delineates how a patch policy or a patch policy exception is applied to a patch installation. This hierarchy is based on whether the patch policy or patch policy exception is attached at the server or device group level.

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 device group.
- Patch policy exceptions that are attached to a public device group take precedence over patch policies that are attached to a public device group.
- If a server is in multiple public device groups, a Never Installed patch policy exception type always take precedence over an Always Installed patch policyexception type for the same patch.

## Remediation Process

To ensure patch compliance, Patch Management identifies vulnerable managed servers and simultaneously deploys patches to many servers when a remediation process is performed. The remediation process examines and applies an entire patch policy (including multiple policies) to the managed servers that it is attached to. A policy must be attached to a server or a group of servers before you can remediate the policy with that server or group.



---

The remediation process requires that the selected managed server is running Opware Agent 5.5 and a Windows 2000 Service Pack 3 (or higher) operating system or a Windows 2003 operating system. You cannot use the remediation 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 Opware Agent 5.5. Use the Install Patch window to install patches on servers that are running these operating systems or Opware Agents 4.5 or earlier.

---

As a best practice, each time you review the latest Microsoft patch releases and subsequently update a patch policy (by adding new patches to a policy), you should perform remediation. In these situations, a remediation process provides demand forecasting information. This allows you to determine how patch policy changes will impact servers that this policy is attached to.


If the remediation process discovers any (applicable) missing patches, these patches will be installed on the servers.

If a patch was installed as part of a patch policy, the remediation process will not uninstall it. However, if a patch was installed as part of a software policy and it is no longer in the software policy, the remediation process will uninstall it.

After Opware SAS determines the packages that need to be installed to complete the remediation process, remediation uses a set of standard system utilities to complete the operation. See “Supporting Technologies for Patch Management” on page 332.

To help you optimally manage the remediation conditions, Patch Management allows you to specify remediate options and pre and post actions, and set up ticket IDs and email notifications that alert you about the status of the remediate process. The Remediate window guides you through setting up these conditions.

Figure 5-5: Remediate Window



## Remediating Patch Policies

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 remediation process for a group of servers, patches will only be remediated if any server in the 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 Remediate 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 remediate a patch policy, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► 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 Server Usage.
- 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 **Remediate**. The first step of the Remediate window appears: Servers and Device Groups.

For instructions on each step, see the following sections:

- Setting Remediate Options
- Setting Reboot Options for Remediation
- Specifying Pre and Post Install Scripts for Remediation
- Scheduling a Patch Installation for Remediation
- Setting Up Email Notifications for Remediation
- Previewing a Remediation

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 remediation job.

After you launch the job, you cannot change its parameters, even if the job is scheduled to run at a later time.

If you leave the Remediate window 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 Remediate Options

You can specify the following remediate policy option:

“Do not interrupt the remediate process even when an error occurs with one of the policies.”

To set this option, perform the following steps:

- 1** From the Remediate window, click **Next** to advance to the Remediate Options step.
- 2** Select one of the following Staged Install Options:
  - Continuous:** Run all phases as an uninterrupted operation.
  - Staged:** Allow download and installation to be scheduled separately.
- 3** Select the Error Options check box if you want the remediation 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 Remediate window.

## Setting Reboot Options for Remediation

To minimize the downtime that server reboots can cause, you can control when servers reboot during a patch installation.

You can specify the reboot options in the following two places in the SAS Client:

- Install Parameters tab of the patch properties window
- Pre & Post Actions step of the Remediate window



When you are selecting reboot options in the Remediate window, Opware, Inc. recommends that you use Microsoft’s reboot recommendations, which is the “Reboot servers as specified by patch properties” option. 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. 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 Opware control).

---

The following options in the Remediate window determine whether the servers are rebooted after the patch is installed. These options apply only to the job launched by the Remediate 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.
- **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.)
- **Hold all server reboots until after all packages are installed and/or uninstalled:** 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 Remediate window, click **Next** to advance to the Pre & Post Actions step.
- 2** Select one of the Reboot Options.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Remediate window.

### Specifying Pre and Post Install Scripts for Remediation

For each patch remediation, you can specify a command or script to run before or after remediation. 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 remediation 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 Remediate 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 Remediate Options step.
- **Pre-Install:** A script that runs before patches are installed on the managed server.
- **Post-Install:** 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 Remediate 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 SAS Web Client. 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 System, 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 Remediate window

## Scheduling a Patch Installation for Remediation

You can schedule when you want patches installed and when you want patches downloaded.

To schedule a patch installation, perform the following steps:

- 1** From the Remediate 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 installation phase. If you selected Staged in the Remediate 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 installation immediately.
  - **Run Task At:** This enables you to specify a date and time that you want the download or installation performed.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Remediate window.

## Setting Up Email Notifications for Remediation

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 Remediate window, click **Next** to advance to the Notifications step.
- 2** To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the installation phase. If you selected Staged in the Remediate Options step, the notification status for the download phase is also displayed.
- 4** Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5** Click **Next** to go to the next step or click **Cancel** to close the Remediate window.



If you previously selected Staged in the Remediate Options step, the Notifications pane displays notification options for both the download and installation phases.

---



## Previewing a Remediation

The remediate preview process provides an up-to-date report about the patch state of servers. The remediate preview is an optional step that lets you see the patches that 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.

In the Preview, the servers, device groups, and patches that are listed in the Summary Step window will be submitted to remediation 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 shows patches and their associated servers (regardless of any patch policy and server group membership changes that may have occurred). If you preview a remediation, this same list of servers, device 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 Remediate 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 device groups, you must click **Preview** again for results that include your changes.



---

The remediation preview does not report on the behavior of the server as though the patches have been applied.

---

To preview a remediation, perform the following steps:

- 1** From the Remediate window, click **Next** to advance to the Summary Review step.
- 2** Verify the information displayed for the Servers, Device 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 Remediate 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 installation, 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 installation.
- **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 the *Opsware® SAS User's Guide: Server Automation* for more information on browsing job logs.

- 7** Click **Stop Job** to prevent the job from running or click **Close** to close the Remediate window. You can stop a job only if it is scheduled.

## Verifying Patch Policy Compliance

To determine whether a managed server complies with patch policies and exceptions, perform the following steps:

- 1** From the Navigation pane, select Devices ► All Managed Servers.
- 2** From the Content pane, select Patches from the View drop-down list.
- 3** Examine the Patch column at the top of the pane. This column indicates the overall patch compliance for a server.

- 4 Select a server at the top of the Content pane and examine the Compliance column at the bottom. This column indicates the compliance status of each individual patch for the selected 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 From the Navigation pane, select Library ► By Type ► 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 Opware SAS but does not remove or uninstall patches from managed servers. You cannot delete a patch policy if it is attached to servers or groups of servers. You must first detach the policy from the servers or groups of servers before removing it from Opware SAS.

To delete a patch policy from Opware SAS, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► 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 remediated.

To add a patch to a patch policy, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patch Policies.
- 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** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Select a specific Windows operating system and view the list of Windows patches.
- 3** From the Content pane, select a patch.
- 4** From the View drop-down list, select Patch Policies.
- 5** From the Show drop-down list, select Policies with Patch Added.
- 6** Select a patch. From the **Actions** menu, select **Remove from Patch Policy**.
- 7** 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 group of servers). You must perform this action before you remediate a policy with a server (or group of servers).

To attach the policy, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► 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 Server Usage (or Device Group Usage).
- 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** From the Navigation pane, select Library ► By Type ► 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 Server Usage (or Device Group Usage).
- 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 remediation process. (The Install Patch and Uninstall Patch actions ignore patch policy exceptions.) A patch policy exception overrides the policy. You can specify an exception for a particular patch and server (or group of servers), but not for a patch policy.

To set a patch policy exception, perform the following steps:





- 1** From the Navigation pane, select Devices ► All Managed Servers.
- 2** Select a server.
- 3** From the Content pane, select a server.
- 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:
  - **Never Install**: The patch should not be installed on the server, even if the patch is in the policy.
  - **Always Install**: The patch should be installed on the server even if the patch is not 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. The Patches with Exceptions option must be selected.
- 9** Click **OK**.

## Finding an Existing Patch Policy Exception

You can search for managed servers that 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** From the Navigation pane, select Devices ► 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 group of servers/patch association.
  -  A never install exception on a patch/server association.
  -  A never install exception inherited to a server from a group of servers/patch association.

### Copying a Patch Policy Exception

To copy an exception between servers or groups of servers, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand the Patches and select a specific Windows operating system.
- 3 From the Content pane, select a patch.
- 4 From the View drop-down list, select Server Usage (or Device Group Usage).
- 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 device 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 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand the Patches and select a specific Windows 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 device groups to determine whether all patches in a policy and a policy exception were installed successfully. 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 show you the servers that are in compliance (have all required patches installed) and the servers that are out of compliance (do not have all required patches installed).

You should 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) Opware SAS, a compliance scan is required because the Opware model has been changed and the compliance information must now be recalculated. Patch Management indicates these types of conditions by displaying Scan Needed. In this case, instead of waiting for the scan schedule to iterate, you can start compliance scan on one or more servers.

### Ways to Start a Patch Compliance Scan

You can start a patch compliance scan in the following ways:



- Immediately, by selecting servers or groups and then selecting a menu item. See “Starting a Patch Compliance Scan Immediately” on page 369.
- Periodically, by setting up a schedule. See “Scheduling a Patch Compliance Scan” on page 375. By default, the scans are not scheduled.
- As a result of another task. Opsware SAS performs a patch compliance scan on a managed server at the end of the tasks described in the following sections:
  - “Installing a Windows Patch” on page 383
  - “Uninstalling a Windows Patch” on page 393
  - “Remediating Patch Policies” on page 355

### Starting a Patch Compliance Scan Immediately

To start a scan on selected servers, perform the following steps:


- 1** From the Navigation pane, select Devices .
- 2** Select an entry from either the Managed Servers or Device Groups list.
- 3** Right-click and select **Scan ► Patch Compliance**.

### Refreshing the Compliance Status of Selected Servers

You can refresh the compliance status of all Windows servers by selecting **View ► Refresh**. However, this global refresh operation can take a long time when scanning a large number of servers. To save time, you can refresh the compliance status of selected servers by performing the following steps:

- 1** From the Navigation pane, select Devices.
- 2** Drill down to the servers you want to check.
- 3** In the Contents pane, select one or more servers
- 4** Right-click and select **Refresh Server Status**.
- 5** Note any changed values in the Patch column.





### Viewing Scan Failure Details

If the scan operation fails, you cannot determine whether a server is in compliance. A scan failure is indicated by the  icon. To find out why a patch compliance scan failed, perform the following steps:

- 1** From the Navigation pane, select Devices.
- 2** Drill down to the server you want to check.
- 3** In the Contents pane, select a server.
- 4** Right-click and select **Scan ► Show Patch Compliance Scan Failure Details**.
- 5** In the Patch Compliance Scan Failure Details window, select a server and examine the detailed error message that appears in the lower part of the window.

### Patch Compliance Icons

Patch Management displays the following icons:

-  The server is compliant for all patches. Patches in policies attached to the server are all installed on that server.
-  The server is partially compliant for patches. An exception has been set for these patches.
-  The server is not compliant for patches. Patches in policies attached to the server are not installed on that server.
-  The scan operation failed. Patch Management is unable to check the compliance of the server.

### 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 the compliance icons that 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 Compliance Reports

To help troubleshoot problems, you can run and examine several patch compliance reports that are based on Sarbanes-Oxley (SOX) standards. These reports identify whether all patches in a policy and a policy exception were installed successfully on managed servers. The Reports feature of the SAS Client provides the following patch compliance reports.

- **Defined Patch Policies:** Lists patch policies by name, customer, and operating system, and includes the total number of patch policies.

- **Patch Policy Compliance (All Servers):** Groups all managed servers by their patch policy compliance level to show compliant and non-compliant servers.
- **Patch Policy Compliance by Customer:** Lists all servers by the customer they belong to and then by the patch policy compliance level.
- **Patch Policy Compliance by Facility:** Groups all managed servers by the facility they belong to and then by the patch software policy compliance level.
- **Servers in Compliance With Their Patch Policies:** Lists all managed servers that are in compliance with all of their attached patch policies.
- **Servers Not in Compliance With Their Patch Policies:** Lists all managed servers that are not in compliance with their attached patch policies.
- **Servers With Attached Patch Policies:** Lists all managed servers that have one or more patch policies attached, and includes the total number of servers with attached patch policies.
- **Servers Without Attached Patch Policies:** Lists all managed servers that do not have any patch policies attached, and includes the total number of servers without any attached patch policies.



See the *Opware® SAS User's Guide: Server Automation* for information about how to run, export, and print these reports.

---

## 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 SAS Client.
- You can import the Microsoft patch database (on demand) by using a command-line script or the SAS Client.
- You can track (and import) only patches that apply to certain Microsoft products or particular locales.

- You can import and export Windows patch utilities.
- 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 Patch Availability

You can set the default patch availability with either the SAS Client or a command-line script. The default used by the script overrides the default set by the SAS Client. For information about the script, see “Automatically Importing Windows Patches” on page 344.

To set the default value for the Availability of a newly imported patch, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Settings.
- 3** For the Patch 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 Administration Guide* for an explanation of these permissions.

### Importing the Microsoft Patch Database

You can import the Microsoft Patch Database by using a command-line script or the SAS Client. For information about the script, see “Automatically Importing Windows Patches” on page 344.

To import the database with the SAS Client, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Settings.
- 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 the default 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 `wsusscan.cab` (MBSA 2.0) file and click **Import**. This file 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 Opware SAS to specific Windows products. After performing this operation, the next time the Microsoft Patch Database is imported, any new patches listed by Opware SAS are limited to the products that you select. Patches that were previously listed by Opware SAS are still tracked. You can also track patches for all MBSA 2.0 products.

To limit the patches tracked to specific Windows operating systems, run the command-line script that automatically imports patches. For more information about the script, see “Automatically Importing Windows Patches” on page 344.

To select the Windows products to track for patching, perform the following steps:

- 1** From the Navigation pane, select Opware Administration.
- 2** Select Patch Settings.
- 3** Select the Windows MBSA tab.
- 4** Click **Edit**.
- 5** In the Edit Patch Properties window, use the include and exclude arrows to select the products whose patches you want to track and then click **Select**.

## Scheduling a Patch Compliance Scan

To schedule a patch compliance scan on all Windows managed servers, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Compliance Settings.
- 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** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Compliance Settings.
- 3** Select one of the following compliance levels: Policy and Exception, Policy Only, or Customized.

## Importing Windows Patch Utilities

You can import the following Windows utilities from your local file system into Opsware SAS:

- mbsacli.exe
- parsembacli20.exe
- qchain.exe
- WindowsUpdateAgent20-x86.exe
- WindowsUpdateAgent20-x64.exe
- wusscan.dll

Initially, these files are imported into Opsware SAS during the installation of the core. To import a Windows patch utility, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Settings.
- 3** In the Patch Utilities section, select a utility and then click **Import Utility Update**.

## Exporting Windows Utility Files

You can export the following Windows patch utilities from Opsware SAS to your local file system:

- mbsacli.exe
- parsembacli20.exe
- qchain.exe



- WindowsUpdateAgent20-x86.exe
- WindowsUpdateAgent20-x64.exe
- wusscan.dll

To export a Windows patch utility, perform the following steps:

- 1** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Settings.
- 3** In the Patch Utilities section, select one or more utilities and then click **Export Utility**.

### 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** From the Navigation pane, select Opsware Administration.
- 2** Select Patch Compliance Settings.
- 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**.

### Locales for Windows Patching

The locale of a patch identifies the language of the Windows servers that should receive the patch. A patch with the same name might be available for different locales. For example, a patch named Q123456 might be available for servers running the English and Japanese versions of Windows. Although they have the same name, the patches installed on the English and Japanese servers are different binaries.

Patch Management supports multiple locales in the same Opsware multimaster mesh. To install a patch on Windows servers with different locales, you specify the patch by name. During the installation (or policy remediation), Opsware SAS matches the locale of the patch with the locale of each managed server. You do not need to repeat the installation for each locale.

## Supported Locales

Patch Management supports Windows patches of the following locales:

- English (en)
- Japanese (ja)
- Korean (ko)

## Overview of Locale Configuration Tasks

By default, Patch Management supports only the English locale. To set up Patch Management for non-English locales, step through the instructions in the following sections:

- “Configuring the Opsware Core for Non-English Locales” on page 378
- “Selecting the Locales of Patches to Import” on page 379
- “End User Requirements for Non-English Locales” on page 380

## Configuring the Opsware Core for Non-English Locales

This task requires `root` access to core servers and a restart the OCC core component. To configure the core for non-English locales, perform the following steps on each core server running the OCC component:

- 1** Log onto the server as `root`.
- 2** With a text editor, in `/etc/opt/opsware/occ/psrvr.properties`, change the line for `pref.user.locales` to the following:  
`pref.user.localesAllowed=en;ja;ko`
- 3** Restart the OCC component of the core:  
`/etc/init.d/opsware-sas restart occ.server`
- 4** In a text editor, open the following file:  
`/opt/opsware/occclient/jnlp.tmp1`

- 5** For the Japanese language, In the `<resources>` section of the `jnlp.tmp1` file, add the following XML element:  

```
<property name="com.opsware.ngui.font.japanese" value="Arial Unicode MS"/>
```
- 6** For the Korean language, In the `<resources>` section of the `jnlp.tmp1` file, add the following XML element:  

```
<property name="com.opsware.ngui.font.korean" value="Arial Unicode MS"/>
```
- 7** In the `/opt/opsware/occclient` directory, if the following files exist, delete them:  

```
$HOST_ja.jnlp  
$IP_ja.jnlp  
$HOST_ko.jnlp  
$IP_ko.jnlp
```
- 8** Follow the steps in “Selecting the Locales of Patches to Import” on page 379.

### Selecting the Locales of Patches to Import

Follow the instructions in “Configuring the Opsware Core for Non-English Locales” on page 378 before performing the steps in this section.

This operation selects the locales of the Windows patches to import into Opsware SAS. The selections take effect the next time patches are imported into Opsware SAS. After the patches have been imported, they can be installed on managed servers. If you remove locales from the list with this operation, patches with those locales that have already been imported are not removed from Opsware SAS.

To select the locales of the Windows patches to import into Opsware SAS, perform the following steps:

- 1** In the SAS Client, from the Navigation pane, select Opsware Administration.
- 2** Select Patch Settings.
- 3** On the Windows MBSA tab, select Patch Locales.
- 4** Click **Edit**.
- 5** In the Edit Patch Locales window, use the include and exclude arrows to select the locales whose patches you want to import. If you want to select a locale that is not listed in “Supported Locales” on page 378, contact Opsware Inc. Support.

- 6** Click **Select**.
- 7** Follow the instructions in “End User Requirements for Non-English Locales” on page 380.

### **End User Requirements for Non-English Locales**

To view non-English fonts in the SAS Client, end users must perform the following steps:

- 1** The end user verifies that the Windows desktop running the SAS Client uses the Arial Unicode MS font.
- 2** After the Opsware Administrator performs the steps in “Configuring the Opsware Core for Non-English Locales” on page 378, the end user logs onto the SAS Web Client and goes to the My Profile page,
- 3** On the My Profile page, the end user updates the Locale field on the User Identification tab. For example, if the Opsware Administrator configured the core for Japanese, then the end user sets the Locale field to Japanese.

### **Patch Installation**

Patch Management provides the following two phases in the patch installation process:

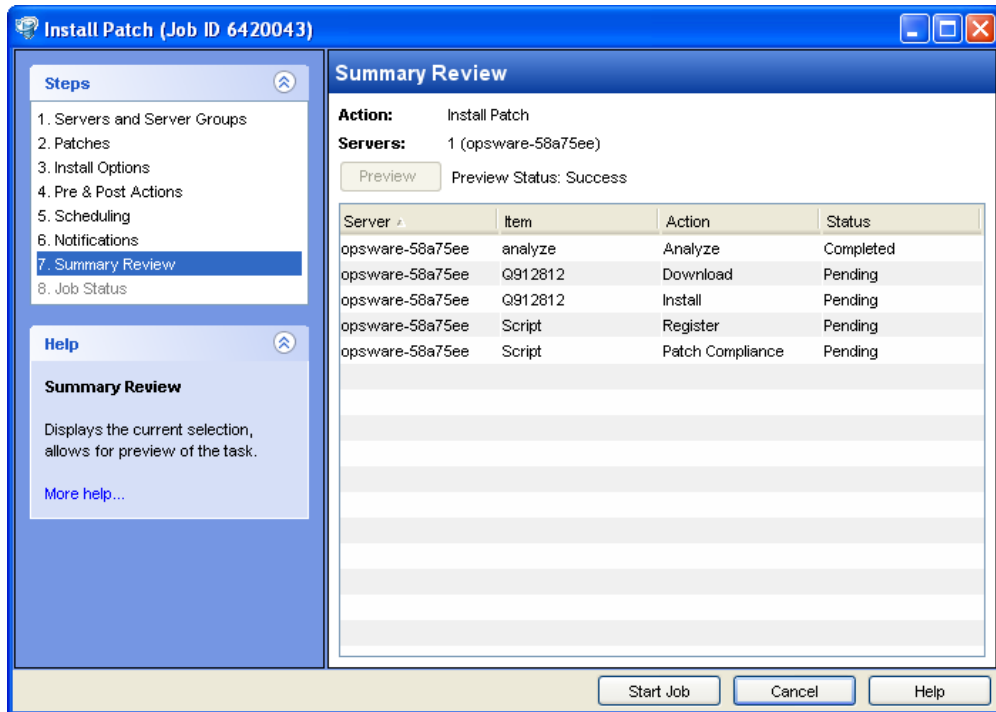
- **Download Phase:** This is when the patch is downloaded from Opsware SAS to the managed server. This phase is commonly referred to as the staging phase.
- **Installation Phase:** This is when the patch is installed on the managed server. This phase is commonly referred to as the deployment phase.

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 these default command-line arguments.

To help you optimally manage Windows patch installation, Patch Management allows you to manage server reboot options, specify 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 window guides you through setting up these conditions.

Figure 5-6: Install Patch Window



## Installation Flags

You can specify installation flags that are applied whenever a Windows patch is installed. However, Opware 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 by Opware SAS. See “Setting Windows Install Options” on page 384 for information about how to specify commands and flags.



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. This prevents the Patch Management feature from passing in the `-z` or `-q` or `-z -q` flag. By default, Opware 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 Opware SAS uses.

Table 5-4: Default Installation Flags

WINDOWS PATCH TYPE	FLAGS
Windows Hotfix	<code>-q -z</code>
Windows Security Rollup Package (treated identically to a Hotfix by the Patch Management feature)	<code>-q -z</code>
Windows OS Service Pack	<code>-u -n -o -q -z</code>

## Application Patches

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, Patch Management does not automatically filter out servers that do not have the corresponding application installed. Although Patch Management does not prevent you from doing so, you should not attempt to apply application patches to servers that do not have the necessary applications installed. If a patch is for an application that is not installed on the server, the patch will not be applied and an error message will display, such as "There was an error with package <name of the package>".

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.

### **Service Packs, Update Rollups, and Hotfixes**

When you try to install a Service Pack, Update Rollup, or a Hotfix, there is a known delay when a confirmation dialog displays. Since the Opsware Agent is installing or uninstalling the patch, it cannot respond to the confirmation dialog. The Agent will time out an installation or uninstallation process if you do not click **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. For Service Packs and Update Rollups, the Agent will time out if 60 minutes have lapsed and you have not clicked **OK** in the confirmation dialog.

To prevent this from happening, patch install and uninstall commands should have arguments that invoke silent mode installs and uninstalls. By default, the `-q` flag is set.

### **Installing a Windows 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 Administration Guide*.

---

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** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand the Patches and select a specific Windows operating system.
- 3** From the Content pane, select a patch.
- 4** From the View drop-down list, select Servers (or Device Groups).

- 5 From the Show drop-down list, select Servers without Patch Installed (or Device 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 window appears: Servers and Device Groups. For instructions on each step, see the following sections:

- Setting Windows Install Options
- Setting Reboot Options for a Windows Patch Installation
- Specifying Install Scripts for a Windows Patch Installation
- Scheduling a Windows Patch Installation
- Setting Up Email Notifications for a Windows Patch Installation
- Previewing a Windows Patch Installation
- Viewing Job Progress of a Windows Patch Installation

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 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.

See “Remediating Patch Policies” on page 355 for another method of installing a patch.

## Setting Windows 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 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 installation 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 window.

### Setting Reboot Options for a Windows Patch Installation

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 window, Opsware, Inc. recommends that you use Microsoft's reboot recommendations, which is the "Reboot servers as specified by patch properties" option. 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. Failure to do this can result in MBSA incorrectly reporting the patches that 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 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.)
- **Hold all server reboots until after all packages are installed and/or uninstalled:** 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 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 window.

### Specifying Install Scripts for a Windows Patch Installation

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-Install:** 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 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 SAS Web Client. 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 System, 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 window.

## Scheduling a Windows Patch Installation

Since the two phases of patching can be decoupled, you can schedule that you want patches installed independently of when patches are downloaded.

To schedule a patch installation, perform the following steps:

- 1** From the Install Patch window, click **Next** to advance to the Scheduling step.  
By default, the Scheduling step displays only the scheduling options for the installation 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 installation or download performed.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Install Patch 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 Windows Patch Installation

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 window, click **Next** to advance to the Notifications step.
- 2** To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.

- 3 To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the installation phase.
- 4 Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.



---

If you previously selected Staged in the Install Options step, the Notifications pane displays notification options for both the download and installation phases.

---

### Previewing a Windows 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 the patches that will be installed on managed servers and the type of server reboots that are required. This preview process verifies whether the servers that 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 supersede other patches or are superseded 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 Opware 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 remediate preview will display a “Will Not Install” error message to indicate this discrepancy. The Install Patch 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 window, click **Next** to advance to the Summary Review step.
- 2** Verify the information displayed for the Servers, Device 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 window without launching the installation.

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 Windows Patch Installation

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 window, click **Next** to advance to the Job Progress step. This will start the installation 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 installation, 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 uninstallation.
- **Install & Reboot:** When a patch is installed, the server is also 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 the *Opware<sup>®</sup> SAS User's Guide: Server Automation* for more information about browsing job logs.
- 3** Click **Stop Job** to prevent the job from running or click **Close** to close the Install Patch 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 Opware 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 window guides you through setting up these conditions.

Figure 5-7: Uninstall Patch Window



## Uninstallation Flags

You can specify uninstallation flags that are applied whenever a Windows patch is uninstalled. However, Opware 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 by Opware SAS.



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, Opware 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 5-5: Default Uninstallation Flags

WINDOWS PATCH TYPES	FLAGS
Windows Hotfix	-q -z
Security Rollup Package	-q -z
Windows OS Service Pack	Not uninstallable

## Uninstalling a Windows Patch

To remove a patch from a managed server, perform the following steps:

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand the Patches and select a specific Windows 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 (Servers) in the Uninstall Patch window appears.

For instructions on each step, see the following sections:

- Setting Uninstall Options
- Setting Uninstall OptionsSetting Reboot Options for a Windows Patch Uninstallation
- Specifying Install Scripts for a Windows Patch Uninstallation
- Scheduling a Windows Patch Uninstallation
- Setting Up Email Notifications for a Windows Patch Uninstallation
- Viewing Job Progress of a Patch Uninstallation

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 uninstallation 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 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 uninstallation.

To set these options, perform the following steps:

- 1 From the Uninstall Patch 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, Opware 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 window.

### Setting Reboot Options for a Windows Patch Uninstallation

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 window, Opware, Inc. recommends that you use Microsoft's reboot recommendation. This is the "Reboot servers as specified by patch properties" option. 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. Failure to do this can result in MBSA incorrectly reporting which patches are installed on the server until the next reboot occurs (outside of Opware 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 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.)
- **Hold all server reboots until after all packages are installed and/or uninstalled:** 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 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 window.

## Specifying Install Scripts for a Windows Patch Uninstallation

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 uninstallation:

- **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 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 Opware SAS with the SAS Web Client. To specify the script, click **Select**.

An Ad-Hoc script runs only for this operation and is not saved in Opware 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 Windows Patch Uninstallation

You can remove a patch from a server immediately, or at a later date and time.



To schedule a patch uninstallation, perform the following steps:

- 1** From the Uninstall Patch 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 uninstallation in the Summary Review step.
  - **Run Task At:** This enables you to specify a later date and time that you want the uninstallation performed.
- 3** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

## Setting Up Email Notifications for a Windows Patch Uninstallation

You can set up email notifications to alert users when the patch uninstallation operation completes successfully or with errors.

To set up email notifications, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Notifications step.
- 2** To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the uninstallation phase.
- 4** Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

## Previewing a Windows Patch Uninstallation

The uninstallation preview process provides an up-to-date report about the patch state of servers. The uninstallation preview is an optional step that lets you see the patches that will be removed from managed servers. This preview process verifies whether the servers you selected for the patch uninstallation have that patch installed (based on the MBSA).



---

The uninstallation preview process does not report or simulate the behavior of a system with patches removed from the server.

---

To preview a patch uninstallation, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Summary Review step.
- 2** Verify the information displayed for the Servers, Device 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 window without launching the uninstallation.

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 Uninstallation

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 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:** Opware SAS examines the patches needed for the uninstallation, 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 uninstallation.
  - **Uninstall & Reboot:** When a patch is installed, the server is also 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 the *Opware<sup>®</sup> SAS User's Guide: Server Automation* for more information on browsing job logs.
  - 3** Click **Stop Job** to prevent the job from running or click **Close** to close the Uninstall Patch window.





# Chapter 6: 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 Properties
- Software Policies
- Patch Administration for Unix
- Patch Installation
- Patch Uninstallation

## Overview of Patch Management for Unix

The Patch Management for Unix feature enables you to identify, install, and remove patches, and maintain a high level of security across managed servers in your organization. With the SAS Client user interface, you can identify and install patches that protect against security vulnerabilities for the AIX, HP-UX, and Solaris operating systems.

This section contains information about how to install and uninstall Unix patches using software policies. It also contains information about generating patch policy compliance reports.

Opware 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 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 Opware SAS. It leverages the Opware SAS server automation features. Opware 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 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. Opware SAS also allows you to schedule patch activity, so that patching occurs during off-peak hours.

### **Patch Management for Unix Features**

Opware 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 that identify servers that require patching
- Auditing abilities that enable security personnel to track the deployment of important patches

These features enable you to browse patches by a certain operating system, schedule patch downloads and installations, set up email notifications, preview a patch installation, use software policies and remediation to install and uninstall patches, and export patch information to a reusable file format.

### **Types of Patch Browsing**

The Opware SAS Client interface organizes Unix patches by operating systems and displays detailed vendor security information about each patch. You can browse patches by patch type, availability, platform version, and so on. You can also browse all patches that are installed on a server, and view and edit patch metadata.

### **Scheduling and Notifications**

In Patch Management, you can separately schedule when you want patches uploaded into Opware 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 reboot preferences to adopt, override, postpone, or suppress the vendor's reboot options.

### **Software Policies**

Software policies enable you to customize patch distribution in your environment. They define the Unix patches that should be installed or not installed on certain managed servers. See "Software Management" on page 437 for more information about creating software policies to install Unix patches.

### **Patch Installation Preview**

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 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 that 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. The preview process provides an up-to-date report of the patch state of servers.

### **Software Policy Remediation**

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 installation, you can also preview a patch uninstall. See "Software Management" on page 437 for more information about remediating software policies.

### **Exporting Patch Data**

To help you track the patch state of servers or groups of servers, 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.

### **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 a server is initially provisioned with Opsware SAS, the same data is immediately recorded.

When a new patch is issued, you can use Opsware SAS to immediately identify the servers that require patching. The Opsware SAS Client provides a software repository where you upload patches and other software. Users access this software from the Opsware SAS Client 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, reboot instructions, and how to handle error codes from the pre-install and post-install scripts.

## Library

The SAS Client Library provides flexibility in searching for and displaying Unix patches by name, type of patch, operating system, relationship to other packages, and so on. See Figure 6-1. The number in parenthesis is the total number of patches (for that operating system version) that were uploaded from the Unix web site. Use the column selector to control which columns of patch metadata data to display.

Figure 6-1: Unix Patches in the Opware SAS Client Library



## Search Feature

In the SAS Client, you can search for any information about your operational environment that is available in Opsware SAS using the SAS Client Search feature. The Search feature enables you to search for patches, software policies, servers, and so on. See “SAS Client Search” in the *Opsware® SAS User's Guide: Server Automation*.

## Patch Management Roles for Unix

Opsware SAS provides support for rigorous change management by assigning the functions of patch management to the patch administrator and the system administrator:

- The patch administrator (often referred to as the security administrator) has the authority to upload, test, and edit patch options.
- The system administrator applies the patches (that have been approved for use) uniformly, according to the options that the patch administrator specifies.



Only the patch administrator should have the Patches permission, which gives access to advanced features. To obtain these permissions, contact your Opsware administrator. See the Permissions Reference appendix in the *Opsware® SAS Administration 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 to 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 as available in the Opware SAS Client, 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, installing the patch, and verifying that the patches are installed successfully.

## **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 Opware SAS.

### **Supported Unix Versions and Patch Types**

The Patch Management feature supports all of the operating system versions that Opware 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 Opware SAS manages are therefore not viewable through the Patch Management feature interfaces. New Linux packages and updates should be managed and applied through the software policy. See the *Opware® SAS Policy Setter's Guide*, section "RPM Deployment" for information about importing and installing RPMs using a software policy.



The following table shows the Unix versions and the patch types that the Patch Management feature supports.

Table 6-1: 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 5.2	AIX Update Fileset APARs
AIX 5.3	AIX Update Fileset APARs
HP-UX 11.00	HP-UX Patch Fileset HP-UX Patch Product
HP-UX 11.11	HP-UX Patch Fileset HP-UX Patch Product
HP-UX 11.23	HP-UX Patch Fileset HP-UX Patch Product
Solaris 6	Solaris Patch Solaris Patch Cluster
Solaris 7	Solaris Patch Solaris Patch Cluster
Solaris 8	Solaris Patch Solaris Patch Cluster
Solaris 9	Solaris Patch Solaris Patch Cluster

Table 6-1: Supported Unix Versions and Patch Types (continued)

UNIX VERSIONS	PATCH TYPES
Solaris 10	Solaris Patch Solaris Patch Cluster

### Underlying Technologies for Patch Management on Unix

Although the utilities vary, Opware SAS enables you to perform patching tasks by using a single interface. Opware SAS models the way it treats patches by the way the underlying utility treats patches. 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. Opware 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 6-2: Supporting Technologies for Patch Management on Unix

SOLARIS	AIX	HP-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
Showrev lists installed Solaris patches	Instfix lists installed APARS	Swremove removes a depot
Pkgadd installs Solaris packages		

Table 6-2: Supporting Technologies for Patch Management on Unix (continued)

SOLARIS	AIX	HP-UX
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, Opware 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, Opware SAS still associates the earlier version of the fileset with the APAR.

When uploading an LPP, Opware 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 install all LPPs that APAR specifies, you must make certain to upload all of the specified LPPs into the Patch Management feature.

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 Opware 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.



Since update filesets can be included in folders, global read permissions are required to view and edit AIX update filesets. See "Software Management Setup" in the *Opware® SAS Policy Setter's Guide* for information about how to use folders.

---

## 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, Opware 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. If you installed the patch cluster with the Patch Management feature, you can uninstall individual patches in the cluster. You cannot uninstall a patch cluster.

## 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 Opware 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.

### **Patch Uploads for Unix**

Before a Unix patch can be installed on a managed server with Opsware SAS, the patch must be uploaded into the SAS Client Library. Uploading patches is the responsibility of the patch administrator. See the *Opsware® SAS Administration Guide* and the *Opsware® SAS Policy Setter's Guide* for information about how to upload Unix patches and the importing software process.

### **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 on.)

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.

## Patch Properties

Patch Management displays detailed information (properties) about a patch.

Figure 6-2: Unix Patch Properties



Patch properties include the following information:

- **Name:** The Unix name for the patch.
- **Type:** The type of Unix patch. Table 6-1 identifies these patch types.
- **OS:** The Unix operating systems that are known to be affected by this patch.
- **Size:** The size of the patch file, in kilobytes (KB) or in megabytes (MB). Size is not shown for AIX APARs.
- **Opware ID:** The Opware SAS unique ID for the patch.
- **Availability:** The status of a patch within Opware SAS, which can be one of the following:

- **Limited:** The patch has been imported into Opware SAS but cannot be installed. This is the default patch availability.
- **Available:** The patch has been imported into Opware 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.
- **Containing** (Optional): Depending on the selected patch type, this is the relationship to other packages. For example, for AIX update filesets, this field displays Containing LPPS/APARS.
- **Description:** A brief description of the Solaris patch cluster.

### Viewing Unix Patches

The SAS Client displays information about Unix patches that have been imported into Opware SAS.

To view information about a patch, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand Patches and select a specific Unix operating system.
- 3** (Optional) Use the column selector to sort the patches according to Name, Type, Availability, and Description.
- 4** In the Content pane, open a patch to view its properties in the Patch window.

### Editing Unix 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.

The Availability property indicates the status of the patch in Opware SAS.

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** From the Navigation pane, select Library ► By Type ► Patches.

- 2** Expand Patches and select a specific Unix 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.

### **Finding Servers That Have a Unix Patch Installed**

To find out which servers have a particular patch installed, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.
- 2** Expand Patches and select a specific Unix 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 Server Usage.
- 5** From the Show drop-down list for the selected patch, select Servers with Patch Installed.

### **Finding Servers That Do Not Have a Unix Patch Installed**

To find out which servers do not have a particular patch installed, perform the following steps:

- 1** From the Navigation pane, select Library and then select Patches.
- 2** Expand Patches and select a specific Unix 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 Server Usage.
- 5** From the Show drop-down list, select Servers without Patch Installed.

### **Exporting a Patch**

To export a patch from Opware SAS to the local file system, perform the following steps:

- 1** From the Navigation pane, select Library ► By Type ► Patches.



- 2 Expand Patches and select a specific Unix 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**.

### 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.



Do not delete all of the patches from Opsware SAS. If you do so accidentally, contact your Opsware, Inc. support representative for assistance in uploading all of the patches back into Opsware SAS.

- 1 From the Navigation pane, select Library ► By Type ► Patches.
- 2 Expand Patches and select a specific Unix 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**.
- 5 In the Delete Patches windows, click **Delete**.

### Software Policies

In Patch Management for Unix, software policies enable you to customize patch distribution in your environment. Software policies define which Unix patches should be installed or not installed on certain managed servers.

If you use software policies and you also perform ad hoc patch installs, you must run the remediate process to install all applicable patches on servers. See “Software Management” on page 437 for more information about creating and remediating software policies to install Unix patches.

## Patch Compliance Reports

To troubleshoot and resolve patch compliance problems, you can run and examine several patch compliance reports by using the Reports feature in the SAS Client. The following patch compliance reports identify whether all patches in a software policy were installed successfully on managed servers in your environment.

### **Patch Policy Compliance (All Servers)**

This report groups all managed servers by their patch policy compliance level to show compliant and non-compliant servers.

### **Patch Policy Compliance by Customer**

This report lists all servers by the customer they belong to and then by the patch policy compliance level.

### **Patch Policy Compliance by Facility**

This report groups all managed servers by the facility they belong to and then by the patch software policy compliance level.



---

See the *Opware® SAS User's Guide: Server Automation* for information about how to run, export, and print these reports.

---

## Patch Administration for Unix

You can customize patch administration for Unix to best support your environment by setting the availability flag.

### **Setting the Default Patch Availability**

You can set the default patch availability with the SAS Client. The default used by the script overrides the default set by the SAS Client. See the *Opware® SAS Administration Guide* for information about the script.

To set the default value for the Availability of a newly imported patch, perform the following steps:

- 1** From the Navigation pane, select Opware 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<sup>®</sup> SAS Administration Guide*.

## Patch Installation

The patch installation process consists of the following two phases:

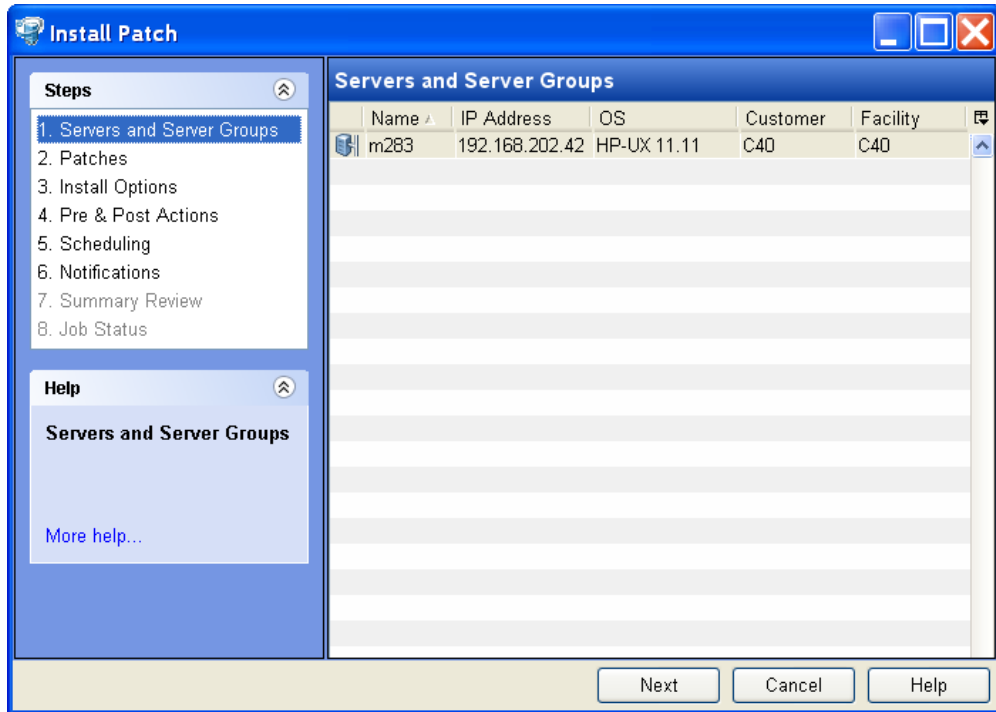
- **Download Phase:** This is when the patch is downloaded from Opsware SAS to the managed server. This phase is commonly referred to as the staging phase.
- **Installation Phase:** This is when the patch is installed on the managed server. This phase is commonly referred to as the deployment phase.

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 that installs the patch. The Opsware Agent runs this command on the managed server. You can override the default command-line arguments that you want to perform the installation.

To optimally manage Unix patch installations, 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 window guides you through setting up these conditions.

Figure 6-3: Install Patch Window



## Installation Flags

You can specify installation flags that are applied whenever a Unix patch is installed. However, Opware 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 Opware SAS. See "Setting Unix Install Options" on page 423 for information about how to specify commands.

The following table lists the default installation flags that Opware SAS uses.

Table 6-3: Default Installation Flags

UNIX PATCH TYPE	FLAGS
AIX	-a -Q -g -X -w
HP-UX	None

### Application Patches

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, Patch Management does not automatically filter out servers that do not have the corresponding application installed. Although Patch Management does not prevent you from doing so, you should not attempt to apply application patches to servers that do not have the necessary applications installed. If a patch is for an application that is not installed on the server, the patch will not be applied and an error message will display, such as “There was an error with package <name of the package>”.

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.

### Installing a Unix Patch

Before a patch can be installed on a managed server, it must be imported into Opware 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 Opware administrator. See the *Opware® SAS Administration Guide*.

---

You can perform the installation by explicitly selecting patches and servers.

To install a patch on a managed server, perform the following steps:

- 1** From the Navigation pane, select Library and then select Patches.
- 2** Expand the Patches and select a specific Unix 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 window appears: Servers and Server Groups.

For instructions on each step, see the following sections:

- Setting Unix Install Options
- Setting Reboot Options for a Unix Patch Installation
- Specifying Install Scripts for a Unix Patch Installation
- Scheduling a Unix Patch Installation
- Setting Up Email Notifications for a Unix Patch Installation
- Previewing a Unix Patch Installation
- Viewing Job Progress of a Unix Patch Installation

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 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 Unix 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 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 installation 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 that is displayed.
- 5** Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.

### Setting Reboot Options for a Unix Patch Installation

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 window, Opsware, Inc. recommends that you use the Unix reboot recommendations, which is the “Reboot servers as specified by patch properties” option. If you cannot use the Unix reboot setting, select the single reboot option, which is the “Do not reboot servers until all patches are installed” option.

---

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 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 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 window.

### Specifying Install Scripts for a Unix Patch Installation

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-Install:** 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 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 SAS Web Client. To specify the script, click **Select**.
- 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 window.

## Scheduling a Unix 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 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 installation or download performed.

**3** Click **Next** to go to the next step or click **Cancel** to close the Install Patch 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 Unix Patch Installation

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 window, click **Next** to advance to the Notifications step.

**2** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the installation phase.

**3** Enter a Ticket ID to be associated with a Job in the Ticket ID field.

- 4 Click **Next** to go to the next step or click **Cancel** to close the Install Patch window.



---

If you previously selected Staged in the Install Options step, the Notifications pane displays notification options for both the download and installation phases.

---

### Previewing a Unix 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. 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 information, such as patches that require certain Unix 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.



---

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 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 window without launching the installation.

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 Unix Patch Installation

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 window, click **Next** to advance to the Job Progress step. This will start the installation 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:** Opware SAS examines the patches needed for the installation, checks the managed servers for the most recent patches installed, and determines other actions that it must perform.
- **Download:** The patch is downloaded from Opware 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 uninstallation.
- **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 the *Opware® SAS User's Guide: Server Automation* for more information about browsing job logs.

- 3** Click **Stop Job** to prevent the job from running or click **Close** to close the Install Patch window.

## Patch Uninstallation


Patch Management provides granular control over how and under what conditions Unix patches are uninstalled (removed) from managed servers. To minimize problems, you can only uninstall one patch at a time. You cannot use Opware 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 window guides you through setting up these conditions.

Figure 6-4: Uninstall Patch Window



## Uninstallation Flags

You can specify uninstallation flags that are applied whenever a Unix patch is uninstalled. However, Opware 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 by Opware SAS.

The following table lists the default uninstallation flags that Opware SAS uses.

Table 6-4: Default Uninstallation Flags

OPERATING SYSTEM/PATCH TYPES	FLAGS
AIX	-u -g -X
AIX Reject Options	-r -g -X
HP-UX	None

### Uninstalling a Unix Patch

To remove a patch from a managed server, perform the following steps:

- 1** From the Navigation pane, select Library and then select Patches.
- 2** Expand the Patches and select a specific Unix 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 window appears: Servers.

For instructions on each step, see the following sections:

- Setting Reboot Options for a Unix Patch Uninstallation
- Specifying Pre and Post Install Scripts for a Unix Patch Uninstallation
- Scheduling a Unix Patch Uninstallation
- Setting Up Email Notifications for a Unix Patch Uninstallation
- Viewing Job Progress of a Patch Uninstallation

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 uninstallation 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 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 uninstallation.

To set these options, perform the following steps:

- 1** From the Uninstall Patch 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 that is displayed.
- 4** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

### Setting Reboot Options for a Unix Patch Uninstallation

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 window, Opsware, Inc. recommends that you use the Unix reboot recommendations, which is the "Reboot servers as specified by patch properties" option in the window. If it is not possible to use the Unix reboot setting, select the single reboot option, which is the "Do not reboot servers until all patches are installed" option in the window.

---

The following options determine whether the servers are rebooted after the patch is uninstalled. These options apply only to the job launched by the Uninstall Patch 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 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 window.

### **Specifying Pre and Post Install Scripts for a Unix Patch Uninstallation**

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 uninstallation:

- **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 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 SAS Web Client. To specify the script, click **Select**.

- 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 Unix 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 uninstallation, perform the following steps:

- 1** From the Uninstall Patch 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 uninstallation in the Summary Review step.
- **Run Task At:** This enables you to specify a later date and time that you want the uninstallation performed.

- 3** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

## Setting Up Email Notifications for a Unix Patch Uninstallation

You can set up email notifications to alert users when the patch uninstallation operation completes successfully or with errors.

To set up email notifications, perform the following steps:

- 1** From the Uninstall Patch window, click **Next** to advance to the Notifications step.
- 2** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon. By default, the Notification step displays only the notification status for the uninstallation phase.
- 3** Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 4** Click **Next** to go to the next step or click **Cancel** to close the Uninstall Patch window.

## Previewing a Unix Patch Uninstallation

The uninstallation preview process provides an up-to-date report about the patch state of servers. The uninstallation 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 uninstallation have that patch installed.



The uninstallation preview process does not report or simulate the behavior of a system with patches removed from the server.

---

To preview a patch uninstallation, perform the following steps:

- 1** From the Uninstall Patch 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 window without launching the uninstallation.

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 Uninstallation

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 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 uninstallation, checks the managed servers for the most recent patches installed, and determines other actions that 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 uninstallation.
  - **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 the *Opsware<sup>®</sup> SAS User's Guide: Server Automation* for more information on browsing job logs.
- 3** Click **Stop Job** to prevent the job from running or click **Close** to close the Uninstall Patch window.



# Chapter 7: Software Management

## IN THIS CHAPTER

This section contains the following topics:

- Overview of Software Installation
- Software Installation Process
- Ways to Install Software in Opware SAS
- Installing or Uninstalling Software on a Server
- Installing Software Using a Software Policy
- Uninstalling Software Using a Software Policy
- Overview of Software Template
- Overview of Running ISM Controls
- Software Policy Compliance
- Software Policy Reports

## Overview of Software Installation

Opware SAS automates the time-consuming process of installing and uninstalling software on managed servers. In the SAS Client, using software policies, you can install software and configure applications across a large number of managed servers with a minimum amount of downtime. In a software policy you can specify the software resources such as packages, patches, scripts, server objects to be installed, and the application configurations to be applied to the managed servers.

When you apply a software policy to a server, the software resources in the software policy are installed and the application configurations are applied on the managed server in a single step. In a software policy, you can also set the installation order among the software resources in a software policy, and set custom attributes and ISM controls for servers. See the *Opware<sup>®</sup> SAS Policy Setter's Guide* for information about creating software policies.

To install software in Opware SAS, you must attach a software policy to servers or groups of servers. When you remediate a server or group of servers, the software resources and application configurations specified in the attached policy are automatically installed and applied respectively. During remediation, you can separate the analysis, download, and installation stages of software deployment, specify the reboot operations, schedule the download and installation stages, set email notifications, and associate a ticket ID with the job. The remediation process allows you preview the installation of software before you actually install the software on servers. See “Overview of Software Policies Remediation” on page 453 in this chapter for more information.

In the SAS Client, you can also install or uninstall software directly on a managed server without using a software policy. See “Installing or Uninstalling Software on a Server” on page 441 in this chapter for more information

You can uninstall any software that you installed by using the SAS Client using a software policy. To uninstall a software, you must detach a software policy from a server and then remediate the server against that software policy. See “Detaching a Software Policy from a Server” on page 460 in this chapter for more information.

The Software Management feature also enables you to run software compliance scans to determine the compliance status of managed servers with respect to a software policy and then remediate non-compliant servers. See “Software Policy Compliance” on page 467 in this chapter for more information.

The Reporting feature in Opware SAS allows you to generate reports that provide summaries of the software policy compliance across servers. After you generate reports, you can print the reports, export the reports to .html and .xls, and perform actions on the results. See “Software Policy Reports” on page 468 in this chapter for more information.

This section contains information about how to install software using a software policy. It also contains information about running software compliance scans and generating software policy compliance reports. See *Opware® SAS Policy Setter's Guide* for information about uploading packages, and creating and managing software policies.

## Software Installation Process

The software installation process consists of installing software directly on a managed server as shown in Figure 7-1 or by attaching software policies to managed servers and then remediating the servers against those software policies as shown in Figure 7-2. This

phase includes tasks such as running software compliance scans to determine the compliance status of servers to remediate non-compliant servers, and generating software compliance reports across servers.

Figure 7-1: Software Management Process- Install Software

## SOFTWARE MANAGEMENT PROCESS (INSTALL SOFTWARE)

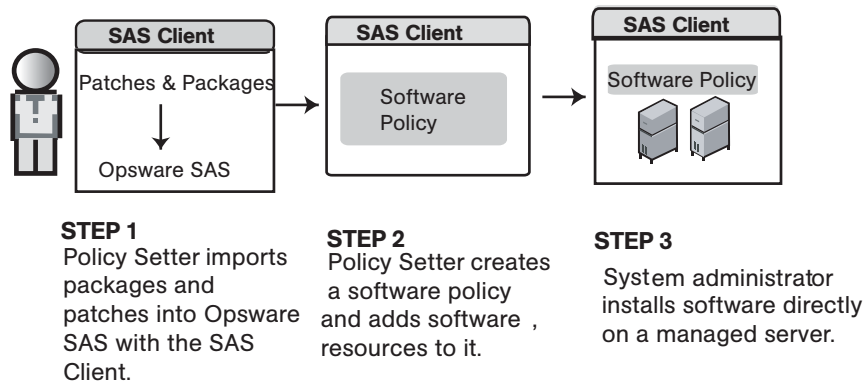
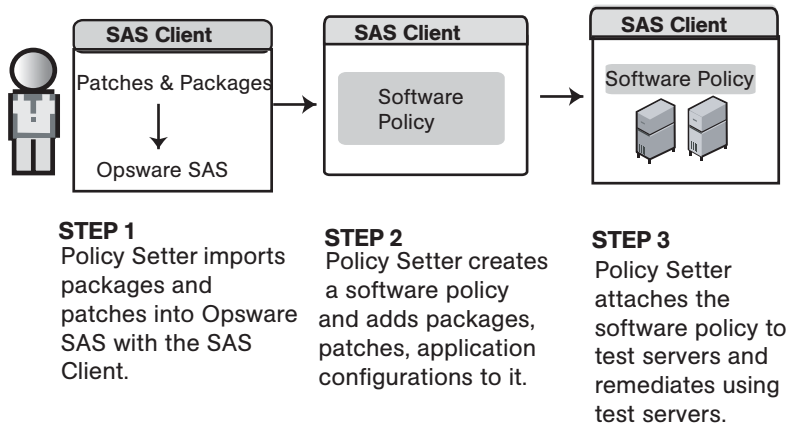


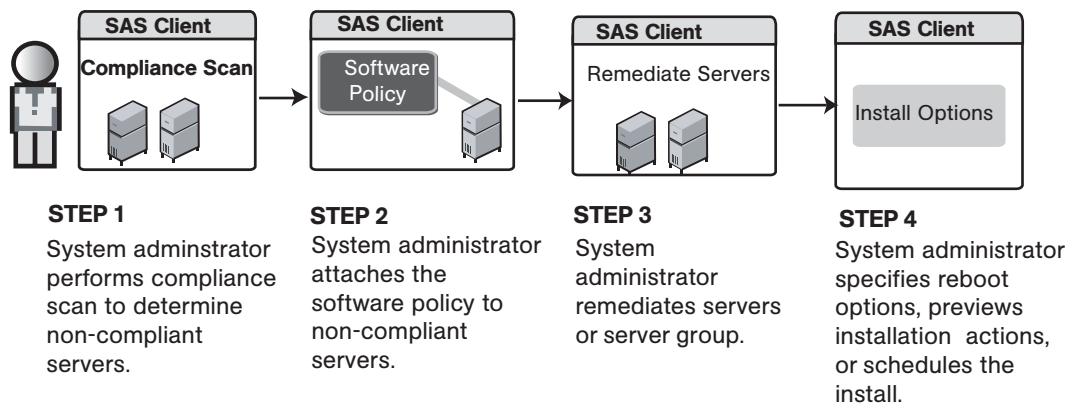
Figure 7-2: Software Management Process - Remediate

## SOFTWARE MANAGEMENT PROCESS (REMEDiate)

### Part A: Set Up Software Policies



### Part B: Attach Software Policies to Servers and Remediate



## Ways to Install Software in Opware SAS

Opware SAS provides several ways to install software and configure applications. In the SAS Client, you can perform the following tasks:



- Install software directly on a managed server. See “Installing or Uninstalling Software on a Server” on page 441 in this chapter for more information.
- Use a software policy to install software and configure applications on a managed server. See “Installing Software Using a Software Policy” on page 449 in this chapter for more information.
- Select a single patch and install it directly on a managed server. See “Patch Management for Windows” on page 323 in this chapter for more information.
- Use Application Configuration Management to configure applications on a managed server. See “Application Configuration Management” on page 495 in this chapter for more information.

## Installing or Uninstalling Software on a Server



You must have the Allow Install/ Uninstall Software permission to install or uninstall software on a managed server.. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide* for more information.

In the SAS Client you can install or uninstall software directly on managed servers without using a software policy.

### Ways to Open the Install or Uninstall Software Window

#### From the server list:

- 1 From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices** ► **Device Groups**. The device group list appears in the Content pane.

- 2 From the Content pane, select a server or device group.
- 3 From the **Actions** menu, select **Install** ► **Software**. The Install Software window appears.

or

From the **Actions** menu, select **Uninstall** ► **Software**. The Uninstall Software window appears

**From the By Type view in the Library:**

- 1** From the Navigation pane, select **Library** ► **By Type** ► **Type of Software**. The software list appears in the Content pane.
- 2** From the Content pane, select a software.
- 3** From the **Actions** menu, select **Install Software**. The Install Software window appears.

**From the Server Explorer**

- 1** From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices** ► **Device Groups**. The device group list appears in the Content pane.

- 2** From the Content pane, select a server or a server from the device group.
- 3** From the **Actions** menu, select **Open**. The Server Explorer window appears.
- 4** From the Views pane, select **Inventory** and then select the type of software.
- 5** From the **Actions** menu, select **Install Software**. The Install Software window appears.

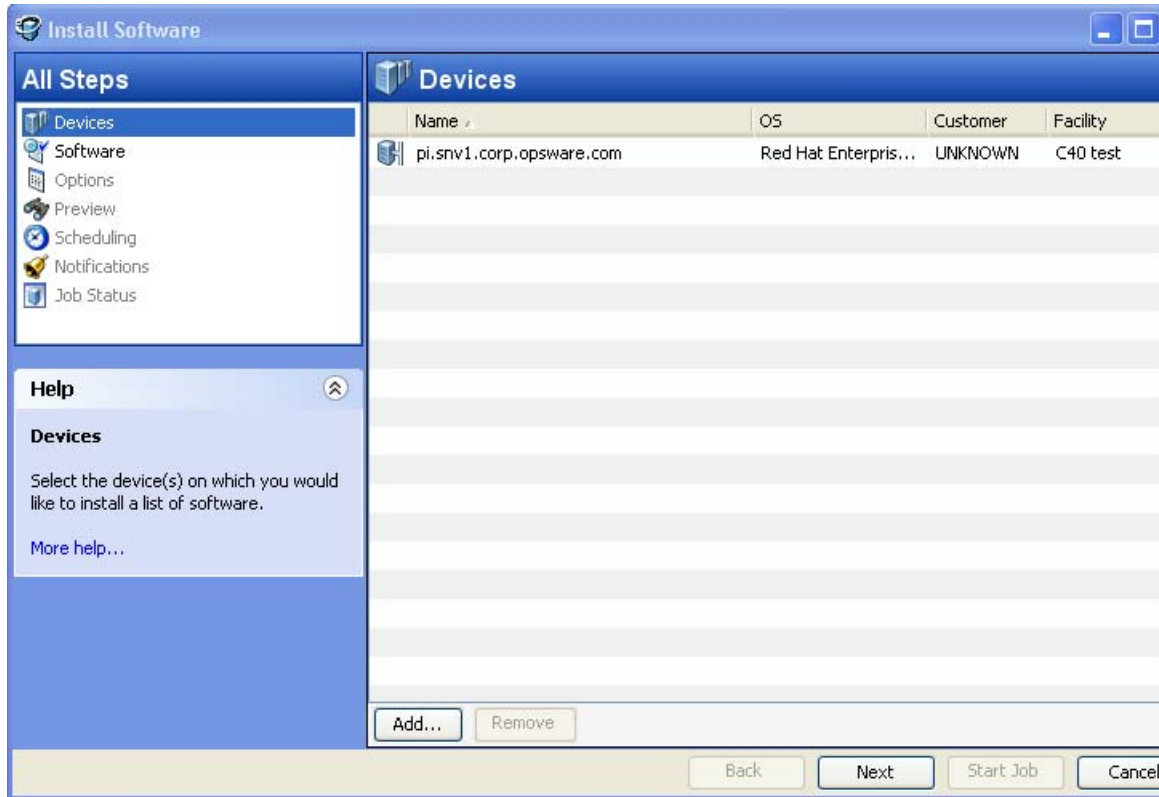
or

From the **Actions** menu, select **Uninstall Software**. The Uninstall Software window appears.

## Installing or Uninstalling Software

The Install Software window as shown in Figure 7-3 allows to you install software directly on a managed server and consists of the following steps:

Figure 7-3: Install Software Window



- Selecting Devices
- Selecting Software
- Specifying Options
- Preview
- Scheduling
- Setting Email Notifications
- Viewing Job Status

### Selecting Devices

In this step, you can specify the servers for installing or uninstalling the software.





Perform the following steps to select the servers:

- 1** Open the Install Software window or Uninstall Software window from one of the methods described in “Ways to Open the Install or Uninstall Software Window” on page 441
- 2** In the Install Software Policy window, select the Devices step.
- 3** (Optional) Click **Add** to add additional servers to the list or click **Remove** to remove servers from the list.
- 4** Select the servers.
- 5** Click **Next** to proceed to the Software step.

### Selecting Software

In this step, you can specify the software such as packages, patches, application configurations to install or uninstall servers. You can also specify the order in which you want to install or uninstall the software on the server.

Perform the following steps to select the software:

- 1** From the Install Software window or Uninstall Software window, click **Next** to advance to the Software step.
- 2** In the Install Software window or Uninstall Software window, click . The Select Library window appears.
- 3** In the Select Library window, select the software to be installed or uninstalled on the servers and click **Select**.
- 4** (Optional) Click  to remove any software.
- 5** (Optional) Click  or  to order the software resources.
- 6** Click **Next** to proceed to the Options Step.

### Specifying Options

In this step you can set the following installation or uninstallation options:

- You can specify the reboot actions required for the installation or uninstallation process. You can control when to reboot servers during installation or uninstallations to minimize the downtime caused by server reboots.

- You can choose to continue with the installation or uninstallation process if an error occurs during the installation or uninstallation of any software.
- You can specify the scripts to run on a server before or after installation or uninstallation. The scripts include:
  - **Pre-Download:** (Only for Installing Software) A script that runs before packages or patches are downloaded from Opware SAS to the server.
  - **Post-Download:** (Only for Installing Software) A script that runs after packages or patches are downloaded from Opware SAS to the server and before the package or patch is installed
  - **Pre-Install/ Pre-Uninstall:** A script that runs before packages or patches are installed or uninstalled on the server.
  - **Post-Install/ Post-Uninstall:** A script that runs after packages or patches are installed or uninstalled on the server.

Perform the following steps to specify the options for installation or uninstallation:

- 1** From the Install Software window or Uninstall Software window, click **Next** to advance to the Options step.
- 2** Select one of the following Reboot options:
  - Reboot servers as specified by individual software items  
This option allows you to reboot servers depending on the reboot option specified in the software resources properties window.
  - Reboots servers after each installation or uninstallation  
This option allows you to reboot servers after installing or uninstalling software.
  - Hold all server reboots until all actions are complete  
If the reboot option is selected in the software resources properties , this option allows you to reboot the servers after all the software resources are installed and uninstalled. If the reboot option is not selected in the software resources properties, this option does not reboot the server after all the software resources are installed and uninstalled.
  - Suppress all reboots  
This option allows you to suppress the reboots even if the reboot option is selected in the software resources properties.

- 3** Select "Attempt to continue running if an error occurs", if you want the installation or uninstallation process to continue even when an error occurs with any of the package, patches or scripts. By default, this check box is not selected.
- 4** In the Scripts section, select the Pre-Download, or Post-Download, or Pre-Install, or Post-Install tab. You may specify different scripts and options on each of the tabs. You require certain Script permissions to select these options. See the *Opware® SAS Administration Guide* for more information about the permissions required.
  1. Select **Enable Script**. Selecting Enable Script enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
  2. Select Saved Script or Ad-Hoc Script from the drop-down list. A Saved script is stored in Opware SAS after you upload the script to Opware SAS. An Ad-Hoc script is intended only for one operation and is not stored in Opware SAS.
  3. If you selected Saved Script from the drop-down list, click **Select** to specify the script. The Select Script window appears. Select the scripts to run and click **Select**.
  4. If you selected Ad-Hoc Script from the drop-down list, select the type from the Type drop-down list and then enter the contents of the script in the Script field.
  5. Enter the command-line flags in the Command field, if required.
  6. Enter a script time-out value in minutes in the Script Timeout field.
  7. In the User section, select Root to execute the script as root. To execute the script as a specified user, select Name and enter the user name and then the password.

To enter a Windows Domain Name in the pre-download, post-download, pre-install, post-install scripts, use the following format in the Name field:

DomainName\UserName and then enter the password in the password field.
  8. Select "Stop job if script returns an error" to stop the installation if the script returns an error.
- 5** Click **Next** to proceed to the Preview step.

### **Preview**

In this step, you have the option to preview the installation or uninstallation process.

The preview option allows you to view a detailed list of actions performed on a server as a result of installation or uninstallation of software. It displays information for each server that is selected for installation or uninstallation. Preview shows the software resources that will be installed on or uninstalled from a server, the application configurations that will be applied to a server, the dependency information required for the packages or patches to be installed, the reboots required during the installation or uninstallation process, and the scripts that will be executed.

Perform the following steps to preview the installation or uninstallation process:

- 1** From the Install Software window or Uninstall Software window, click **Next** to advance to the Preview step.
- 2** (Optional) Click **Preview** to view the separate actions that will be performed during the installation or uninstallation process. To view the details of each of the actions, select a row in the table. The details for each action appear.
- 3** Select **Output** to view the job output or select **Errors** to view the error details.
- 4** Click **Next** to proceed to the Scheduling step.

### **Scheduling**

In this step, you can schedule the analysis, download and installation or uninstallation stage to be run immediately or at a specified date and time.

Perform the following steps to schedule the installation or uninstallation process:


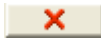
- 1** From the Install Software window or Uninstall Software window, click **Next** to advance to the Scheduling step.
- 2** In the Schedule Analysis section, select one of the following options:
  - Run at Job Start: This option allows you to run the job immediately.
  - Start time: This option allows you to specify the date and time to schedule the job.
- 3** In the Schedule Download section, select one of the following options: (Only for Installing Software)
  - Run Immediately After Analysis: This option allows you to download software immediately.
  - Start time: This option allows you to specify the date and time to download software.

- 4** In the Schedule Install or Schedule Uninstall section, select one of the following options:
  - Run Immediately : This option allows you to install or Uninstall software immediately.
  - Start time: This option allows you to specify the date and time to install or uninstall software.
- 5** Click **Next** to proceed to the Notifications step.

### Setting Email Notifications

In this step, you can set email notifications to alert users on the success or failure of the of the installation or uninstallation process. You can also associate a Ticket ID with the installation or uninstallation process.

Perform the following steps to set email notifications:

- 1** From the Install Software window or Uninstall Software window, click **Next** to advance to the Notification step.
- 2** To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon.
- 4** Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5** Click **Next** to go to the Job Status step.

### Viewing Job Status

In this step, you can view the summary information for the progress of a job and the individual status of each action required to be performed for the job to be completed.

Perform the following steps to view the job status:

- 1** From the Install Software window or Uninstall Software window, click **Next** to advance to the Job Status step.
- 2** If you choose to run the job immediately in the Scheduling step, the job begins immediately. If you scheduled the job for a later time, the job will run at the scheduled time. The job progress appears in the Install Software window.



- 3 To view the details of each action, select a row in the table. The details for each action appear.
- 4 Select **Output** to view the job output or select **Errors** to view the error details.
- 5 Click **End Job** to stop the job from running or click **Close** to close the Install Software window.

## Installing Software Using a Software Policy

Installing software by using a software policy includes the following steps:

- Attaching a software policy to a server or attaching a server to a software policy
- Remediating a server against a software policy

### Attaching a Software Policy to a Server

When you attach a software policy to a server or group of servers, the software policy is associated with that server or group of servers. This action does not install the software contained in the software policy. To install the software, you must remediate the server with the software policy. See “Remediating Software Policies” on page 454 in this chapter for more information.



You must have a set of permissions to attach a software policy to a server. To obtain these permissions, contact your Opware administrator. See the *Opware<sup>®</sup> SAS Administration Guide* for more information.

Perform the following steps to attach a software policy to a server:

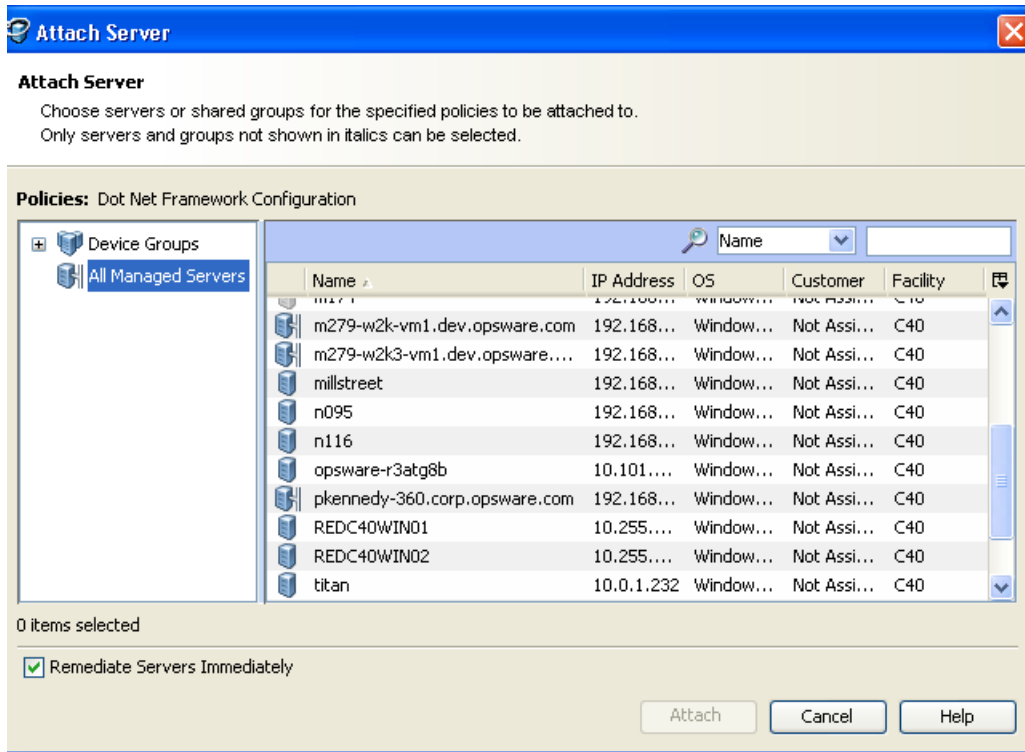
- 1 From the Navigation pane, select **Library** ► **By Type** ► **Software Policies**. The software policies appear in the Content pane.
- 2 (Optional) From the Content pane, select the software policy.
  1. Open the software policy. The Software Policy Window appears.
  2. From the View pane, select Server Usage.
  3. From the Content pane, select a server.

Or

1. From the View drop-down list in the Content pane, select Server Usage.
2. Select a server.

- 3 From the **Actions** menu, select **Attach Server**.The Attach Server window appears as shown in Figure 7-4:

Figure 7-4: The Attach Server Window in the SAS Client



- 4 In the Attach Server window, select servers or device groups and then click **Attach**. You can only select servers that are not in italics. Servers in italics indicate that you do not have the permission to attach a software policy to the server.
- 5 (Optional) Select “Remediate Servers Immediately” to remediate the servers against the software policy. Selecting this option displays the Remediate window. This option is only available if you have the Remediate Servers permission. See “Remediating Software Policies” on page 454 in this chapter for more information.

## Attaching a Server to a Software Policy

When you attach a server or group of servers to a software policy, the software policy is associated with that server or group of servers. This action does not install the software contained in the software policy. To install the software, you must remediate the server with the software policy. See “Remediating Software Policies” on page 454 in this chapter for more information.



---

You must have a set of permissions to attach a server to a software policy. To obtain these permissions, contact your Opware administrator. See the *Opware<sup>®</sup> SAS Administration Guide* for more information.

---

**1** From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.

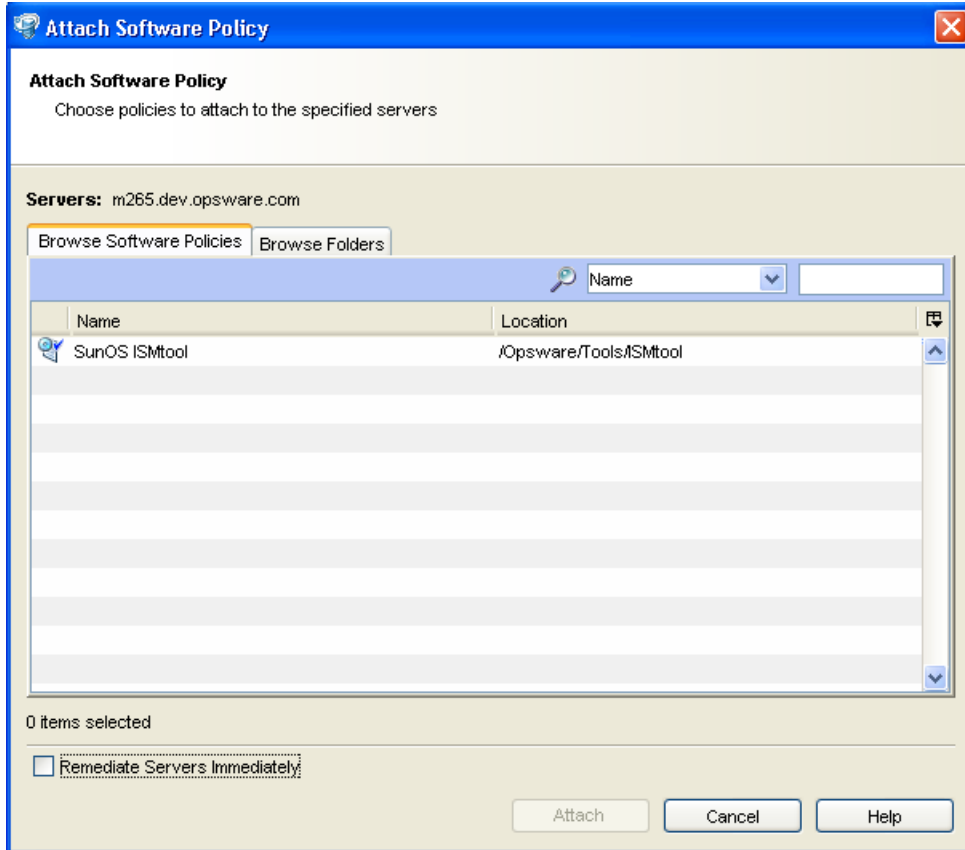
Or

From the Navigation pane, select **Devices** ► **Device Groups**. The device group list displays in the Content pane.

**2** From the Content pane, select a server or a device group.

- From the **Actions** menu, select **Attach** ► **Software Policy**. The Attach Software Policy window appears as shown in Figure 7-5.

Figure 7-5: The Attach Policy Window in the SAS Client



- Select Browse Software Policies and then select the software policies from the list.  
Or  
Select Browse Folders and then select the software policies from the folder hierarchy.
- Click **Attach**.
- (Optional) Select “Remediate Servers Immediately” to remediate the servers against the software policy. Selecting this option displays the Remediate window. This option is only available if you have the Remediate Servers permission. See “Remediating Software Policies” on page 454 in this chapter for more information.

## Overview of Software Policies Remediation

The remediation process installs the software resources, server objects, and applies the configurations specified in a software policy to a server. (A software policy must be attached to a server or a group of servers before you can remediate the software policy with that server or group of servers.) When you detach a software policy from a server and remediate, then the remediation process uninstalls the software in a software policy.

The remediation process allows you to specify remediation options and pre and post installation scripts required for the remediation process, schedule the analysis, download, and the installation phase of the remediation process, set up email notifications to alert you about the status of the remediation process, and associate a Ticket ID with each remediation process.



You must have a set of permissions to remediate policies. To obtain these permissions, contact your Opsware administrator. See the *Opsware® SAS Administration Guide* for more information.

The Remediate window allows you to remediate the servers against the software policies and define the conditions for remediation.

## Ways to Open the Remediate Window

### *From the server list:*

**1** From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices** ► **Device Groups**. The device group list appears in the Content pane.

**2** From the Content pane, select a server or device group.

**3** From the **Actions** menu, select **Remediate**. The Remediate window appears.

### *From the software policies list:*

**1** From the Navigation pane, select **Library** ► **By Type** ► **Software Policies**. The software policy list appears in the Content pane.

**2** From the Content pane, select a software policy.

1. From the View drop-down list, select Server Usage.
2. Select servers and then select **Remediate** from the **Actions** menu. The Remediate window appears.

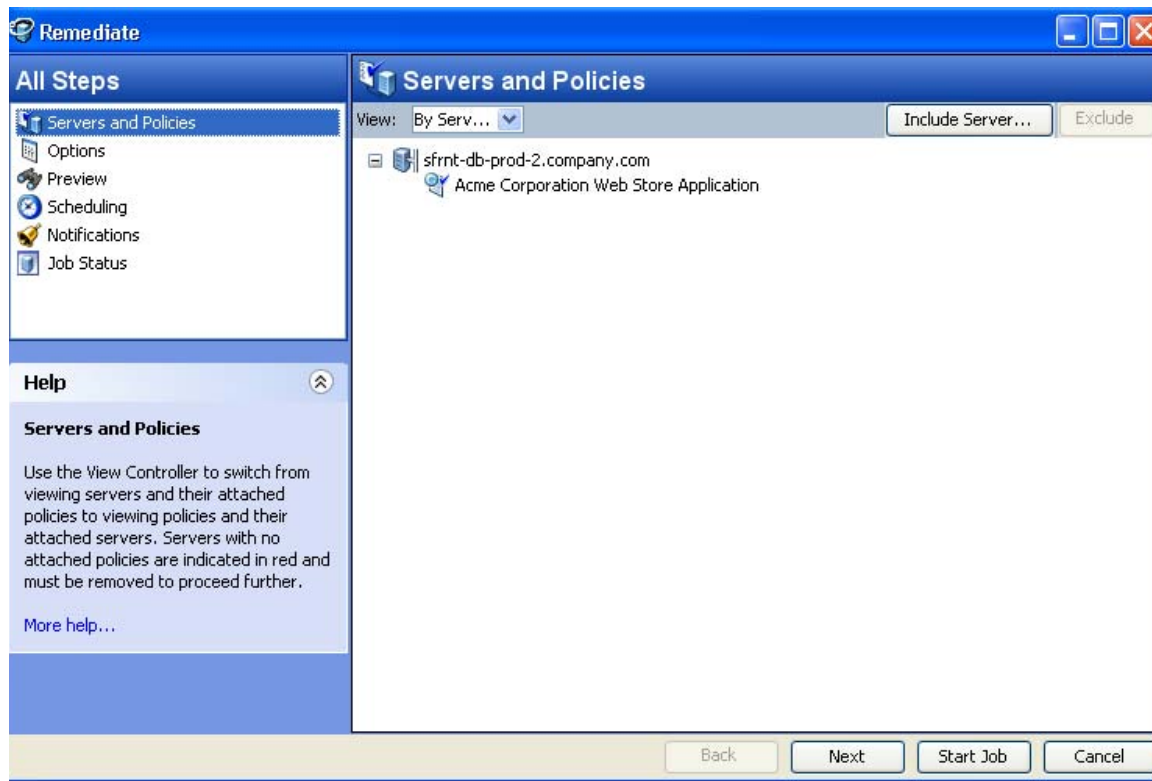
Or

1. From the Content pane, open a software policy. The Software Policy Window appears.
2. From the View pane, select Server Usage.
3. Select servers and then select **Remediate** from the **Actions** menu. The Remediate window appears.

### Remediating Software Policies

The Remediate window as shown in Figure 7-6, allows to you remediate the servers against the software policies and consists of the following steps:

Figure 7-6: The Remediate Window in the SAS Client



- Selecting Servers and Policies for Remediation
- Specifying Options for Remediation
- Preview Software Policy Remediation
- Scheduling Software Policy Remediation
- Setting Email Notifications for Remediation
- Viewing Job Status

### **Selecting Servers and Policies for Remediation**

This step allows you to specify the servers (with software policies attached) for remediation. In this step, you can add and remove servers from the list, view all the software policies attached to a server, and remove software policies attached to servers.

Perform the following steps to select servers and policies for remediation:

- 1** Open the Remediate window from one of the methods described in “Ways to Open the Remediate Window” on page 453.
- 2** In the Remediate window, select the Servers and Policies step. The servers with attached software policies and patch policies appear.

A software policy is represented by the icon .

A patch policy is represented by the icon .

You can also view a list of policies with attached servers by selecting By Policies from the View drop-down list.

- 3** (Optional) Click **Include Server** to add servers to the list or select a server and click **Exclude** to remove servers from the list.
- 4** Select servers with attached software policies.
- 5** Click **Next** to proceed to the Options step.

### **Specifying Options for Remediation**

In this step you can set the following options:

- You can specify the reboot actions required for the remediation process. You can control when to reboot servers during remediation to minimize the downtime caused by server reboots.

- You can choose to continue with the remediate process if an error occurs during the installation or uninstallation of any software contained in the software policy.
- You can specify the scripts to run on a server before or after remediation. The scripts include:
  - **Pre-Download:** A script that runs before packages or patches are downloaded from Opware SAS to the server.
  - **Post-Download:** A script that runs after packages or patches are downloaded from Opware SAS to the server and before the package or patch is installed.
  - **Pre-Remediate:** A script that runs before packages or patches are installed on the server.
  - **Post-Remediate:** A script that runs after packages or patches are installed on the server.

Perform the following steps to specify the options for remediation:

- 1** From the Remediate window, click **Next** to advance to the Options step.
- 2** Select one of the following Reboot options:
  - Reboot servers as specified by individual software items  
This option allows you to reboot servers depending on the reboot option specified in the software resources properties window.
  - Reboots servers after each installation or uninstallation  
This option allows you to reboot servers after installing or uninstalling software.
  - Hold all server reboots until all actions are complete  
If the reboot option is selected in the software resources properties , this option allows you to reboot the servers after all the software resources are installed and uninstalled. If the reboot option is not selected in the software resources properties, this option does not reboot the server after all the software resources are installed and uninstalled.
  - Suppress all reboots  
This option allows you to suppress the reboots even if the reboot option is selected in the software resources properties.





If a software policy contains multiple non RPM type packages with the option "reboot =yes" selected for every package in the Package Properties window, and the option "Reboot as dictated by package properties" selected in the Remediate window, then remediating a sever with the software policy will reboot the server every time a package is installed. If a software policy contains multiple RPM type packages with the option reboot =yes selected for every RPM package in the Package Properties window, and the option "Reboot as dictated by package properties" selected in the Remediate window, then remediating a sever with software policy will reboot the server only once after all the RPM packages are installed.

- 3** Select "Attempt to continue running if an error occurs", if you want the remediate process to continue even when an error occurs with any of the package, patches or scripts. By default, this check box is not selected.
- 4** In the Scripts section, select the Pre-Download, or Post-Download, or Pre-Remediate, or Post-Remediate tab. You may specify different scripts and options on each of the tabs. You require certain Script permissions to select these options. See the *Opware<sup>®</sup> SAS Administration Guide* for more information about the permissions required.
  1. Select **Enable Script**. Selecting Enable Script enables the remainder of the fields on the tab. If Enable Script is not selected, the script will not run.
  2. Select Saved Script or Ad-Hoc Script from the drop-down list. A Saved script is stored in Opware SAS after you upload the script to Opware SAS. An Ad-Hoc script is intended only for one operation and is not stored in Opware SAS.
  3. If you selected Saved Script from the drop-down list, click **Select** to specify the script. The Select Script window appears. Select the scripts to run and click **Select**.
  4. If you selected Ad-Hoc Script from the drop-down list, select the type from the Type drop-down list and then enter the contents of the script in the Script field.
  5. Enter the command-line flags in the Command field, if required.
  6. Enter a script time-out value in minutes in the Script Timeout field.
  7. In the User section, select Root to execute the script as root. To execute the script as a specified user, select Name and enter the user name and then the password.

To enter a Windows Domain Name in the pre-download, post-download, pre-install, post-install scripts, use the following format in the Name field:

`DomainName\UserName` and then enter the password in the password field.

8. Select "Stop job if script returns an error" to stop the installation if the script returns an error.

**5** Click **Next** to proceed to the Preview step.

### **Preview Software Policy Remediation**

In this step, you have the option to preview the remediation process.

The preview option allows you to view a detailed list of actions performed on a server as a result of installation or uninstallation of software. It displays information for each server that is selected for remediation. Preview shows the software resources that will be installed on or uninstalled from a server, the application configurations that will be applied to a server, the dependency information required for the packages or patches to be installed, the reboots required during the remediation process, and the scripts that will be executed during the remediation process.

Perform the following steps to preview the remediation process:

- 1** From the Remediate window, click **Next** to advance to the Review step.
- 2** (Optional) Click **Preview** to view the separate actions that will be performed during the remediation process. To view the details of each of the actions, select a row in the table. The details for each action appear.
- 3** Select **Output** to view the job output or select **Errors** to view the error details.
- 4** Click **Next to** proceed to the to the Scheduling step.

### **Scheduling Software Policy Remediation**

In this step, you can schedule the analysis, download and installation stage to be run immediately or at a specified date and time.

Perform the following steps to schedule the remediation process:



- 1** From the Remediate window, click **Next** to advance to the Scheduling step.
- 2** In the Schedule Analysis section, select one of the following options:
  - Run at Job Start: This option allows you to run the job immediately.
  - Start time: This option allows you to specify the date and time to schedule the job.

- 3** In the Schedule Download section, select one of the following options:
  - Run Immediately After Analysis: This option allows you to download software immediately.
  - Start time: This option allows you to specify the date and time to download software.
- 4** In the Schedule Remediate section, select one of the following options:
  - Run Immediately After Remediate: This option allows you to install software immediately.
  - Start time: This option allows you to specify the date and time to install software.
- 5** Click **Next** to proceed to the Notifications step.

### **Setting Email Notifications for Remediation**

In this step, you can set email notifications to alert users on the success or failure of the the remediation process. You can associate a Ticket ID with the remediation process.

Perform the following steps to set email notifications:

- 1** From the Remediate window, click **Next** to advance to the Notification step.
- 2** To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3** To set the notification status on the success of a Job, select the  icon. To set the notification status on the failure of a Job, select the  icon.
- 4** Enter a Ticket ID to be associated with a Job in the Ticket ID field.
- 5** Click **Next** to go to the Job Status step.

### **Viewing Job Status**

In this step, you can view the summary information for the progress of a job and the individual status of each action required to be performed for the job to be completed.

Perform the following steps to view the job status:

- 1** From the Remediate window, click **Next** to advance to the Job Status step.
- 2** If you choose to run the job immediately in the Scheduling step, the job begins immediately. If you scheduled the job for a later time, the job will run at the scheduled time. The job progress appears in the Remediate window.

- 3** To view the details of each action, select a row in the table. The details for each action appear.
- 4** Select **Output** to view the job output or select **Errors** to view the error details.
- 5** Click **End Job** to stop the job from running or click **Close** to close the Remediate window.



---

You can also view all your jobs from the job logs in the SAS Client. See the *Opware® SAS User's Guide: Server Automation* for information about job logs.

---

## Uninstalling Software Using a Software Policy

Uninstalling software by using a software policy includes the following steps:

- Detaching a Software Policy from a Server
- Remediating Software Policies

### Detaching a Software Policy from a Server

Detaching a software policy from a server does not delete the policy or uninstall the software from a server. To uninstall the software, you must detach the software policy from the server and then remediate the server with the software policy. See "Remediating Software Policies" on page 454 in this chapter for more information.



---

You must have a set of permissions to detach a software policy from a server. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide* for more information.

---

Perform the following steps to detach a software policy from a server:

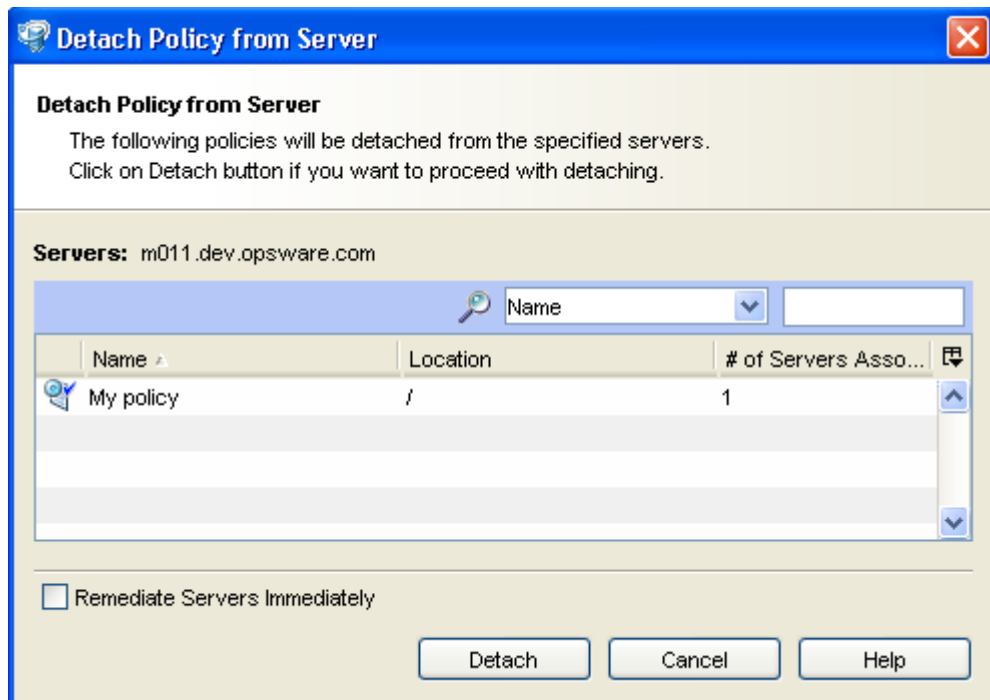
- 1** From the Navigation pane, select **Devices > Servers > All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices > Device Groups**. The device group list appears in the Content pane.

- 2 From the Content pane, select a server or a device group.
- 3 From the View drop-down list, select Software Policies.
- 4 From the **Actions** menu, select **Detach**. The Detach Software Policy window appears as shown in Figure 7-7.

Figure 7-7: The Detach Software Policy Window in the SAS Client



- 5 Click **Detach**.
- 6 (Optional) Select “Remediate Servers Immediately” to remediate the servers against the software policy. Selecting this option will display the Remediate window. See “Remediating Software Policies” on page 454 in this chapter for more information.

## Overview of Software Template

Opsware SAS allows you to install software by using a software template. A software template can only contain other software policies. A software template is not persistently associated with a server or group of servers. When you install a software template to a server or group of servers, the software policies specified in the software template are

installed. If you update a software template, servers that already had the software template applied are not automatically modified to match the updated software template. You must install the software template again to reflect the changes made to the software template on the server.

A software template has the following features:

- A software template is not associated with a server or group of servers.
- A software template contains other software policies.
- A software template is associated with an operating system family.
- Software templates are located in folders.
- Custom attributes can be set on a software template.

Installing software on a server by using a software template consists of the following steps:

- Creating a software template

See the *Opware® SAS Policy Setter's Guide* for information about creating a software template.

- Adding software policies to a software template

See the *Opware® SAS Policy Setter's Guide* for information about adding software policies.

- Installing the software template

See “Installing or Uninstalling Software” on page 443 in this chapter for more information.

## Overview of Running ISM Controls

The Run ISM Control window in the SAS Client allows you to run the control scripts in an ISM (Intelligent Software Module).

To run the control scripts in an ISM, you must add the ISM package to a software policy first and then attach the software policy to a server.

See the *Opware® SAS Policy Setter's Guide* for information about adding an ISM package to a software policy. See “Attaching a Software Policy to a Server” on page 449 in this chapter for more information.



You must have a set of permissions to run an ISM Control. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide* for more information.

## Ways to Open the Run ISM Control Window

### **From the server list:**

**1** From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices** ► **Device Groups**. The device group list appears in the Content pane.

**2** From the Content pane, select a server or device group.

**3** From the **Actions** menu, select **Run** ► **ISM Control**. The Run ISM Control window appears.

### **From the software policies list:**

**1** From the Navigation pane, select **Library** ► **By Type** ► **Software Policies**. The software policy list appears in the Content pane.

**2** From the Content pane, select a software policy containing an ISM.

1. From the View drop-down list, select Server Usage.

2. Select servers and then select **ISM Control** from the **Actions** menu. The Run ISM Control window appears.

Or

1. From the Content pane, open a software policy containing ISM package. The Software Policy Window appears.

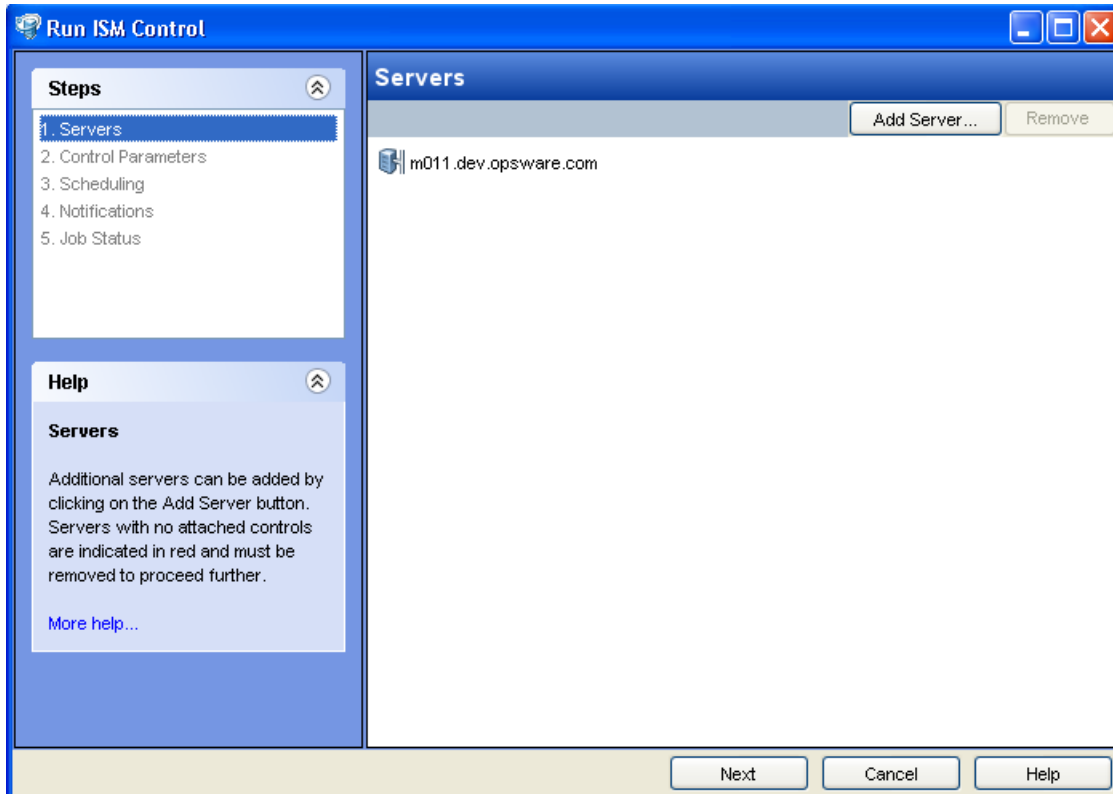
2. From the View pane, select Server Usage.

3. Select servers and then select **Run ISM Control** from the **Actions** menu. The Run ISM Control window appears.

## Running ISM Controls

The Run ISM Control window, as shown in Figure 7-8, allows you to run an ISM Control on a server and consists of the following steps:

Figure 7-8: The Run ISM Control Window in the SAS Client



- Selecting Servers
- Selecting Control Parameters
- Scheduling ISM Control Script Execution
- Setting Email Notifications
- Viewing Job Status

### Selecting Servers

In this step, you can specify the servers for running an ISM Control.

Perform the following steps to select servers:



- 1** In the Run ISM Control window, select the Servers.
- 2** (Optional) Click **Include Server** to add additional servers to the list or click **Exclude** to remove servers from the list.
- 3** Select the servers.
- 4** Click **Next** to proceed to the Control Parameters step.

### **Selecting Control Parameters**

In this step, you can select a control script in an ISM package to be executed.

Perform the following steps to select the control parameters:

- 1** From the Run ISM Control window, click **Next** to advance to the Control Parameters step.
- 2** From the Software Policy drop-down list, select an ISM package.
- 3** From the Control script drop-down list, select a control script. The drop-down list contains only the control scripts assigned to the ISM package selected in the previous step.
- 4** In the Parameters section, the name of a parameter matches the name of its corresponding custom attribute name. The value of a custom attribute determines the value of the parameter.
- 5** Click **Next** to proceed to the Scheduling step.

### **Scheduling ISM Control Script Execution**

In this step, you can schedule an ISM Control script to be run immediately or at a specified date and time.



Perform the following steps to schedule the ISM Control script execution:

- 1** From the Run ISM Control window, click **Next** to advance to the Scheduling step.
- 2** Select one of the following options:
  - **Run Task Immediately:** This option allows you to run the ISM control script immediately.
  - **Run Task At:** This option allows you to specify the date and time to run the ISM control script.
- 3** Click **Next** to proceed to the Notification step.

### Setting Email Notifications

In this step, you can set email notifications to alert users on the success or failure of ISM control script execution. You can associate a Ticket ID with the ISM Control script execution job.

Perform the following steps to set email notifications:

- 1 From the Run ISM Control window, click **Next** to advance to the Notification step.
- 2 To add email addresses, click **Add Notifier** and enter the email addresses in the Notification Email Address field.
- 3 To set the notification status on the success of a job, select the  icon.  
To set the notification status on the failure of a job, select the  icon.
- 4 Enter a Ticket ID to be associated with the job in the Ticket ID field.
- 5 Click **Next** to go to the Summary Review step.

### Viewing Job Status

In this step, you can view the summary information for the progress of a Job and the status of each action required for the Job to be completed.

Perform the following steps to view the job status:

- 1 From the Run ISM Control window, click **Start Job** to advance to the Job Status step.
- 2 If you selected Run Task Immediately in the Scheduling step, the job begins immediately. If you scheduled the job for a later time, the job will run at the scheduled time. The job progress appears in the Run ISM Control window.
- 3 To view the details of each action, select a row in the table. The details for each action will appear.
- 4 Click **End Job** to stop the Job from running or click **close** to close the Run ISM Control window.



You can also view all your jobs from the job logs in the SAS Client. See *Opware® SAS User's Guide: Server Automation* for information about job logs.

---





## Software Policy Compliance

Software compliance indicates whether or not the software policies attached to the selected server are compliant with the actual server configuration. A software policy scan compares the actual configuration of the server with the software policies attached to that server. If the actual server configuration does not match the software policies attached to a server, then the server is said to be out of compliance with the software policies.

A server can be either compliant or non-compliant with respect to a software policy attached to it. If the server's configuration does not match the packages, patches, server objects, software policies and application configurations defined in a software policy (attached to that server), then the server is said to be non-compliant with that software policy. If a software policy contains scripts, then the scripts are not used to calculate the software compliance.

In the SAS Client, when you perform a software compliance scan, the scan indicates the server's overall compliance state as a result of all the software policies attached to the server. Even if only one software policy attached to the server is not compliant, the server is said to be non-compliant. You can then view the non-compliant server and remediate the server against that software policy.

The SAS Client displays the following compliance information for a software policy:

- **Compliant:** If all the software policies attached to a server are compliant, the server is compliant and is represented by the icon .
- **Non-compliant:** If one of the software policies attached to a server is not compliant, the server is non-compliant and is represented by the icon .
- **Scan Started:** The software compliance information is currently being calculated and is represented by the icon .
- **Scan Needed:** The software compliance information needs to be calculated or the compliance information might be inaccurate and is represented by the icon .
- **Not Applicable:** The software compliance information does not apply and is represented by the icon -.

For example, if you detach a software policy from a server and do not remediate the server against the software, policy, then the compliance status is represented as Not Applicable.

In the SAS Client, you can check for software compliance from the server list or from the Compliance View for a device or device group.

See “Checking Software Compliance Scan” on page 468 in this chapter for information about performing a compliance scan from the server list.

See “Server Compliance” on page 257 in Chapter 3 for information about the Compliance View for a device or device group.

## Checking Software Compliance Scan



---

You must have a set of permissions to perform a software compliance scan. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide* for more information.

---

Perform the following steps to scan a server for software compliance:

- 1** From the Navigation pane, select **Devices > Servers > All Managed Servers**. The server list appears in the Content pane.
- 2** From the Content pane, select the server.
- 3** From the **Actions** menu, select **Scan > Software Compliance**. The compliance status of the server appears in the server list.

After you perform a software compliance scan, you can view the software policies that are not compliant and then remediate the server against that software policy. The software compliance scan is automatically updated after you install or uninstall software or remediate a software policy against a managed server. See “Remediating Software Policies” on page 454 in this chapter for more information. See “Installing or Uninstalling Software” on page 443 in this chapter for more information

## Software Policy Reports

The Reporting feature in Opware SAS allows you to generate reports that provide a summary of the software policy compliance across servers. You can also generate reports that provide information about software policies on a given server. After you generate reports, you can print them, export the reports to .html and .xls, and perform actions on the results.

The Opsware SAS allows you to run the following software policy reports:

- **Software Policy Compliance:** This report groups all managed servers by their software policy compliance level to show compliant and non-compliant servers.
- **Software Policy Compliance By Customer:** This report lists all servers by the customer they are associated with and then by the software policy compliance level.
- **Software Policy Compliance By Facility:** This report displays a chart of all servers by the facility they are associated with and then by the software policy compliance level.
- **Defined Software Policies:** This report lists all the software policies by name and their location in the folder hierarchy.
- **Servers With Attached Software Policies:** This report lists all servers that have one or more software policies attached.
- **Servers In Compliance With Their Software Policies:** This report lists all servers that are in compliance with all of their attached software policies.
- **Servers Not In Compliance with their Software Policies:** This report lists all servers that are not in compliance with all of their attached software policies.
- **Servers Without Attached Software Policies:** This report lists all servers that have no software policies attached.

See “SAS Client Reports” on page 307 in Chapter 4 for information about how to run and view reports in the SAS Client.



# Chapter 8: Script Execution

## IN THIS CHAPTER

This section discusses the following topics:

- Overview of Script Execution
- Script Execution Process
- Managing Scripts
- Executing Scripts
- Script Execution Permissions

## Overview of Script Execution

The Script Execution feature allows you to automate the management and execution of scripts in the SAS Client. It also allows you to organize your scripts in folders and define security permissions around them. From the SAS Client, 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 is executed, you can view the results for every server and then export the script results. You can also modify, delete, and rename a script. You can also execute scripts in the Opware Global Shell using the SAS Client.

## Script Execution Features

The Script Execution feature in the SAS Client enables you to perform the following functions:

- Organize your scripts into folders and define security permissions to control access of their contents across different users and user groups.
- Create or upload scripts in the SAS Client.
- Run scripts across multiple Unix or Windows servers or server groups.
- Execute scripts in the Opware Global Shell.
- Schedule one time or recurring script execution jobs.

- Notify the status of the script execution job via email.
- Approve script execution jobs.
- View the script output against multiple servers in a tabular format.
- Export the script execution results.
- Search for scripts and script execution jobs.

## Script Execution Process

The script execution process involves defining permissions, managing scripts, and executing scripts.

### • Defining Permissions

In this phase, an Opsware administrator assigns Folder permissions, Client feature permissions, and Customer constraints to define the security boundaries across various user groups. The permissions determines the actions the users in a user group can perform with the SAS Client.

See the *Opsware® SAS Administration Guide* for more information about defining security permissions.

### • Managing Scripts

In this phase, a policy setter or an advanced system administrator performs script management tasks such as creating or importing scripts, editing script properties, exporting scripts, and deleting scripts. See “Managing Scripts” on page 473 for more information.

### • Executing Scripts

In this phase, a system administrator executes server scripts directly on servers or server groups and OFGS scripts in the Opsware Global Shell. A system administrator can also execute scripts by adding the scripts to a software policy and then remediating the servers against the software policy. See “Executing Scripts” on page 482 and “Installing Software Using a Software Policy” on page 449 for more information.

## Types of Scripts

In the SAS Client, the Script Execution feature supports two main types of scripts : Server scripts and OGFS scripts.



The Server script allows you to execute scripts on Unix and Windows servers managed by Opware SAS. SAS Client supports the following types of Server scripts for Unix and Windows operating systems: Unix/Linux shell, Windows batch (.BAT), Windows Visual Basic (VBScript), and Windows PowerShell.

The OGFS scripts allows you to execute scripts in the Opware Global Shell from the SAS Client. You can specify the directory path in the OFGS to execute the scripts. See the *Opware® SAS User's Guide: Server Automation* for information about the Opware Global Shell.

The server scripts are further classified to Saved Scripts, and Ad-hoc Scripts.

- Saved scripts are accessible to all the users, if they have the appropriate permissions. You are required to have the appropriate permissions to create, view, edit, and execute shared scripts. Private scripts are only accessible to the user who created them. They can only be created, edited, deleted, or executed by the user who created the script.
- Ad-Hoc scripts are created or uploaded for one-time use and is not stored in Opware SAS. Ad- Hoc script is created or uploaded and then immediately executed by a user and during this process , only one user has access to the script.

After you create a script and save it as a specific type of script in Opware SAS, you cannot convert the script to the other type of script. Private scripts cannot be converted to Saved Scripts and vice versa.

In the SAS Client, you can specify to run a Sever script as a Super User or as a specified user. A Super User script allows you to execute the script as root on UNIX or Local System on Windows servers without entering a password. If the script is not designated as a Super User Script, then you need to enter a username and password to run the script. You also require the appropriate permissions to manage and run Super User Server Scripts. See the *Opware® SAS Administration Guide* for information on the permissions required to run the Super User Server Scripts. All the OGFS scripts can only be executed as an Opware User.

## Managing Scripts

The script management tasks include:

- Creating a Script
- Opening a Script in the SAS Client
- Editing Script Properties

- Locating Scripts in Folders
- Exporting a Script
- Renaming a Script
- Deleting a Script



---

You must have a set of permissions to create and manage a script. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide* for more information.

---

### **Creating a Script**

In the SAS Client, you can create a script from either the By Type or the By Folder view in the Library.

#### **Script Creation Guidelines**

Opware SAS supports the following types of Server scripts for Unix and Windows operating systems: Unix/Linux shell, Windows batch (.BAT), Windows Visual Basic (VBScript), and Windows PowerShell.

When creating scripts you must 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 Opware-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, Opware 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.
- 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 a Script from the By Type View in the Library***

To create a script perform the following steps:

- 1** From the Navigation pane, select **Library** ► **By Type** ► **Scripts**. The three main types of scripts appear in the content pane.

- 2 Select the script type and then from the **Actions** menu, select **New script**. The Script window appears as shown in Figure 8-1.

Figure 8-1: Script Window.

**Properties**

Name: Simple BAT

Type: Windows .BAT

Location: / Select

Changes Server:  Yes  No

Run as super user:  Yes  No

Script Contents: Enter the script contents or import a script file Import Script File

```
dir c:\temp
```

Description: dir

Last Modified: Thu Aug 02 18:16:24 2007

Last Modified By: paul

Created: Thu Aug 02 18:16:18 2007

Created By: paul

Opware ID: 195380040

- 3 In the Name field, enter the name of the script.
- 4 (Windows only) Select the script type from the Type drop-down list.
- 5 Click **Select** to specify the location for the script in the folder hierarchy. The Select Folder window appears. Select a folder in the Library to specify the location of the script and then click **Select**.
- 6 In the Changes Server field, select Yes, if the script causes a change in the server configuration when executed.

- 7** In the Run as Super User field, select Yes if the script can be run as a Super User when executed. Selecting yes, allows you to run the script as a Super User without providing a password for the script.  
  
This option is enabled only if u have to appropriate permission. See the *Opware® SAS Administration Guide* for more information about script execution permissions.
- 8** In the Script Contents field, enter the contents of the script or click **Import Script File** to import a script.
- 9** In the Description field, enter text that describes the purpose or contents of the script.
- 10** To save the changes, select **Save** from the **File** menu.

### **Creating a Script from the By Folder View in the Library**

To create a script perform the following steps:

- 1** From the Navigation pane, select **Library ► By Folder**. The folder hierarchy in the Library appears in the Content pane.
- 2** Select the folder that should contain the script.
- 3** From the **Actions** menu, select **New ► Script**. The Script window appears.
- 4** In the Name field, enter the name of the script.
- 5** Select the script type from the Type drop-down list.
- 6** Click **Select** to change the location for the script in the folder hierarchy. The Select Folder window appears. Select a folder in the Library to specify the location of the script and then click **Select**.
- 7** In the Changes Server field, select Yes, if the script causes a change in the server configuration when executed.
- 8** In the Run as Super User field, select yes if the script can be run as a Super user when executed. OGFS Scripts can only be executed as an Opware User.  
  
This option is enabled only if u have to appropriate permission. See the *Opware® SAS Administration Guide* for more information about script execution permissions.
- 9** In the Script Contents filed, enter the contents of the script or click **Import Script File** to import a script. In the Open window, select the script to import and then click **Open**.
- 10** In the Description field, enter text that describes the purpose or contents of the script.

- 11** To save the changes, select **Save** from the **File** menu.



---

The Library in the SAS Client contains a Home directory and each user has a folder in the Home directory. You can save private scripts in this folder and later execute the script on managed servers.


---

### Opening a Script in the SAS Client

In the SAS Client, there are several ways to open a script. You can open a script from:

- The Search option in the Navigation pane
- The By Type view in the Library
- The By Folder view in the Library
- The Device list in the Navigation pane

#### Opening a Script from Search

- 1** From the Navigation pane, select **Search**.
- 2** Select Server Script or OGFS Script from the drop-down list and then enter the name of the script in the text field.
- 3** Select . The search results appear in the Content pane.
- 4** From the Content pane, select the script and then select **Open** from the **Actions** menu. The Script window appears.

#### Opening a Script from the By Type view in the Library

- 1** From the Navigation pane, select **Library** ► **By Type** ► **Scripts**. The scripts appear in the Content pane.
- 2** From the Content pane, select the script and then select **Open** from the **Actions** menu. The Script window appears.

#### Opening a Script from the By Folder view in the Library

- 1** From the Navigation pane, select **Library** ► **By Folder**. The folder hierarchy in the Library appears in the Content pane.
- 2** From the Content pane, select the script in a folder and then select **Open** from the **Actions** menu. The Script window appears.

### Opening a Script from Devices

- 1 From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.  
  
Or  
  
From the Navigation pane, select **Devices** ► **Device Groups**. The device groups list appears in the Content pane.
- 2 From the Content pane, select a server and then from the **Actions** menu, select **Open**. The Server Explorer window opens.
- 3 From the Views pane, select **Management Policies** ► **Software Policies**. The software policies attached to the server appear in the Content pane.
- 4 From the Content pane, select the software policy and then select **Open** from the **Actions** menu. The Software Policy window appears.
- 5 From the Views pane, select Policy Items. The policy items appear in the Content pane.
- 6 From the Content pane, select the script and then select **Open** from the **Actions** menu. The Script window appears.

### Editing Script Properties

After you create a script, you can view and modify its properties. You can view properties such as the Opsware user who created the script, the date when it was created, and the Opsware ID of the script. You can also modify the name, description, contents, the location of the script in the Library and the script options.

To view and edit script properties, perform the following steps:

- 1 Open a script in the SAS Client. See “Opening a Script in the SAS Client” on page 478 for ways to open a script. The Script window appears.
- 2 In the Name field, edit the name of the script.
- 3 Click **Select** to change the location for the script in the folder hierarchy. The Select Folder window appears. Select a folder in the Library to specify the location of the script and then click **Select**.
- 4 In the Changes Server field, select Yes, if the script causes a change in the server configuration when executed.

- 5 In the Run as Super User field, select yes if the script can be run as a Super User when executed. Selecting yes, allows you to run the script as a Super User without providing a password for the script.

This option is enabled only if you have the appropriate permission. See the *Opware® SAS Administration Guide* for more information about script execution permissions.

- 6 In the Script Contents field, edit the contents of the script or click **Import Script File** to import another script. In the Open window, select the script to import and then click **Open**.
- 7 In the Description field, edit the text that describes the purpose or contents of the script.
- 8 To save the changes, select **Save** from the **File** menu.

### Viewing All the Software Policies Associated with a Script


In the SAS Client, Server scripts can be added to a software policy. In the Scripts window, you can view all the software policies that contain the selected Server script. You cannot add OGFS script to a software policy.

To view the policy usage for a script, perform the following steps:

- 1 From the Navigation pane, select **Library > By Type > Scripts**.
- 2 From the Content pane, select the script and open it. The Scripts window appears.
- 3 From the Views pane, select Policy Usage. The list of software policies associated with the scripts appears in the Content pane.

### Viewing Script Version History

To view the version history of a script perform the following steps:


- 1 From the Navigation pane, select **Library > By Type > Scripts**.
- 2 From the Content pane, select the script and open it. The Scripts window appears.
- 3 From the Views pane, select Version History. The events associated with the script will display in the Content pane. You can view the script version, the script description, the user who created the script, and the time when the script was created. The  indicates the current version of the script. You can view the script



execution history such as who executed the script and when from the Server Explorer. See the *Opware® SAS User's Guide: Server Automation* for more information on server history.

- 4 To make any of the previous version of script current, select the script version and from the **Actions menu**, select **Set as Current Version** as shown in figure Figure 8-2.

Figure 8-2: Script Version History



## Locating Scripts in Folders

To locate a script in the folder hierarchy, perform the following steps:

- 1 From the Navigation pane, select **Library > By Type > Scripts**.
- 2 From the Content pane, select the script and then select **Locate in Folders** from the **Actions** menu. The folder hierarchy for the script appears in the Content pane.

## Exporting a Script

To download a script, perform the following steps:

- 1 From the Navigation pane, select **Library > By Type > Scripts**. The scripts appear in the Content pane.

Or

From the Navigation pane, select **Library > By Folder** and then select the folder which contains the script.

- 2 From the Content pane, select a script to export.
- 3 From the **Actions** menu, select **Export Script**. The Export Software window appears.
- 4 In the Browse window, specify the location for the script to be exported to.
- 5 Click **Export**.

## Renaming a Script

To rename a script perform the following steps:

- 1 From the Navigation pane, select **Library** ► **By Type** ► **Scripts**.
- 2 From the Content pane select the script, and then from the **Actions** menu select **Rename**.
- 3 Enter the new name for the script in the Content pane.

## Deleting a Script

To delete a script perform the following steps:

- 1 From the Navigation pane, select **Library** ► **By Type** ► **Scripts**.
- 2 From the Content pane select the script, and then from the **Actions** menu select **Delete**. The Confirmation window appears.
- 3 Click **Delete** to delete the script.

## Executing Scripts

In the SAS Client, you can execute scripts in the following ways:

- Execute a server script directly on servers or server groups and execute scripts in the Opware Global Shell. See “Running a Server Script (Saved Script or Ad-Hoc Script)” on page 484 and “Running an OGFS Script” on page 490 for more information.
- Add a script to a software policy and execute the script by attaching the software policy to the server and then remediating the server against the software policy. See “Installing Software Using a Software Policy” on page 449 for more information.

A software policy allows you to execute multiple scripts on a servers or server groups simultaneously, and execute a sequence of scripts on a server by specifying an install order in the software policy. See the *Opware® SAS Policy Setter's Guide* for information about software policy.



You must have a set of permissions to execute a script. To obtain these permissions, contact your Opware administrator. See the *Opware® SAS Administration Guide* for more information.

---

## Ways to Open the Run Script Window

The Run Script window allows to you execute a script on managed servers. In the SAS Client you can launch the Run Script window in the following ways:

- From the Device List
- From the Device Explorer
- From the Library

### From the Device List

**1** From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.

Or

From the Navigation pane, select **Devices** ► **Device Groups**. The device group list appears in the Content pane.

**2** From the Content pane, select a server or device group.

**3** From the **Actions** menu, select **Run Script**. The Run Script window appears.

### From the Device Explorer

**1** From the Navigation pane, select **Devices** ► **Servers** ► **All Managed Servers**. The server list appears in the Content pane.

**2** From the Content pane, select a server.

**3** From the Action menu, select **Open**. The Device Explorer appears.

**4** From the **Actions** menu, select **Run Script**. The Run Script window appears.

### From the Library

**1** From the Navigation pane, select **Library** ► **By Type** ► **Scripts**. The scripts list appears in the Content pane.

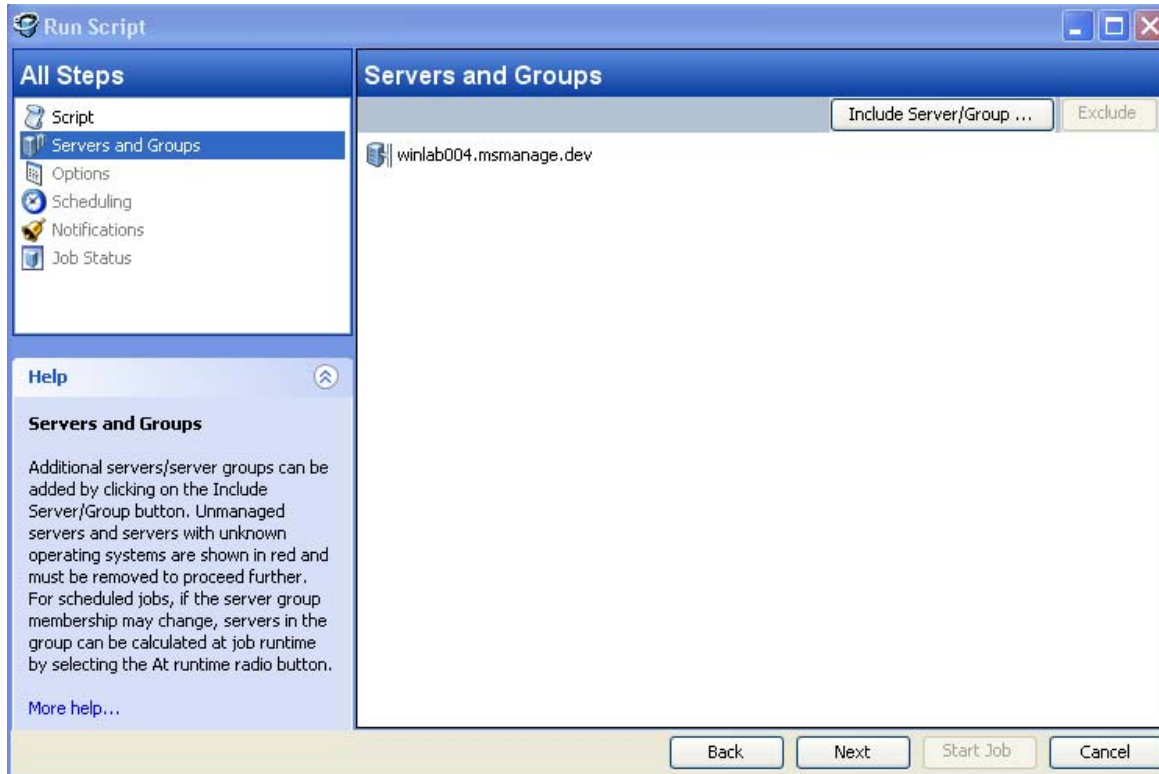
**2** From the Content pane, select a script.

**3** From the **Actions** menu, select **Run Script**. The Run Script window appears.

## Running a Server Script (Saved Script or Ad-Hoc Script)

The Run Script, as shown in Figure 8-3, allows you to run a script on managed servers and consists of the following steps:

Figure 8-3: Run Server Script Window



- Servers and Groups
- Script
- Options
- Scheduling
- Notification
- Job Status

See “Ways to Open the Run Script Window” on page 483 on how to access the Run Script window. If you access the Run Script window from the Device list or Device Explorer, the first step in the window is Script. If you access the Run Script window from the Library, the first step in the Run Script window is Servers and Groups.

## Servers and Groups

This step allows you to specify the servers or server groups for executing the script. In this step, you can add and remove servers or server groups from the list. For a server group you can also specify how the members of the group are determined.

If you choose the option Now, then the membership is determined based on the time when you made the selection. As a result the script is executed on the servers that were in the group when you selected the option. Changes to the group membership does not affect the list of servers on which the script will be executed.

If you choose the option Runtime, then the membership is determined when the script execution job is run. The script is executed on the servers present in the server group when the job is run. Changes to group membership is reflected in the list of servers when is script is executed.



---

To be able to select the Runtime option, the “Allow Run Refresh Jobs” permission is required. See the *Opware® SAS Administration Guide* for more information on permissions.

---

To select servers and groups perform the following steps :

- 1** Open the Run Script window from one of the methods described in “Ways to Open the Run Script Window” on page 483.
- 2** In the Run Script window, select the step Servers and Groups.
- 3** (Optional) Click **Include Server/Group** to add servers or server groups to the list or select a server or server group and click **Exclude** to remove servers from the list.
- 4** For a server group, in the Server Group Calculation field, select the option Now to execute the script on the servers that were in the group when you made the selection. Select the option Runtime to execute the script on the servers when the job is run.
- 5** Click **Next** to proceed to the Script step.

## Script

This step allows to select a saved script or define an ad-hoc script to be executed on managed servers. See “Types of Scripts” on page 472 for information on the script types.

### **Saved Script**

To select a saved script perform the following steps:

- 1** To select a saved script, select the option Select Saved Script.
- 2** From the Name drop-down list select the script or click **Select Script** to open the Select Script window. Select the script from the Select Script window.
- 3** The script properties such as version, type, location are displayed in the content pane. To view the contents on the script, click **View Script**. The contents of the script are displayed in the Run Script window.
- 4** Click **Next** to proceed to the Options step.

### **Ad-Hoc Script**

To define an ad-hoc script perform the following steps:

- 1** To select an ad-hoc script, select the option Define Ad-hoc Script.
- 2** (Windows only) From the Type drop-down list, select the script type.
- 3** Enter the contents of the script in the Script Contents field or click **Import Script File** to import a script.
- 4** Click **Next** to proceed to the Options step.

### **Options**

This step allows you to specify the runtime options and output options for executing a script. In this step you can specify whether to execute the script as root or Local System or as a specified user. You can also specify the script time-out value, any additional parameters for executing the script, and the output options for the script.

To specify the runtime and output options for a script perform the following steps:

- 1** In the Runtime User field select root (for Unix) or Local System (for Windows) to execute the script as root or local system. To execute the script as root or local system, you require the appropriate permissions. See the *Opware® SAS Administration Guide* for information about the permissions required for executing scripts.

Or

1. Select Name and enter user name and password to execute the script as a specified user. To execute the script simultaneously across multiple servers or server groups, you must use the same user name and password across all the servers.
2. (Windows only) Enter the domain name in the domain field.
- 2** In the Script timeout field enter the script timeout value in minutes. The time out value is the amount of time required for a script to complete execution activities on a server. If the script is not executed when the timeout value is reached, then the script is stopped by Opware SAS and a script error occurs. Select a timeout value greater than the time required for execution to complete.
- 3** In the Specify any needed parameters for this script execution field, enter any parameters if required.
- 4** In the Output Options, select Discard all script output to discard script output or else select Retain script output.
- 5** Select the output size of the script from the Size of the output to retain drop-down list.
- 6** Click **Next** to proceed to the Scheduling step.

### **Scheduling**

This step allows you to schedule the script execution job. You can choose to run the script execution job immediately, or on a specified date and time, or on a recurring basis.

To schedule a script execution job, perform the following steps:

- 1** In the Schedule Frequency section, choose to run the script once, daily, weekly, monthly, or on a custom schedule. Select any one of the following options:
  - **Once**: Choose this option to run the job immediately or only once at a specified date and time.
  - **Daily**: Choose this option to run the job on a daily basis at a specified time.
  - **Weekly**: Choose this option to specify the day or days of the week to run the job.
  - **Monthly**: Choose this option to specify the months to run the job, and the days of the month.
  - **Custom**: In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values.

- 2 In the Time and Duration section, for each type of schedule, specify the start time for the job. You must also specify the start date and end date for the job. The Time Zone is set according to the time zone set in your user profile.
- 3 Click **Next** to proceed to the Notifications step.

### **Notifications**

This step allows you to set email notifications to alert users on the success or failure of a job. You can also associate a Ticket ID with the job. This setting is optional.

To set email notifications, perform the following steps:

- 1 Click **Add Notifier**.
- 2 Enter the addresses in the Email Address of Recipient field.
- 3 To send email to the address if the job succeeds, select the checkbox On Success.
- 4 To send email if the job fails, select the select the checkbox On Failure.
- 5 Enter an ID to be associated with this job in the Ticket ID field.
- 6 Click **Next** to proceed to the Job Status step.

### **Job Status**

This step allows you to start the job, view the job progress, the job results, the script output for a managed server, and export the script output from all the servers.

Opsware SAS supports the following file formats for exporting script output results:

- A Zip file with folders for each managed server
- A Zip file containing no folders
- Consolidated raw text file
- Consolidated formatted text file
- Consolidated CSV file

You can also view jobs in the Jobs Log window of the SAS Client. See the *Opsware® SAS User's Guide: Server Automation* for information about Job Logs.

To start a job, perform the following steps:

- 1 To start the job, click **Start Job**.



If you selected Immediately in the Scheduling step, the job will begin now. If you scheduled the job for a later time, the job will run later. You can then view the job in the Jobs Log window of the SAS Client.

- 2 The job's progress information appears in the Job Status window. You can view the server on which the script was executed, the job status, and the exit code. If the exit code is zero, then it indicates that the script is executed successfully. If the exit code is non-zero, then it indicates an error during script execution.

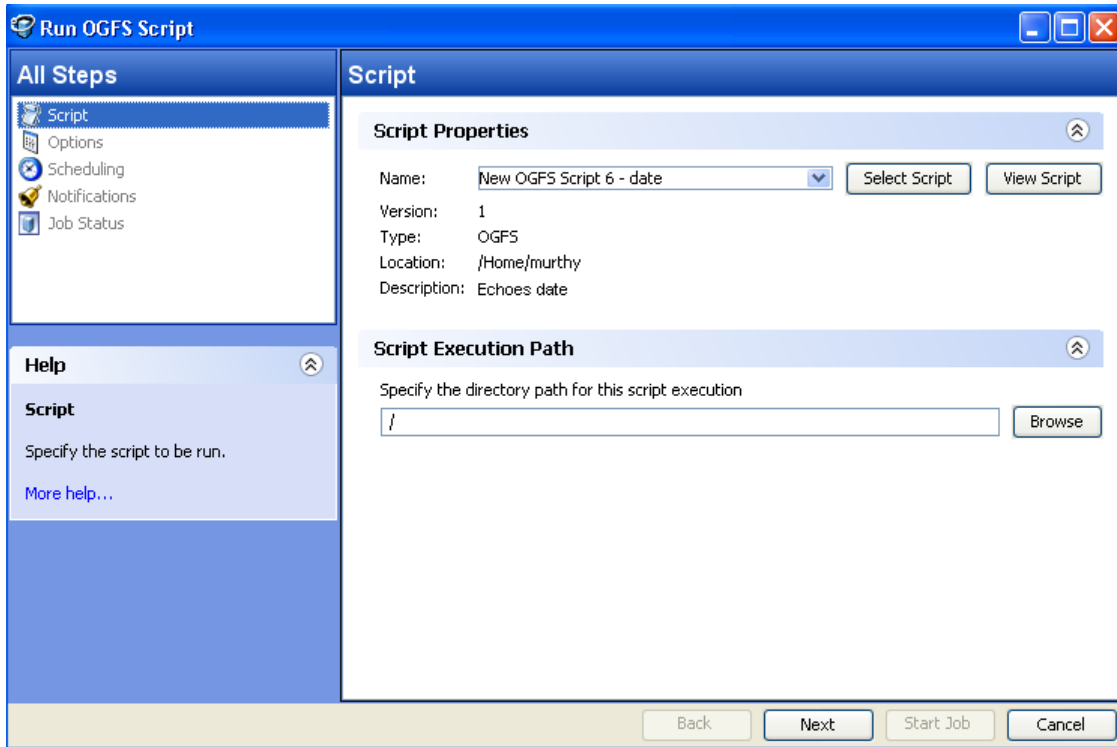
If the job status is displayed as Pending Approval, then the job is blocked until it is approved by a process that is external to Opsware SAS. See the *Opsware® SAS User's Guide: Server Automation* for information about job status.

- 3 (Optional) To view the script output from a managed server, select the managed server and script output appears below the table.
- 4 (Optional) To view the script output from all the managed servers, select the option Show output in table. The output for each server appears in the Output column in the table.
- 5 (Optional) To view the output for all the servers in separate columns, select the option Show output in table and enter the delimiter character in the Delimiter checkbox. The output for each server appears in separate columns in the table.
- 6 (Optional) To export the script output results, click **Export All Results**. In the Browse window specify the location and the file type and click **Export**.
- 7 Click **Close** to exit the Run Script window.

## Running an OGFS Script

The Run OGFS Script, as shown in Figure 8-4, allows you to run an OGFS script and consists of the following steps:

Figure 8-4: Run OGFS Script Window



- Script
- Options
- Scheduling
- Notification
- Job Status

### Script

This step allows to specify an OGFS script for execution.

To select an OGFS script perform the following steps:

- 1** From the Navigation pane, select **Library ► By Type ► Scripts**. The scripts list appears in the Content pane.  
  
Or  
  
From the Navigation pane, select **Library ► By Folder**. The folder hierarchy in the Library appears in the Content pane.
- 2** From the Content pane, select an OGFS script.
- 3** From the **Actions** menu, select **Run Script**. The Run OGFS Script window appears.
- 4** In the Script Properties section, select script from the Name drop-down list or click **Select Script** to open the Select Script window. Select the script from the Select Script window.
- 5** The script properties such as version, type, location, description are displayed in the content pane. To view the contents on the script, click **View Script**. The contents of the script are displayed in the Run OFGS Script window.
- 6** In the Script Execution Path section, enter the OGFS directory path for executing the script or click **Browse** to specify the directory path in the OGFS.
- 7** Click **Next** to proceed to the Options step.

### Options

This step allows you to specify the runtime options and output options for executing a script. In this step you can specify the script time-out value, any additional parameters for executing the script, and the output options for the script.

To specify the runtime and output options for a script perform the following steps:

- 8** In the Script timeout field enter the script timeout value in minutes. The time out value is the amount of time required for a script to complete execution activities on a server. If the script is not executed when the timeout value is reached, then the script is stopped by Opsware SAS and a script error occurs. Select a timeout value greater than the time required for execution to complete.
- 9** In the Specify any needed parameters for this script execution field, enter any parameters if required.
- 10** In the Output Options, select Discard all script output to discard script output or else select Retain script output.
- 11** Select the output size of the script from the Size of the output to retain drop-down list.

**12** Click **Next** to proceed to the Scheduling step.

### **Scheduling**

This step allows you to schedule the script execution job. You can choose to run the script execution job immediately, or on a specified date and time, or on a recurring basis.

To schedule a script execution job, perform the following steps:

- 1** In the Schedule Frequency section, choose to run the script once, daily, weekly, monthly, or on a custom schedule. Select any one of the following options:
  - **Once:** Choose this option to run the job immediately or only once at a specified date and time.
  - **Daily:** Choose this option to run the job on a daily basis at a specified time.
  - **Weekly:** Choose this option to specify the day or days of the week to run the job.
  - **Monthly:** Choose this option to specify the months to run the job, and the days of the month.
  - **Custom:** In the Custom Crontab string field, enter a string that indicates a time schedule. The crontab string can include serial (1,2,3,4) and range (1-5) values.
- 2** In the Time and Duration section, for each type of schedule, specify the start time for the job. You must also specify the start date and end date for the job. The Time Zone is set according to the time zone set in your user profile.
- 3** Click **Next** to proceed to the Notifications step.

### **Notifications**

This step allows you to set email notifications to alert users on the success or failure of a job. You can also associate a Ticket ID with the job. This setting is optional.

To set email notifications, perform the following steps:

- 1** Click **Add Notifier**.
- 2** Enter the addresses in the Email Address of Recipient field.
- 3** To send email to the address if the job succeeds, select the checkbox On Success.
- 4** To send email if the job fails, select the checkbox On Failure.
- 5** Enter an ID to be associated with this job in the Ticket ID field.
- 6** Click **Next** to proceed to the Job Status step.

## Job Status

This step allows you to start the job, view the job progress, view the job results, view the script output for a managed server, and export the script output from all the servers.

Opware SAS supports the following file formats for exporting script output results:

- A Zip file with folders for each managed server
- A Zip file containing no folders
- Consolidated raw text file
- Consolidated formatted text file
- Consolidated CSV file

You can also view jobs in the Jobs Log window of the SAS Client. See the *Opware® SAS User's Guide: Server Automation* for information about Job Logs.

To start a job, perform the following steps:

- 1** To start the job, click **Start Job**.

If you selected Immediately in the Scheduling step, the job will begin now. If you scheduled the job for a later time, the job will run later. You can then view the job in the Jobs Log window of the SAS Client.

- 2** The job's progress information appears in the Job Status window. You can view the server on which the script was executed, the job status, and the exit code. If the job status is displayed as Pending Approval, then the job is blocked until it is approved by a process that is external to Opware SAS. See the *Opware® SAS User's Guide: Server Automation* for information about Job Logs.
- 3** (Optional) To view the script output from a managed server, select the managed server and script output appears below the table.
- 4** (Optional) To view the script output from all the managed servers, select the option Show output in table. The output for each server appears in the Output column in the table.
- 5** (Optional) To view the output for all the servers in separate columns, select the option Show output in table and enter the delimiter character in the Delimiter checkbox. The output for each server appears in separate columns in the table.
- 6** (Optional) To export the script output results, click **Export All Results**. In the Browse window specify the location and the file type and click **Export**.

- 7 Click **Close** to exit the Run OGFS Script window.

# Chapter 9: 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)

Opware 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

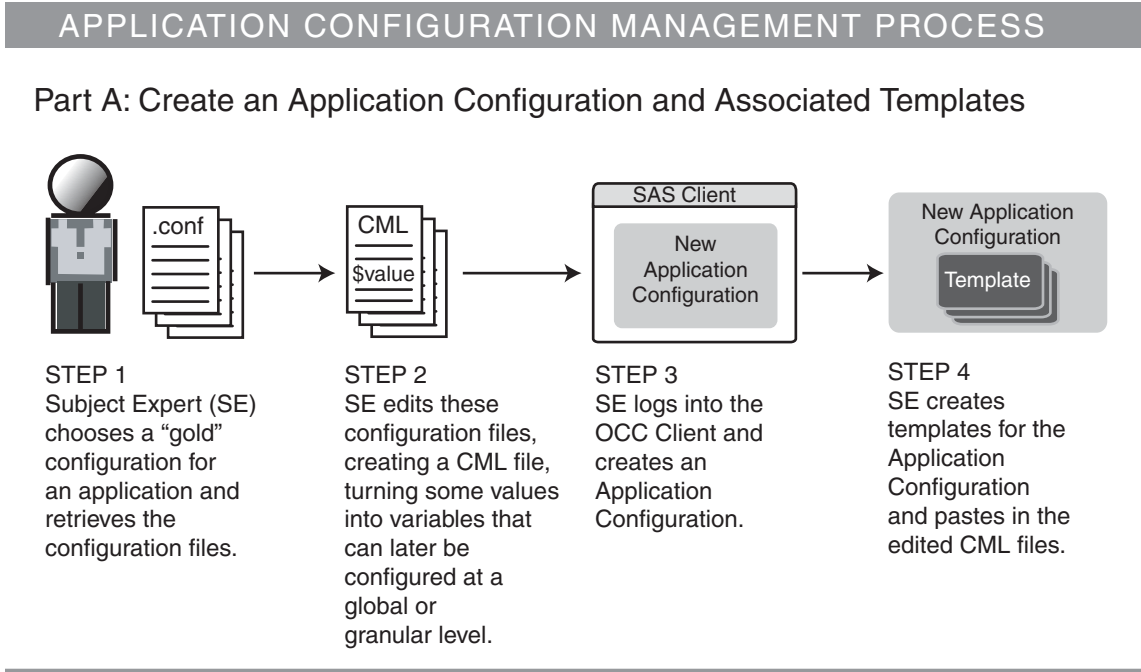
Using an Application Configuration enables you to manage and modify configuration files for applications on your Opware-managed servers. The process of using ACM follows these general steps:

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 SAS 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 group of servers) where you are managing application files.

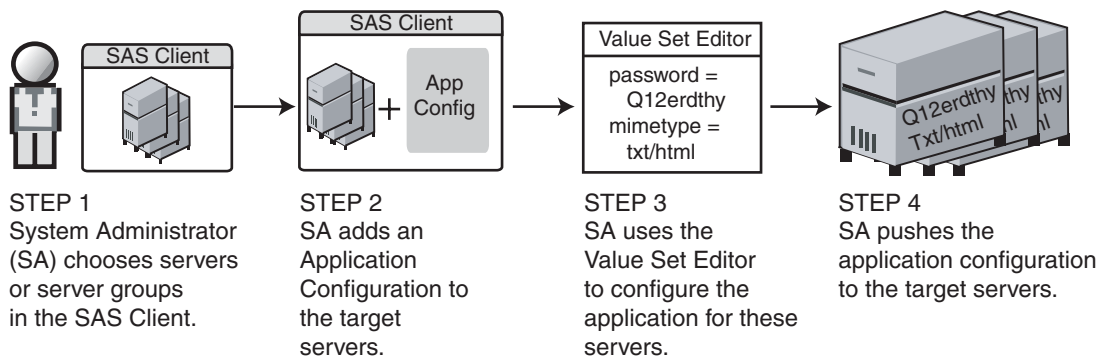


7. **Compare the actual configuration files with the configuration template:** You can easily compare a Application Configuration Template with the actual configuration file on the server and see if any changes have been made. This comparison shows manually changed configuration files or configuration values that have been changed, but not pushed.
8. **Push configuration 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 9-1: Application Configuration Creation and Usage Process



**Part B: Configure and Push Application Configurations to 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 SAS Client, you can edit the values in the configuration template and push those changes to the actual configuration file on the server.

Using Opware'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, Opware 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**

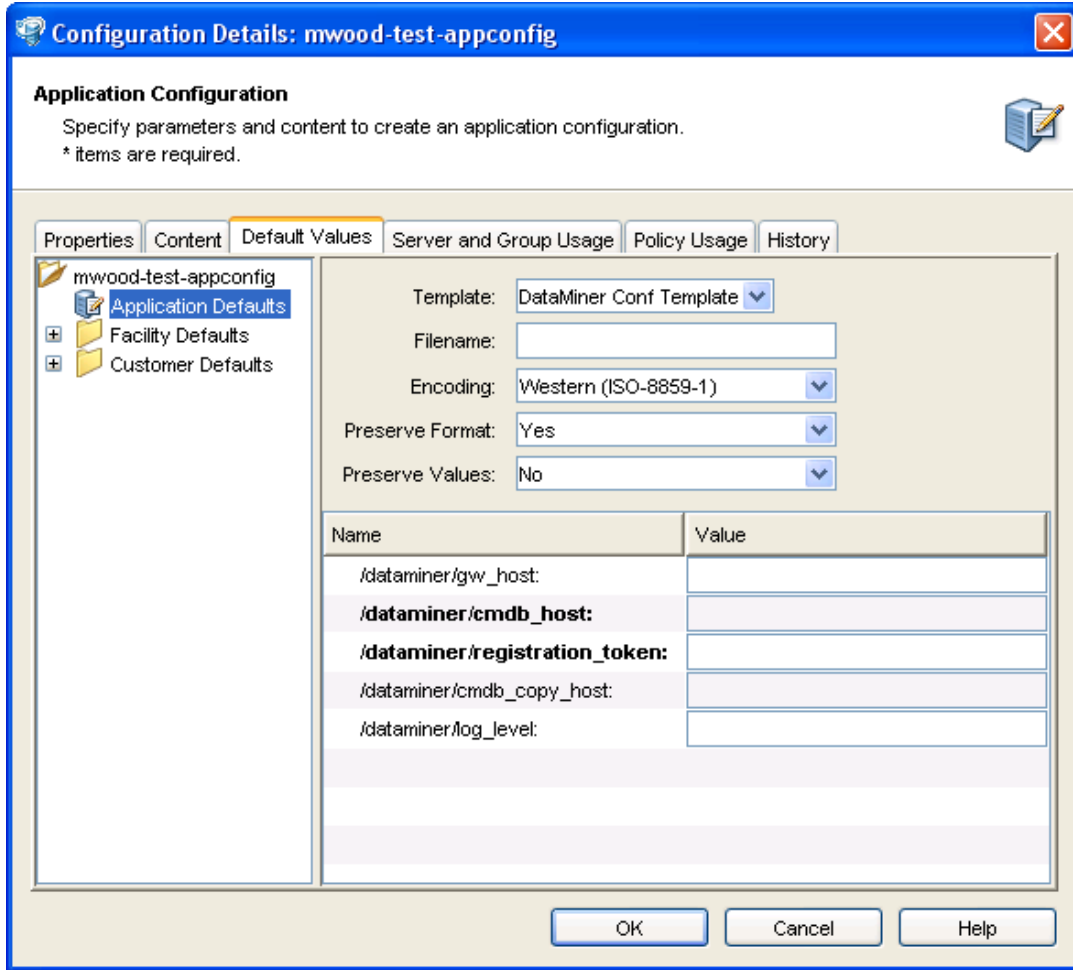
The Value Set Editor enables you to specify the values for each configuration file. Each entry inside a configuration file is represented inside the value set editor as an element, which consists of a name-value pair. The entire collection of elements in a configuration file is referred to as the configuration 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 SAS Client and double-clicking an Application Configuration.

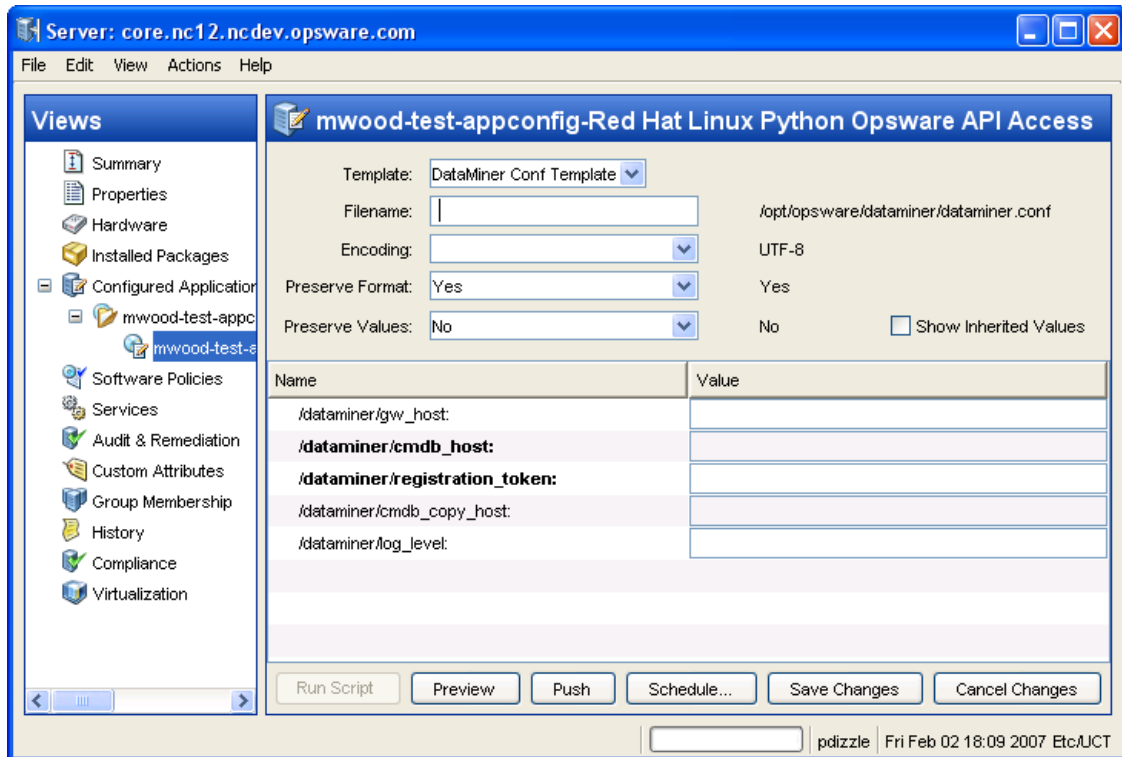
Figure 9-2: Application Configuration Default Values



Inside the Value Set Editor, elements that are required will appear in bold.

- **Device Explorer or Device Groups Explorer:** 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 9-3: Value Set Editor in the Device Explorer



The left side of the Device 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. (Note that UTF-16 encoding is not supported in the SAS Client.)
- **Preserve Format:** Choose this option if you want to both keep comments and preserve as much of the original ordering and spacing of the actual configuration file on the target server. The Application Configuration feature will attempt to preserve as much of the target source file as possible, but may not be able to preserve all comments and formatting. This options is also required if your Application Configuration uses the `@!partial-template@` CML tag. For more information on how to use CML, see the CML tutorial located in the *Opware® SAS Policy Setter's Guide*.
- **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 SAS 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 Device Explorer. Figure 9-2 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 groups of servers. 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.

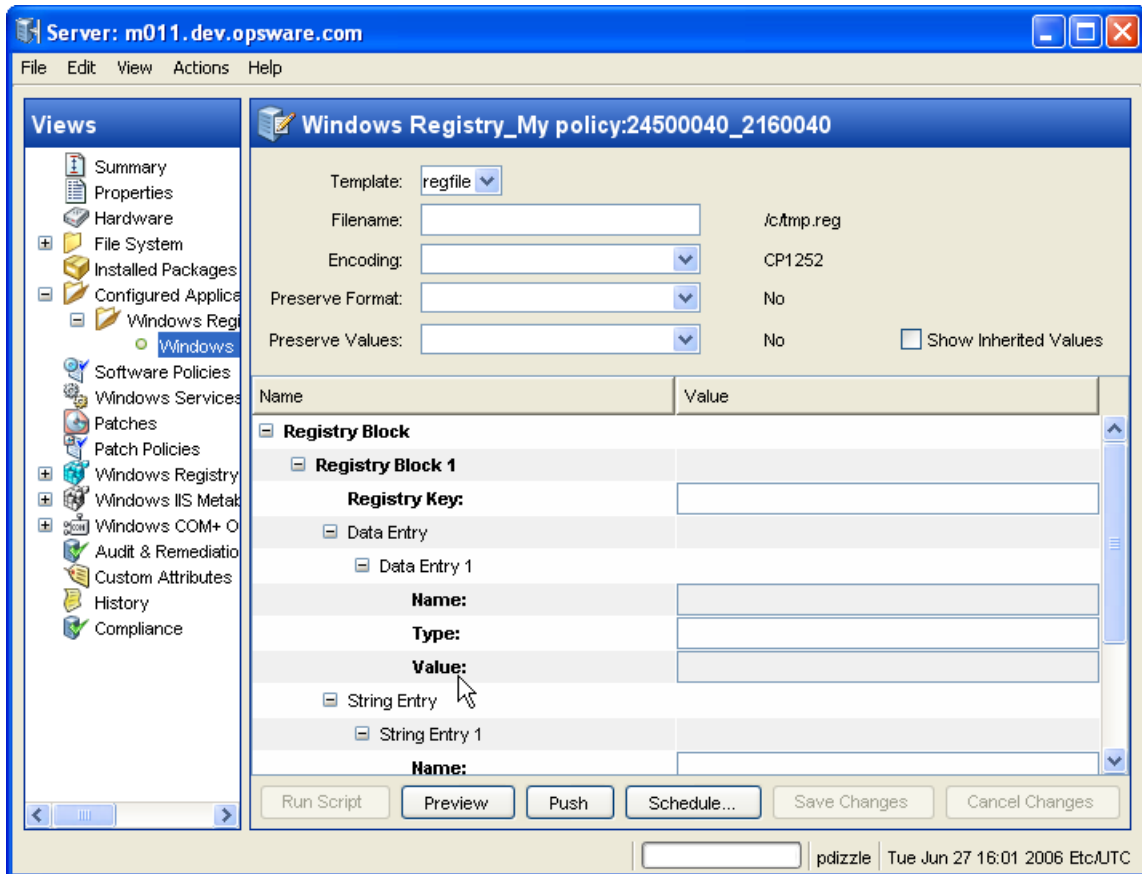


## Application Instance Values

From the Device Explorer (or Device 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 Device 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 9-4 shows the Application Configuration hierarchy at the server level.

Figure 9-4: Application Configuration Inheritance Hierarchy at the Server Level



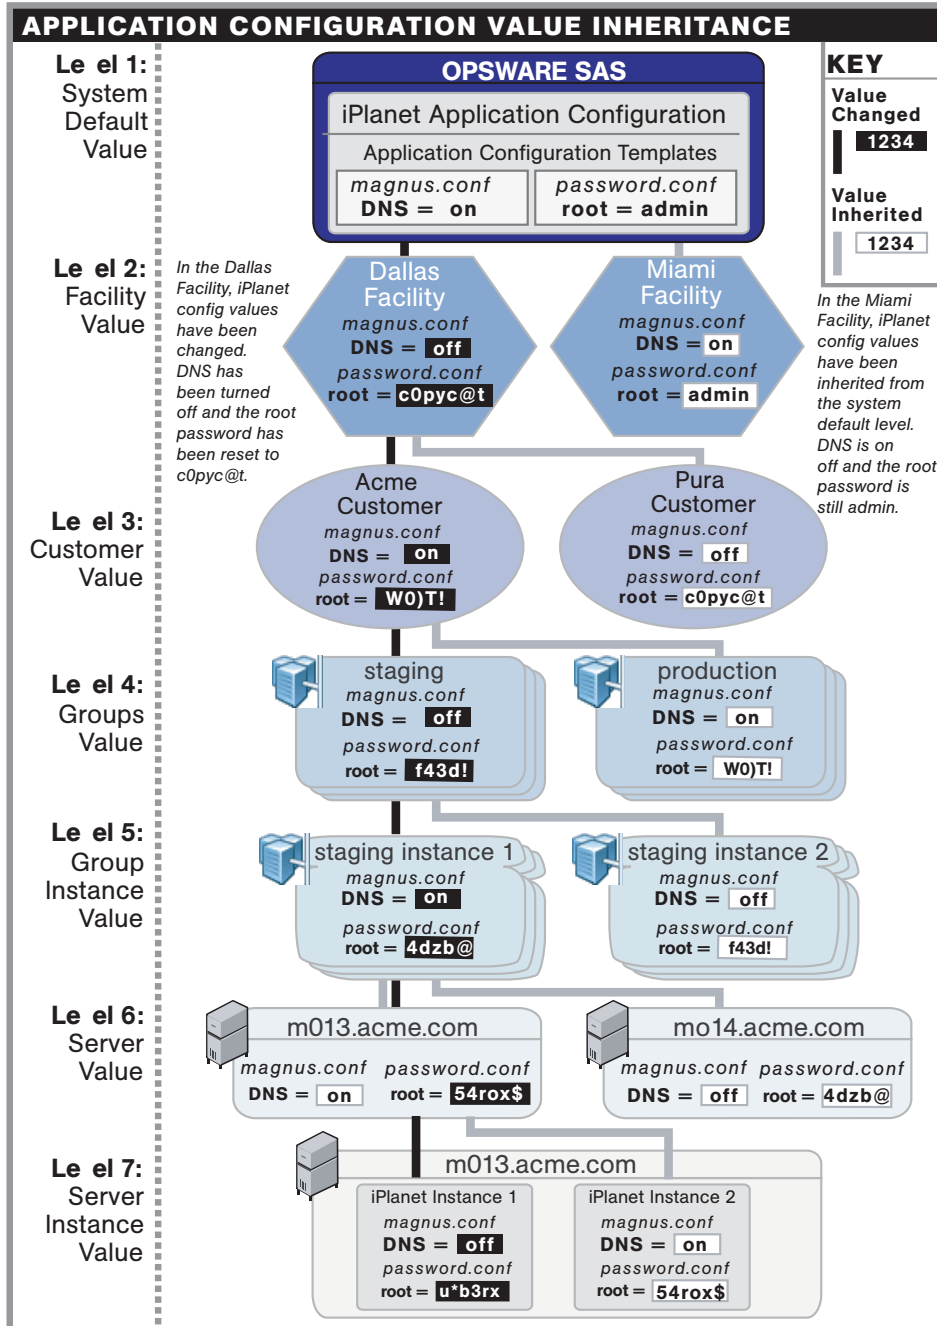
Application configuration inheritance on a managed server adheres to the following hierarchy:

- **Group of Servers:** This applies to all applications on all servers within the specific group of servers. Configuration values are inherited from the Application Configuration default values unless otherwise specified. (For example, if this group of servers belongs to a specific customer, it inherits the values of that customer.)
  - **Group of Servers Application Instance:** This applies to a specific instance of an application on all servers in the specific group of servers. 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 group of servers 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 group of servers the server belongs to (if it belongs to a group), and any other Application Configuration default values.

### Application Configuration Inheritance Visualized

Figure 9-5 illustrates how Application Configuration values are inherited.

Figure 9-5: Application Configuration Inheritance



## 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 to 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 prepended (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 517.

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 535.

## Using ACM

This section contains the following tasks:

- Creating an Application Configuration
- Creating a Configuration Template
- Searching for Application Configurations
- 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 Device Group
- Setting Application Configuration Values on a Server or Device 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
- Scanning Configuration Compliance
- Scheduling a Configuration Compliance Scan
- Restoring 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 SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 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 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 Opware'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 521 in this chapter for more information.

To create an Application Configuration Template, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 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**.

## Searching for Application Configurations

You can use the SAS Client Search tool to find Application Configurations and Application Configuration Templates in your facility. You can search for Application Configurations by name, by the operating system, and many other criteria.

To search for Application Configurations, perform the following steps:

- 1** From inside the SAS Client, make sure the search pane is activated by selecting **Search** from the **View** menu.
- 2** From the top drop-down list, select Application Configuration or Application Configuration Templates.
- 3** Click the green arrow button or ENTER to execute the search.
- 4** The results appear in the Contents pane.
- 5** If you want to extend your search criteria, add new criteria in the search parameters section at the top of the Contents pane. You can also save the search by clicking Save, or export the Search results to HTML or CSV.

## 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 508.

To view Application Configuration Template CML source, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 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 SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 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 SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 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.

UTF-16 encoding is not supported in the SAS Client.

To load a template file, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 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. (Note that UTF-16 encoding is not supported in the SAS Client.)
- 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 SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 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 might want to control templates application and script execution order.

For example, you might want to apply changes to certain configuration files before others. Or, you might 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 SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 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.)

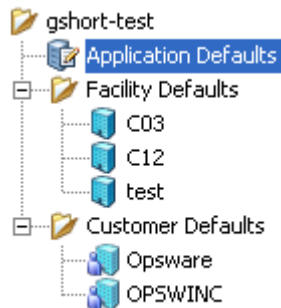
However, you can override the scope of an application configuration's default values by customer or facility. You can also edit specific instances of the application configuration to override the scope of an application configuration's default values. All elements that are required appear in bold font.

To set default values for an application configuration, perform the following steps:

- 1 Launch the SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 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 9-6 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 9-6: 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. (Note that UTF-16 encoding is not supported in the SAS Client.)
  - **Preserve Format:** Choose this option if you want to both keep comments and preserve as much of the original ordering and spacing of the actual configuration file on the target server. The Application Configuration feature will attempt to preserve as much of the target source file as possible, but may not be able to preserve all comments and formatting. This options is also required if your

Application Configuration uses the `@!partial-template@` CML tag. For more information on how to use CML, see the CML tutorial located in the *Opsware® SAS Policy Setter's Guide*.

- **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.
- 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 value sets for the Application Configuration, click **Save Changes**.

## Attaching an Application Configuration to a Server or Device 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 single server or public device group.

For an Application Configuration to manage an application on a server, it must be added to a server or group of servers. Once you add an Application Configuration to a server or group of servers, 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 device group.

---

## Attaching an Application Configuration to a Single Server

To attach an Application Configuration to a single server, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select **Devices** ► **All Managed Servers**.
- 2** From the Content pane, select a server.
- 3** From the **Actions** menu select **Open**. to open the Device Explorer. (Or, you can double-click the server's to open the Device Explorer.)
- 4** From inside the Device Explorer, in the Views pane select Configured Applications.
- 5** From the **Action** menu, select **Add Configuration**.
- 6** In the Select Application Configuration window, select an Application Configuration.

You can 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).


- 7** When you have selected an Application Configuration, click **OK**. The Application Configuration is attached to the server.

- 8** You can now set the Application Configuration's values. For more information on setting up the Application Configuration, see "Setting Application Configuration Values on a Server or Device Group" on page 528.

### ***Attaching an Application Configuration to a Device Group***

To attach an Application Configuration to a device group, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select **Devices** ► **Device Groups**.
- 2** From the Content pane, select a device group.
- 3** From **Actions** menu, select **Open** to open the Device Group Explorer. (Or, you can double-click the group to open the Device Group Explorer.)
- 4** From inside the Device Group Explorer, in the Views pane select Configured Applications.
- 5** From the **Action** menu, select **Add Configuration**.
- 6** 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).

- 7** When you have selected an Application Configuration, click **OK**. The Application Configuration is attached to the Device Group. For more information on setting up the Application Configuration, see "Setting Application Configuration Values on a Server or Device Group" on page 528.

## Setting Application Configuration Values on a Server or Device Group

Once an Application Configuration has been attached to a server or device group, you can edit its values. You can also override the default values set at the Application Configuration level. If the server (or device 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 503 in this chapter for more information.

For information on how to attach an Application Configuration to a server or device group, see “Attaching an Application Configuration to a Server or Device Group” on page 526

To set Application Configuration values on a server or group, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select either Device Groups or All Managed Servers.
- 3** Select a server or device group in the Content pane, and from the **Actions** menu select **Open**.
- 4** From inside the Device Explorer or Device Group Explorer, select Configured Applications.
- 5** Expand 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 523 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. (Note that UTF-16 encoding is not supported in the SAS Client.)
  - **Preserve Format:** Choose this option if you want to both keep comments and preserve as much of the original ordering and spacing of the actual configuration file on the target server. The Application Configuration feature will attempt to preserve as much of the target source file as possible, but may not be able to preserve all comments and formatting. This options is also required if your Application Configuration uses the `@!partial-template@` CML tag. For more information on how to use CML, see the CML tutorial located in the *Opware® SAS Policy Setter's Guide*.
  - **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.

- 7** (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.
- 8** (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**.
- 9** 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 SAS Client. From the Navigation pane, select Devices.
- 2** Select either Device Groups or All Managed Servers.
- 3** Select a server or device group in the Content pane, and from the **Actions** menu select **Open**.

You now see the Device Explorer (or Device Group Explorer), with the Installed Configurations tab selected. All Application Configurations that have been attached to the server (or group) will be displayed.

- 4** From the View pane, click the plus (+) symbol to expand Application Configuration folder and select an Application Configuration instance to edit.
- 5** From the Content pane, choose an Application Configuration Template from the Template drop-down list.
- 6** In the File name field, enter the absolute file name of the configuration file that contains the values that you want to import.
- 7** 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.
- 8** All of the values for the Application Configuration Template are replaced with the values from the actual configuration file.
- 9** 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 508.

---



---

If your push times out before the push succeeds (default is ten minutes), it could be that the default timeout value set in the Application Configuration is less than the time it takes to push the Application Configuration. See your Opware administrator for help in extending the duration allowed for an Application Configuration push to occur.

---

To push Application Configuration changes to a server or group, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select either Device Groups or All Managed Servers.
- 3** Select a server or device group in the Content pane, and from the **Actions** menu select **Open**.

You now see the Device Explorer (or Device Groups Explorer), 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.

- 4** From the Views pane of the Device Explorer (or Device Groups Explorer), select an Application Configuration instance to edit.
- 5** If you wish, make edits to the Application Configuration. (See “Setting Application Configuration Values on a Server or Device Group” on page 528 in this chapter for more information.)

- 6** 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.
- 7** 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 SAS Client. From the Navigation pane, select Devices.
- 2** Select either Device Groups or All Managed Servers.
- 3** Select a server or device group in the Content pane, and from the **Actions** menu select **Open**.

You now see the Device Explorer (or Device Groups Explorer), 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.

- 4** From the View pane, click the plus (+) symbol to expand Application Configuration folder and select an Application Configuration instance.
- 5** Click **Schedule**.
- 6** 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 Opware 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.
- 7** 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.
- 8** 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 successfully.
  - **On Failure:** Email addresses that will receive notifications of jobs that failed to complete.
- 9** In the Ticket Tracking section, enter a ticket ID from your own job trackins system here.
- 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 SAS Client. From the Navigation pane, select Library and then select the By Type tab.
- 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:
  - **Blue**: Indicates that information has been added, modified, or delete between the two templates
  - **Red**: Indicates that information is different between both templates.
  - **Black**: 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 SAS Client. From the Navigation pane, select Devices.
- 2** Select either All Managed Servers or Device Groups.
- 3** Select a server or device group in the Content pane, and from the **Actions** menu select **Open**.
- 4** From the Device Explorer (or Device Groups Explorer), select either an individual server or a group of servers.
- 5** Depending upon if you selected a server or a group of servers, from the Views pane of the explorer:
  - From the Device Explorer (individual server), select **Management Policies** > **Configured Applications**.

- From the Device Group Explorer, select **Configured Applications**.
- 6** From the Content pane, select the Installed Configurations tab. All Application Configurations that have been attached to the server (or group) will be displayed.
  - 7** The Installed Configurations tab will be selected. From the Views pane of the Device Explorer (or Device Groups Browser), select an Application Configuration instance.
  - 8** 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.
  - 9** 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.
  - 10** When you are finished viewing the differences, click **Close** to close the Comparison dialog box.



## Scanning Configuration Compliance

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 scan for configuration compliance 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 SAS Client:




For example, open the SAS Client and select the Servers feature icon. A list of all managed servers in your environment is displayed and you can see if any servers show the out of sync icon.

If a server shows this icon, scan for configuration compliance to find out which configuration files on the server are out of sync with the Application Configuration.

For information on how to schedule a configuration compliance scan, see “Scheduling a Configuration Compliance Scan” on page 538.

To scan for configuration compliance, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select either Device Groups or All Managed Servers. (If you selected a Device Groups, select the group in the Navigation pane to see the servers that belong to it.)
- 3** From the Content pane, select a server that shows the out of sync icon .
- 4** From the **Actions** menu, select **Scan > Configuration Compliance**. (You can also multiple-select and scan more than one out of sync server.)
- 5** You will be asked if you are sure you want to scan the Application Configuration on the selected managed server. Click **Yes** to run the scan.
- 6** The Job dialog box appears, showing the details of the scan. 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 scan 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.

- 7** To view the configuration compliance scan 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.
- 8** 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 535 in this chapter for more information

### Scheduling a Configuration Compliance Scan

You can schedule an configuration compliance scan to run a single time, or on a recurring schedule, such as daily, weekly, or monthly.

To schedule a configuration compliance scan, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select a server or group of servers from the Navigation pane, and then select a server from the Content pane.
- 3** From the **Actions** menu, select **Schedule Configuration Compliance Scan**.
- 4** In the Schedule Job window, 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.
- 5** 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.
- 6** 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 successfully.
  - **On Failure:** Email addresses that will receive notifications of jobs that failed to complete.
- 7** In the Ticket Tracking section, enter a ticket ID from your own job trackins system here.
- 8** When you have finished setting the parameters, click **OK**.

## Restoring to a Previous State

Every time you push an Application Configuration to a server, that push is saved in a configuration push history list. At any time, you can restore to a previous state of an Application Configuration push in this list.

To restore an Application Configuration to a previous state, perform the following steps:

- 1** Launch the SAS Client. From the Navigation pane, select Devices.
- 2** Select either Device Groups or All Managed Servers.
- 3** Select a server or device group in the Content pane, and from the **Actions** menu select **Open**.

You now see the Device Explorer (or Device 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.

- 4** 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.

---

- 5** To restore to a saved Application Configuration push, select a item in the list, and click **Restore**. 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 10: 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 SAS Client

Using server build policies, Opware SAS OS Provisioning allows you to provision operating systems onto bare-metal servers quickly and consistently. It also helps ensure that each operating system installed has the right build information.

OS Provisioning supports a large variety of hardware models from different manufacturers out of the box, and it can be configured to support additional hardware models. See the *Opware® SAS Policy Setter's Guide* for more information.



---

Before you can install operating systems on servers with the OS Provisioning feature, the operating systems must be defined and the OS media must be made available in Opware SAS. Additionally, OS installation profiles can be created in the Opware SAS Web Client and in the Opware Server Automation System Client. Please see the *Opware® SAS Policy Setter's Guide*.

---

## Supported Operating Systems for OS Provisioning

The OS Provisioning feature supports installation of the following versions of Linux, Sun Solaris, VMware ESX, and Microsoft Windows operating systems (architecture support is listed in parentheses):

- Red Hat Linux Enterprise Linux 2.1 AS (x86)
- Red Hat Linux Enterprise Linux 2.1 ES (x86)
- Red Hat Linux Enterprise Linux 2.1 WS (x86)
- Red Hat Linux Enterprise Linux 3 AS (x86, x86\_64, and ia64)
- Red Hat Linux Enterprise Linux 3 WS (x86, x86\_64, and ia64)
- Red Hat Linux Enterprise Linux 3 ES (x86, x86\_64, and ia64)
- Red Hat Linux Enterprise Linux 4 AS (x86 and x86\_64)
- Red Hat Linux Enterprise Linux 4 WS (x86 and x86\_64)
- Red Hat Linux Enterprise Linux 4 ES (x86 and x86\_64)
- Red Hat Enterprise Linux Server 5 (x86 and x86\_64)
- Red Hat Enterprise Linux Desktop 5 (x86 and x86\_64)
- Sun Solaris 8 (SPARC)
- Sun Solaris 9 (SPARC)
- Sun Solaris 10 (SPARC x86 and x86\_64)
- Fujitsu Solaris 8 (SPARC)
- Fujitsu Solaris 9 (SPARC)
- Fujitsu Solaris 10 (SPARC)
- Fujitsu Solaris 10 Update 2 (SPARC and Niagara)
- SUSE Linux Standard Server 8 (x86)
- SUSE Linux Enterprise Server 8 (x86)
- SUSE Linux Enterprise Server 9 (x86 and x86\_64)
- SUSE Linux Enterprise Server 10 Update 1 and Update 2
- SUSE Linux Enterprise Server 10 (x86\_64)
- VMware ESX 3

- VMware ESX 3 (x64)
- Windows 2000
- Windows 2003
- Windows 2003 x64
- Windows XP Professional



---

Opware OS Provisioning does not support Windows OS provisioning on Itanium-based systems.

---



---

In order to provision a bare metal SPARC sun4u server with any of the supported SPARC Solaris versions, the server must support Solaris 10 U3 (06/06).

---

The OS Provisioning feature works with a floppy disk for Windows via the DOS preinstallation environment, a CD-ROM for Windows via the WinPE preinstallation environment, a CD-ROM for Linux, or network booting for all supported operating systems. Non-network booting is not supported for Sun Solaris (SPARC and x86).



---

The OS Provisioning feature does not provision HP-UX or AIX operating systems. However, you can integrate Opware SAS with Network Installation Management (NIM) to provision AIX and Ignite-UX to provision HP-UX. See the *Opware<sup>®</sup> SAS Administration Guide* for more information on how to integrate the Opware SAS with HP-UX and AIX OS provisioning systems.

---

## OS Provisioning Overview

In Opware SAS, OS Provisioning is installation-based instead of image-based. The OS Provisioning feature uses Red Hat Linux Kickstart (used for Linux and ESX), Sun Solaris JumpStart, and Microsoft Windows unattended installation to install operating systems on servers.

The OS Provisioning feature is fully integrated with Opware SAS; you 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 helps you install operating systems on servers in the following ways:

- Each OS installation profile 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 the operating systems that 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 groups of servers. To obtain these permissions, contact your Opsware administrator. For more information, see the Permissions Reference appendix in the *Opsware® SAS Administration Guide*.

---

### **Server Lifecycle for OS Provisioning**

Opsware SAS enables multiple teams to work together and provision servers. The OS Provisioning feature allows IT teams to separate the tasks of readying servers for provisioning (such as racking servers or 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 in the Opsware SAS Web Client or in the Unprovisioned server list in the SAS Client – 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.

During OS provisioning, servers progress through Planned, Unprovisioned, Available, Installing OS, and Managed states. See Table 10-1 for more information on these lifecycle values.



Table 10-1: Opware SAS Lifecycle Values for Servers

OPSWARE LIFECYCLE VALUE	DESCRIPTION
<b>Server Pool Values</b>	
Planned	<p>Indicates that a device record has been created for the server, but an Opware OS Build Agent has not yet been installed. (The OS Build Agent is a small agent that can run in the memory of the bare metal server.)</p> <p>Servers in this stage cannot be provisioned until an OS Build agent is installed.</p> <p>For more information on this unprovisioned lifecycle value, see your Opware Administrator.</p>
Available	<p>Indicates a server on which the OS Build Agent was installed and is running, but that does not have an OS installed.</p> <p>See “Installing OS Build Agents” on page 557 in this chapter for more information.</p>
Installing OS	<p>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 Managed Server list.</p>
Build Failed	<p>Indicates a server where 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 seven days before Opware SAS deletes the entry.</p> <p>See “Recovering when an OS Build Agent Fails to Install” on page 558 in this chapter for more information.</p>
<b>Managed Server Values</b>	
Managed	<p>Indicates a server that Opware SAS is managing. Opware SAS performs reachability checks for managed servers.</p> <p>After a server reaches this lifecycle state, the entry for the server moves from the Server Pool list to the Managed Servers list.</p>
Deactivated	<p>Indicates a server that was previously managed by Opware SAS, but is no longer managed by Opware SAS. However, the server’s history still exists in Opware SAS. Deactivated servers are not reachable.</p>

## OS Provisioning

This section provides information about the OS provisioning process within Opware SAS and contains the following topics:

- Overview of OS Provisioning
- Solaris OS Provisioning
- Windows OS Provisioning
- Network Setup for OS Provisioning

### Overview of OS Provisioning

The process for provisioning new servers on supported operating systems includes the following steps:

**1** A system administrator unpacks a server, mounts it in a rack, and attaches the server to power and a network that can reach Opware SAS.

**2** The system administrator prepares the hardware for OS provisioning.

See “Hardware Preparation” on page 548 in this chapter for more information.

**3** If necessary, the system administrator inserts a bootable floppy or CD provided with Opware 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 servers are capable of booting over a network.)

See “New Server Booting” on page 550 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 Opware OS Build Manager responds to this network boot request by delivering the Opware 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 Opware OS Build Agent is on the bootable floppy or CD.)

The Opware 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 Opware OS Build Manager.

- 5** In the SAS Web Client, 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 557 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 installs the OS or a complete server baseline on the server at that time or schedules the installation for some time in the future.

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 *Opsware® SAS User's Guide: Server Automation* for more information.

Additionally, the system administrator may choose to configure servers running Red Hat Linux or Sun Solaris operating systems by using the OS Provisioning feature.

See “Reprovisioning a Managed Server” on page 571 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 installation profile 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 Policy Setter's Guide* for more information on the Solaris build process.

### **Windows OS Provisioning**

In the OS Provisioning feature, system administrators can perform unattended, scripted installations as well as WinPE-based image installations of Windows NT, Windows 2000, Windows 2003, and Windows XP Professional 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 known hardware 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.

#### *WinPE Memory Requirements*

In order to perform PXE booting of a VMWare ESX Windows 2003 x86 or x86\_64 VM using WinPE, the minimum required RAM is 512MB (higher than the VMWare recommended RAM minimum).

See the *Opware® SAS Policy Setter's Guide* for more information on the Windows build process.

### **Network Setup for OS Provisioning**

In order for OS Provisioning to work correctly and perform properly, it is essential that you correctly configure any network switch ports used for provisioning. These switch ports must have PortFast mode enabled and must be set for speed/duplex auto-negotiation. While provisioning using manually configured interface speed and duplex settings is possible for the Solaris and RedHat Linux-based boot images, Opware recommends using auto-negotiation as it has been found to work the most consistently.

Opware also recommends that the OS Build Agent contact the OS Provisioning Build Manager using the IP address as opposed to a DNS name. If you must use DNS names, but you must specify the DNS name during the Opware Installer interview and save it in the Response File (`boot_server.buildmgr_host` parameter). You must also configure DNS properly resolve that name to clients being provisioned. The hostnames of all OS media servers must also be resolvable in DNS.

### **Hardware Preparation**

Before you use OS Provisioning to install an OS on a server, the server must meet certain requirements, or the hardware must be prepared in certain ways.

### **Windows Hardware Preparation Requirements**

Before you install Windows on a server, you need to prepare the hardware by performing the following tasks:

- If the hardware has a RAID controller, you may have to extend the Windows OS media distribution based on vendor-specific requirements. The Microsoft Windows OS media might not (depending on the version of Windows) include the necessary drivers for many RAID controllers. Also, certain newer types of SATA controllers may require additional drivers.
- When using the DOS-based PXE or floppy boot images to install the Windows OS, you need to create a FAT16 or FAT32 partition on the primary boot (hard) drive to install the Windows OS on. The boot images contain the necessary functionality to create the required partition.
- If you use the WinPE-based PXE or CD-ROM boot images to install the Windows OS, disk partitioning is performed during the OS installation. You can control the disk partitioning by editing the OS installation profile in Opsware. (For more information, see the *Opsware<sup>®</sup> SAS Policy Setter's Guide*.)
- If you use the WinPE-based PXE or CD-ROM boot images provided by Opsware, and you are using a RAID or SATA controller, you might need to supply a build customization. This will enable you to load the necessary drivers before the OS installation can commence. It is unnecessary to do the same in a DOS-based installation environment, since all disk access is done via the BIOS until the Windows installation is in progress. For more information, see the *Opsware<sup>®</sup> SAS Policy Setter's Guide*.

### **Sun Solaris Hardware Preparation Requirements**

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 and VMware ESX Hardware Preparation Requirements**

Before you install Linux on a server, you need to prepare the hardware by configuring valid, logical drives for RAID.

VMware ESX hardware requirements are the same as Linux.



If you plan to use OS provisioning, you must also change the configuration of the managed switch for Redhat Linux, and enable PortFast on the managed switch. If this isn't done, when the Redhat Linux installer uses NFS to mount the media, the DHCP request might time out. (This problem is fixed in the packages listed in the advisory RHEA-2004:518-06.)

---

## 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 (DOS), Linux, or VMware ESX Server with PXE
- Booting a Windows Server with PXE Using WinPE
- Booting a Solaris Server Over the Network
- Installing OS Build Agents
- Verifying Installation of an OS Build Agent
- Recovering when an OS Build Agent Fails to Install

### Bootting 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 booting.

For Windows and Linux or VMware ESX servers, the Opsware Boot Floppy and CD respectively contain a small operating system, network drivers, the 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, but you cannot boot new Solaris servers using a floppy or CD.



---

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 preparing a server for provisioning from the task of provisioning the server with an OS. This de-coupling of tasks is made possible by 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 (DOS), Linux, or VMware ESX Server with PXE**

The following instructions show you how to boot a Windows (DOS), Linux, or VMware ESX Server with PXE.

For information on how to boot a server with WinPE, see “Booting a Windows Server with PXE Using WinPE” on page 554. For more information about hardware support, see the *Opsware<sup>®</sup> SAS Policy Setter's Guide*.



---

When booting a Linux, VMware ESX, or Windows server by using PXE, the DHCP relay must be running on the router of the build network for PXE to function properly. Alternatively, if the build script is plugged directly into the boot server providing DHCP service, a DHCP relay is not necessary.

---

- 1** After you mount the new server in a rack and connect it to the Opware 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.
- 3** Choose an Opware boot image by entering the appropriate text at the boot prompt.

```
windows    - Windows Build Agent (DOS 7.01)
undi       - Windows Build Agent (DOS 7.01 + UNDI)
winpe      - Windows Build Agent (WINPE based)
win-bcom   - Windows Build Agent (DOS 7.01 w/Broadcom driver
v9.07)
linux      - Linux Build Agent (RHEL 3.0-based)
linux4     - Linux Build Agent (RHEL 4.0-based)
solaris    - Solaris x86 Build Agent
localdisk  - Normal boot from localdisk (default after 10
second
```



---

If you are booting an VMware ESX server, select one of the Linux options. If you are provisioning a machine into NT 4, at the boot prompt (DOS menu), enter "windows-old" to continue. This option is hidden on the displayed menu, but is still accessible by typing it in yourself.

---

If you are booting a Windows server, the type of hardware being provisioned determines the version of the DOS Windows OS Build Agent. The images for the DOS Windows OS Build Agents vary in terms of the memory management software, disk partitioning capabilities, and network drivers.

If an incompatible boot image is selected for the hardware, an error message might appear at the console during the provisioning process; for example, it might appear when the Windows OS Build Agent is booting and DOS is loading or it might appear later in the process when the Windows Installer is running.

(For more information on WinPE booting, see "Bootting a Windows Server with PXE Using WinPE" on page 554).



See Table 10-2 for the differences between images for the Windows OS Build Agents.

Table 10-2: Differences Between Images for the Windows OS Build Agents

BOOT IMAGE	NETWORK DRIVERS	PREINSTALLATION ENVIRONMENT	DISK PARTITIONING CAPABILITIES
windows	Native DOS	DOS 7.0.1	FAT16 or FAT32
undi	UNDI	DOS 7.0.1	FAT16 or FAT32
winPE	Windows XP/2003/ Vista drivers	WinPE 2.0	NTFS or FAT32
winbcom	Native DOS with Broadcom v9.07	DOS 7.0.1	FAT16 or FAT32

For Windows, if you select a DOS option to boot the server, an additional set of Opsware SAS menus appear on the console. This enables you to partition the hardware disk. (If you boot using WinPE, you can specify disk partitioning later, in the OS Installation Profile properties, which will create the partitions when the OS is installed.)

If you do not select an option within 10 seconds, the server defaults to booting from the local disk. To stop the default selection, type something (other than ENTER) at the command line.

- 4** 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 SAS Web Client.

- 5** (Optional) Record the MAC address and/or the serial number of the server so that you can locate the server in the Server Pool list in the SAS Web Client or in the Unprovisioned Servers list in the SAS Client.

You should verify that the newly racked server shows up in the SAS Client Unprovisioned Servers or SAS Web Client Server Pool, and is ready to hand off for OS installation.

See “Verifying Installation of an OS Build Agent” on page 557 in this chapter for more information.

### **Booting a Windows Server with PXE Using WinPE**

Opware OS Provisioning now supports booting a bare metal server with PXE into a WinPE preinstallation environment. You can choose between either a WinPE x86 32 bit environment or a WinPE x64 64 bit environment.

WinPE provides greater flexibility than DOS, because it does not require that you format your hard drive during the boot process. You can define the disk partition configuration later in an OS installation profile, when you create or edit the OS installation profile.

For more information on how to create and edit an OS installation profile, see “Create an OS Installation Profile” on page 559.

WinPE also allows WIM-based image installation, as an alternative to unattended Windows installations.



---

When booting a Windows server by using PXE, the DHCP relay must be running on the router of the build network for PXE to function properly. Alternatively, if the build script is plugged directly into the boot server providing DHCP service, a DHCP relay is not necessary.

---

To boot a bare metal server with PXE into a WinPE preinstallation environment, perform the following steps:

- 1** Mount the new server in a rack and connect it to the Opware build network.
- 2** Set up the server to boot by using PXE.

See the hardware vendor’s documentation on how to prepare a server to boot using PXE.

- 3** Power on the server and select the option to boot the server with PXE.

The Opware SAS menu appears and prompts you to select the type of OS Build Agent to install on the server.

All of the Opware boot image options are displayed:

```
Windows    - Windows Build Agent (DOS 7.01)
undi       - Windows Build Agent (DOS 7.01 + UNDI)
```

```
winpe      - Windows Build Agent (WINPE based)
win-bcom   - Windows Build Agent (DOS 7.01 w/Broadcom driver
v9.07)
linux      - Linux Build Agent (RHEL 3.0-based)
linux4     - Linux Build Agent (RHEL 4.0-based)
solaris    - Solaris x86 Build Agent
localdisk  - Normal boot from localdisk (default after 10
second
```

- 4** At the boot prompt enter:

```
winpe
```



---

If you do not select an option after 10 seconds, the server defaults to booting from the local disk. To prevent the default selection, type in the command line.

---

- 5** A new menu displays the option to boot a WinPE x86 32 bit environment or a Windows x64 64 bit environment. Make a selection by using the arrow keys to highlight your choice, and then press ENTER.

The server will now be booted with the WinPE preinstallation environment. This may take a few minutes to complete, depending upon the speed of the network and the machine.

Once the booting has finished, a new window will appear indicating that the server has had an Opware Build Agent installed and registered with the Opware core.

- 6** (Optional) Record the MAC address and/or serial number of the server so that you can locate the server in the Server Pool list in the SAS Web Client or in the Unprovisioned Servers list in the SAS Client.
- 7** Verify that the newly racked server shows up in the SAS Client Unprovisioned Servers, or SAS Web Client Server Pool, and is ready for OS installation. See “Verifying Installation of an OS Build Agent” on page 557 in this chapter for more information.

## Booting a Solaris Server Over the Network

When Opware SAS is installed, the OS Provisioning feature is configured so that the Opware Boot Server listens for broadcast requests from new servers and responds 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 Opware DHCP server on the Opware core network. If the installation client is running on a different network than the Opware 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 Opware 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 *Opware® SAS Policy Setter's Guide* for more information.

## Ways that the OS Build Agent Locates the Opware Build Manager

For Solaris OS provisioning, the JumpStart build script runs the OS Build Agent, which contacts the Opware Build Manager (via the Agent Gateway in the core). The Solaris `begin` script attempts to locate the Opware Build Manager in the following ways:

- By using information that the Opware 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 Opware 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
```

## Installing OS Build Agents

You can install an OS Build Agent on a server by booting the server with PXE or an Opware Boot Image (Windows, Linux, or VMware ESX), or by using the network (Solaris). After a successful installation, the server appears in the Server Pool list.

You should verify that the newly racked server shows up SAS Client Unprovisioned Servers, or SAS Web Client Server Pool, and is ready to hand off for OS installation.

The SAS Client's Unprovisioned Servers list and the SAS Web Client Server Pool list display the servers that have registered their presence with Opware SAS but do not yet have the target OS installed.

You can start the OS installation process in either one of two ways:

- From the SAS Client's Unprovisioned Servers list, right click on the server in the content pane, and choose Run OS Sequence. Please See "OS Installation with the SAS Client" on page 559 for details.
- From the SAS Web Client Server Pool, select the server and click **Install OS**. This option is only available for SAS version 6.1 cores and later.

## Verifying Installation of an OS Build Agent

Perform the following steps to verify the installation of an OS build agent:

- 1 Log into the SAS Web Client.

- From the Navigation pane, select Servers ► Server Pool. The Server Pool page appears, as Figure 10-1 shows.

Figure 10-1: Server Pool List in the SAS Web Client

The following servers have registered their presence with Opware but do not have a full operating system installed.

All Manufacturers		All Models		Update						
Delete...		Install OS...								1 Total
	Name	MAC Address	Manufacturer	Model	Reported OS	Registered	Lifecycle	Facility	Customer	
<input type="checkbox"/>	m101.tr3.opsware.com	00:11:43:CE:19:4A	DELL COMPUTER CORPORATION	POWEREDGE 750	DOS	05-25-2005	Available	TR3	Not Assigned	

- (Optional) From the drop-down lists, select the manufacturer, model, or facility of the server and click **Update**.
- For Intel x86 and Sun SPARC processor-based servers, locate the MAC address and Host ID of the server that you just booted.

The Lifecycle column indicates the progress or success of the OS Build Agent installation. If the OS Build Agent was successfully installed, the Lifecycle column indicates that the server is available for OS provisioning.

See “Server Lifecycle for OS Provisioning” on page 544 in this chapter for more information.

To obtain information on a server in the SAS Web client, click on the server name. If you are viewing an unprovisioned server in the Unprovisioned Servers list in the SAS Client, double-click on the server to open the Device Explorer. This will display detailed 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 Opware 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 Lifecycle 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 Opware administrator for troubleshooting assistance.

## OS Installation with the SAS Client

This section describes how to install an operating system on an unprovisioned server using the SAS Client.



---

For information on how to set up OS provisioning, see the *Opware<sup>®</sup> SAS Policy Setter's Guide*.

---

In order to install an OS and provision a server using the SAS Client, you must create, define, and run an OS sequence. An OS sequence defines what to install on an unprovisioned server, including OS build information from the OS installation profile, selected software and patch policies, and remediation settings. An OS sequence represents a server build policy, and it defines how a server should be provisioned, affecting its software and operating systems. When the OS sequence is defined, it can be used to provision additional servers with the same OS and software.

This section explains how to:

- **Create an OS Installation Profile:** Define the OS, configuration or response file, build customization scripts, OS media, customer association, and packages.
- **Create an OS Sequence:** Choose the OS installation profile, software policies, patch policies, and remediation policies.
- **Select Servers in the Unprovisioned Servers List:** Choose the server(s) that you would like to provision.
- **Run an OS Sequence:** Launch the OS sequence to provision the selected unprovisioned server(s).

### Create an OS Installation Profile

An OS installation profile defines all necessary parameters of an OS, including the OS type and version, the OS Media Resource Locator (MRL), the configuration or response file, the build customization script, and the packages related to the OS installation.

For information about creating an OS installation profile, see the *Opware<sup>®</sup> SAS Policy Setter's Guide*.

## Create an OS Sequence

An OS sequence defines what to install on an unprovisioned server, including OS build information from the installation profile, selected application and patch policies, and the target servers you want to install the OS on to.



---

When you create an OS sequence, it is saved into the Folder list in the Library. You must have permissions to the folder where you want to save the OS sequence. For more information on how folder permissions work, see User and Group Setup in the *Opsware® SAS Administration Guide*.

---

## Elements of an OS Sequence

An OS sequence consists of the following components that must be configured before you run the OS sequence:

- **Properties:** Allows you to name the OS sequence and choose a location to save it in a library folder. You must have permissions to write to the folder where you save the OS sequence, otherwise you will be unable to save it in the selected location in the library.
- **Install OS:** Allows you to choose an OS installation profile. If the OS installation profile already has a customer associated with it, you will be unable to select a customer for the OS sequence. If the OS installation profile does not have a customer associated with it, then you can select one here. Once you choose a customer, then all servers on which you install the OS using this OS sequence will be associated with that customer.
- **Attach Software Policies:** Allows you to add a software policy to the OS sequence. When the OS sequence is run, if the remediate option is enabled (in Remediate Policies), then all the software in the software policy will be installed on the server during OS installation. If the remediation option is disabled, then none of the software will be installed on the server. In order to perform OS provisioning with remediation, you must have at minimum read access to all server module policies.

The software policies that you can attach to an OS sequence are restricted by the OS type. You can only attach software policies when their OS matches the OS installation profile chosen for the OS sequence.

For more information on software policies, see Chapter 7, “Software Management”.



- **Attach Patch Policies:** Allows you to select a patch policy to attach to the OS sequence. When run OS sequence is run, if the remediate option is enabled (in Remediate Policies), then all the patches in the patch policy will be installed on the server. If the remediate option is disabled, then none of the patches will be installed on the server.

Attach Patch Policies is available only for Windows OS Sequences.

For more information Chapter 5, "Patch Management for Windows".

- **Attach Device Group:** Allows you to select a device group (group of servers) for a the server once the OS sequence has been run. You can select any public static group to attach to the OS sequence. Also, a group of servers can have software and patch policies associated with it. If you enable remediation in the OS sequence (in Remediate Policies), then all software and patches associated with the group of servers will also be installed on the server when you run the OS sequence. If you disable remediation, then none of the software or patches in the policies attached to the group of servers will be installed on the server.

For information on groups of servers, see Server Management in the *Opware® SAS User's Guide: Server Automation*.

- **Remediate Polices:** Allows you to choose to enable or disable remediation when the server is provisioned with the OS sequence. The Default is **Disabled**.

When remediation is disabled, running an OS sequence installs the OS however no policies in the OS sequence are remediated –that is, no software or patches in any of the policies attached to the OS sequence are installed when the sequence is run.

If you enable remediation, then all software and patches in all policies attached to the server will be installed when the OS sequence is run. This is also true for any policies attached to the group of servers selected for the OS sequence. You can also set reboot and pre and post installation script options.



In order to perform OS provisioning with remediation, you must have at minimum read access to all server module policies.

---

## Create an OS Sequence

To create an OS Sequence, perform the following steps:

- 1** In the SAS Client, from the Navigation pane, select Library and then select OS Sequences.
- 2** Choose an OS folder.
- 3** From the **Actions** menu, select **Create New**.
- 4** In the Views pane of the OS Sequence window, select Properties and enter a name for the OS sequence.
- 5** Click **Change** in the Content pane to choose a location in the folder library to save the OS sequence. You must have permissions to write to the folder where you save the OS sequence.
- 6** From the Views pane, click **Tasks** then **Install OS** to choose an OS installation definition.
- 7** If the OS installation profile does not have a customer associated with it, then select a customer from the Assign Customer drop-down list. If the OS installation profile already has a customer associated with it, you will be unable to select a customer for the OS sequence. All servers provisioned with this OS installation profile will be associated with the specified customer (if a customer has been assigned).
- 8** From the Views pane, select **Attach Software Policy**.
- 9** At the bottom of the Content pane, click **Add** and select a software policy to add to the OS sequence.
- 10** From the Views pane, select **Attach Patch Policies**.
- 11** At the bottom of the Content pane, click **Add** and select a patch policy to add to the OS sequence.
- 12** From the Views pane, select **Attach Device Group**.
- 13** At the bottom of the Content pane, click **Add**. Select a device group to place the server into, after the OS sequence has been run. You can only select a public static group for this option.
- 14** From the Views pane, select **Remediate Polices**.

- 15** In the Content pane, choose to enable or disable remediation when the server is provisioned with the OS sequence. If you select Disable Remediation, then when you run the OS sequence, the OS will be installed but no policies in the OS sequence will be remediated – this means that no software in any of the policies attached to the OS sequence will be installed when the sequence is run.
- 16** If you select Enable Remediation, then you will need to configure the Rebooting and Scripts parameters. For the rebooting options, you can select one of the following:
- **Reboot servers as dictated by properties on each installed item:** Selecting this option will allow any reboot settings to run that might be set in any software or patch policies attached to the OS sequence.
  - **Hold all server reboots until after all items are installed:** This option will override any pre-install reboot options that might be set in any software or patch policies attached to the OS sequence. If any post-install reboots have been set, then they will execute after the OS has been installed.
  - **Suppress all server reboots:** This option will override reboot options set in any software or patch policies attached to the OS sequence.
- 17** Next, in the Scripts section, select either a Pre-Install/Post-Install Script. These tabs allow you to set a pre- or post-install script to be executed before the OS sequence has been run and after the OS has been installed. Click **Enable Script** to enable a the script parameters.
- 18** From the Select drop-down list, select either Saved Script or Ad Hoc Script. Each script type has its own settings:

#### **Saved Script**

- **Command:** Add any commands or arguments to be executed here.
- **Script Timeout:** Enter a numerical value for the number of minutes to pass until the script will timeout.
- **User:** Enter a user name and password, or choose to run the script as Local System. (If using a Unix OS, choose root as the user.)
- **Error:** Select if you want the OS sequence job to stop if the script returns an error.

### Ad Hoc Script

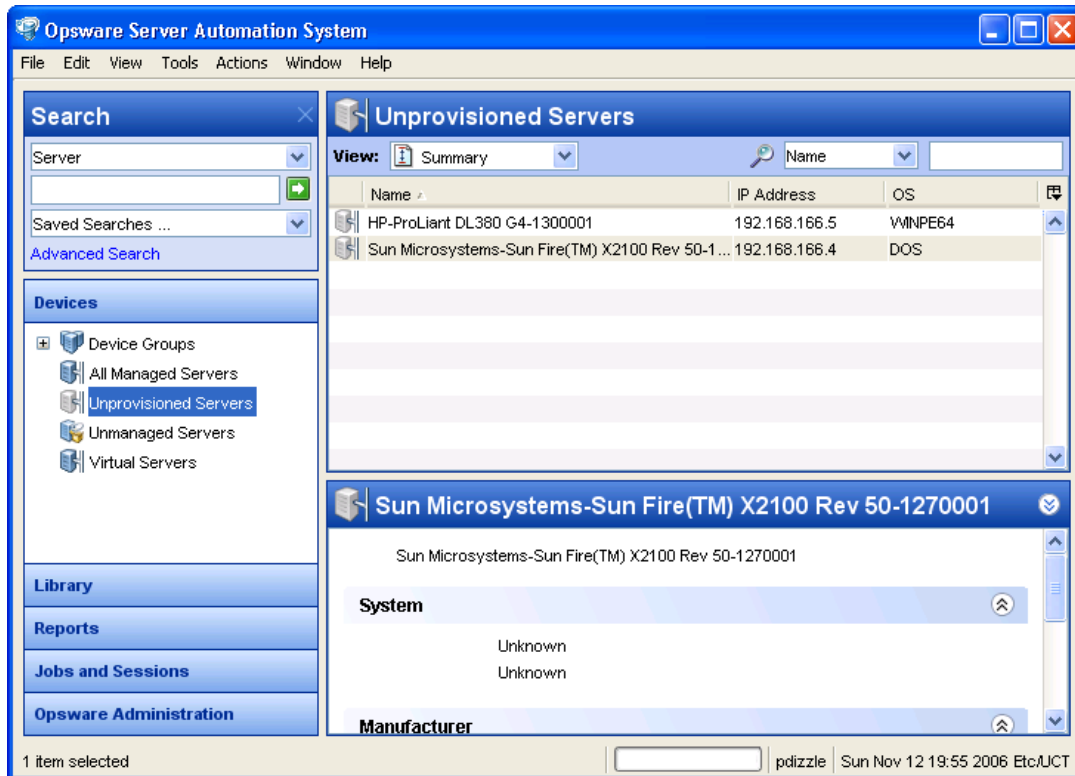
- **Type:** Choose UNIX shell for Unix systems, or for Windows, select BAT or VBSCRIPT.
- **Script:** Enter the text of the script. An Ad-Hoc script runs only for this operation and is not saved in Opware SAS. In the Script box, enter the contents of the script.
- **Command:** If the script requires command-line flags, enter the flags here.
- **Script Timeout:** Enter a numerical value for the number of minutes to pass until the script will timeout.
- **User:** Enter a user name and password, or choose to run the script as Local System account. (If using a Unix OS, choose root as the user.)
- **Error:** Select if you want the OS sequence job to stop if the script returns an error.

**19** When you have finished making your selections, from the **File** menu, select **Save** to save the OS Sequence.

## Select Servers in the Unprovisioned Servers List

To provision a server and install an OS, select an unprovisioned server from the Unprovisioned Servers list in the SAS Client. Servers in the Unprovisioned Servers list have registered their presence, but do not have an OS installed. From this location, you can install an OS by selecting an unprovisioned server. See Figure 10-2.

Figure 10-2: Unprovisioned Servers in the SAS Client




Select an unprovisioned server in the list and the Content pane will display detailed information about the unprovisioned server that was gathered by the OS Build Agent after a network boot.

The View drop-down list enables you to view the server in the following ways:

- **Summary:** Provides information about the host name set by booting the server the first time over the network or by using an Opware Boot Floppy or CD. It also displays the OS of the OS Build Agent (Windows, Red Hat Linux, or Solaris), processor type, manufacturer and model of the server, and Opware registration information.

- **Properties:** Displays placeholders for various management and reported information which will be filled in later once the server is provisioned.
- **Hardware:** Displays details about the hardware on the server, such as a processor type, physical and virtual memory, storage and network interfaces.
- **Custom Attributes:** Allows you to read and manage custom attributes.
- **History:** Indicates the first event associated with the server.

You can also search for an unprovisioned server using the search tool  in the upper right corner of the Content pane. You can choose a filter, then enter text to search for the server.




You also have the option of running an OS sequence from the Library and then selecting a server or servers as you configure the Run OS Sequence window.

---



Some servers in the Unprovisioned Servers list are in a server lifecycle state called *Planned*, which means the server has been partially prepared for OS provisioning. (It has a device record created for it, but no Opware mini agent installed yet.) An OS Sequence can't be run on Servers in the Planned state.

To display the server lifecycle stage value in the Unprovisioned Servers list, in the upper right corner of the Content pane, select the column selector  and from the list select Lifecycle. For more information, see your Opware administrator.

---

## Before Running an OS Sequence

### **Firewall Considerations**

The following operating systems come with default firewall settings that must be modified during the OS installation process in order to allow the Opware agent to be properly installed and configured on the target server.

- VMware ESX Server 3.0
- Windows 2003 x64 and Windows 2003 R2
- Windows XP SP2

OS Provisioning makes minor modifications to the firewall configurations on the managed server such that communication between the Opsware core and the Opsware agent can succeed.

### **VMware ESX 3.0 Firewall Settings**

VMware ESX 3.0 ships by default with an iptables firewall that will block communication between the core and the mini-agent or agent. In order for communications to and from the Opsware core to succeed, rules are added to the VMware ESX firewall by the build scripts and the Opsware agent.

### **Windows 2003 SP1 and Windows XP SP2 Firewall Settings**

For Windows 2003 SP1 and Windows XP SP2, in order for OS provisioning and ongoing management to succeed, Opsware must ensure that the Windows firewall settings are configured to bypass the default “Security Out Of the Box” experience and allow communication over the Opsware ports. Thus the Opsware OS provisioning process updates the Windows Firewall settings in the `unattend.txt` answer file as necessary for provisioning and management to work.

OS provisioning looks for the following Windows Firewall configurations in `unattend.txt`:

- There is no WindowsFirewall configuration in `unattend.txt`.
- There is a WindowsFirewall configuration, but it does not allow the ports needed by Opsware.
- There is a WindowsFirewall configuration that does allow the ports needed by Opsware (no changes will be made).

In any of the cases, after running an OS sequence and installing the OS (and agent), any predefined firewall settings remain in tact, with the exception that the Opsware agent will have been installed and all of its required ports will have been opened.

### **Red Hat EL 4 and Red Hat EL 5 Firewall Settings**

For Red Hat EL 4 and Red Hat EL 5, the following line in your `ks.cfg` profile will enable the firewall and allow the Opsware Agent to function correctly:

```
firewall --enabled --port 1002:tcp,1002:udp,1001:tcp,1023:tcp
```

### **Model Base Packages Functionality**

Opsware OS provisioning provides the ability to create software policies that model the base set of packages installed during OS provisioning.

During OS provisioning – after the base OS install, agent install, and reachability test, but before reconcile/remediate – a new script triggers software registration on the newly provisioned server, then models the installed packages as a software policy.

To activate this functionality, the server being provisioned must have a custom attribute defined (or inherited) named `model_base_packages`. The value for this attribute must either be empty or an absolute folder path to the name of the software policy to be created (or updated) with the package list.

If the `model_base_packages` value is empty, a software policy is created (or updated if it already exists) in the same folder as the OS Sequence. The software policy name will be the OS Sequence name plus `Base Packages`.

Each installed package that is successfully found in Opware SAS is added to the list of software policy items. A list of package names and versions that were not found in SAS will be available as a custom attribute named `missing_packages` in the software policy. This policy is attached to the OS Sequence which has remediation enabled. Because the above occurs before remediation, this policy is included in the remediation, thus adopting the modeled packages since they are by definition already installed.

You should only specify the `model_base_packages` custom attribute value as empty when running OS Sequences from the SAS Client. When running OS Provisioning from the SAS Web Client, the `model_base_packages` custom attribute value must be the path to the Software Policy.

The only valid value for the `model_base_packages` custom attribute is the path to a Software Policy. For example:

```
/Customer/OS Baselines/Solaris 10 baseline Q4 2007
```

In this case, the Software Policy will be created at the specified path and with the specified name. Any folders that are missing will automatically be created. If the Software Policy already exists, it will be updated.



When run from the SAS Web Client Install OS wizard, the Software Policy will be attached to the server being provisioned. However, since the Install OS wizard triggers a legacy reconcile, remediate is bypassed so the policy will not be remediated.

---



Note that it is not necessary to use the Model Base Packages feature for every OS Provisioning job. It needs only to be used once after an OS Profile changes. From that point on, the Software Policy will be attached to the OS Sequence unless you remove it, and will be available for other servers as they are provisioned.

### **Model Base Packages Script Usage**

The `model_base_packages.py` Command Engine script will function when called from another Command Engine script such as `provisionOS.py`. You can also run it as a standalone python2 `pytwist` script. The following are valid arguments when invoking the script:

```
model_base_packages.py --opsware-username you [--opsware-
password yourpass] --server <serverID> --ossequence
<ossequenceID> [--policy_path "/Some/Folder Path/Some Policy"]
:
```

Table 10-3: Options

ARGUMENT	DESCRIPTION
<code>--version</code>	Show the program version number and exit
<code>-h, --help</code>	Show this help message and exit
<code>-u OPSWAREUSERNAME, --opsware-username=OPSWAREUSERNAME</code>	Login username for SAS
<code>-p OPSWAREPASSWORD, --opsware-password=OPSWAREPASSWORD</code>	Login password for SAS
<code>-s SERVER, --server=SERVER</code>	Numeric Server ID of server to model
<code>m POLICYPATH, --policy_path=POLICYPATH</code>	Absolute path to the software policy that will model the packages
<code>-e OSSEQUENCE, --ossequence=OSSEQUENCE</code>	Numeric OS Sequence ID to link to the model software policy. If you specify an OS Sequence but not a policy path, the software policy will be created in the folder that contains the OS Sequence with the OS Sequence's name plus "Base Packages".

## Run an OS Sequence

To install an OS on an unprovisioned server, select a server from the Unprovisioned Servers list and run an OS sequence, or start an OS sequence and choose a target server in the Run OS Sequence window.




After you run an OS sequence job, if the OS Sequence does not have remediation enabled, the newly provisioned servers will not immediately perform a full software registration. Full software registration occurs after a small variable delay usually less than one hour. Thus when provisioning without remediation, the server's installed software packages and patches might not be listed immediately after the OS Sequence job completes. If this occurs, check again after one hour.

---

To run an OS Sequence and install an OS on an unprovisioned server, perform the following steps:



- 1 Choose a way to install an OS on an unprovisioned server:
  - From the Navigation pane, select **Devices** ► **Unprovisioned Servers**. Select a server and from the **Actions** menu, select **Run OS Sequence**.
  - Or
  - From the Navigation pane, select **Library** ► **OS Sequences**. Select the OS of the OS sequence, then select the OS sequence that you want to run and from the **Actions** menu, select **Run OS Sequence**.



If the **Run OS Sequence** menu item is grayed out, one or more of the unprovisioned servers is in a server lifecycle stage of Planned. Servers in this stage cannot be provisioned. You can display the server lifecycle stage value in the Unprovisioned Servers list. In the upper right corner of the Content pane, select the column selector . From the list, select **Lifecycle**. For more information, see your Opware administrator.

---

- 2 In the Select OS Sequence pane, click **Add** to add an OS sequence or click **Next** if OS Sequence is already listed.
- 3 In the Run OS Sequence window, step one requires that you add an unprovisioned server or servers to provision. To add a server, click **Add**.

- 4 Click **Next**, and in the Scheduling pane choose if you want to run the OS sequence, immediately, or at a later date and time.
- 5 Click **Next** and in the Notifications pane, select an email notifier. Click **Add Notifier** and enter an email address.
- 6 You can specify if you want the email to be sent upon success of the OS sequence job (  ) or failure of the OS sequence job (  ).
- 7 The ticket ID field is only used when Opsware Professional Services has integrated SAS with your change control systems. It should be left blank otherwise.
- 8 Click **Next**, and review the OS sequence information before you run the job.
- 9 Click **Start Job** to run the OS sequence. When the OS sequence job begins to run, click on the Job in the Job Status window or click **Close** to exit the Job Status window. You can also check the status of the Job by clicking on Job Logs under Jobs and Sessions in Navigation Pane.
- 10 When the OS sequence job has completed successfully, you can check the Devices
  - All Managed Servers list to see the newly provisioned server.



If you scheduled the OS sequence job to run at a later date and would like to cancel it, from the Navigation pane, select **Jobs and Sessions** ➤ **Recurring Schedules**. Then, select the job, right-click and select **Stop**.

### Reprovisioning a Managed Server

You have the option of reprovisioning a managed server, but keep in mind that reprovisioning a server completely removes all data on the server.

While all data is lost when you reprovision a server, you have the option of preserving the network configuration of the server. Also, some attributes are saved when you reprovision the server, which are defined in the build script for each OS. For more information on OS provisioning build scripts, see OS Provisioning Setup in the *Opsware® SAS Administration Guide*.



You can only reprovision a server that runs the Solaris or Linux operating system (but not Solaris x86).

To reprovision a managed server, perform the following steps:

- 1** From the Navigation pane, select **Devices ► All Managed Servers**.
- 2** Select a managed server to reprovision and from the **Actions** menu, select **Run OS Sequence**.
- 3** You will be shown a warning message that you are about to reprovision a managed server. By doing so, you will lose all data on the server. Click **Yes** to proceed.
- 4** In the Run OS Sequence window, please select the appropriate option before you begin the reprovisioning:
  - Yes, I understand the OS installation process will erase all data on the selected servers. (Mandatory. You must select this option in order to proceed.)
  - Please preserve the network configuration for the selected servers. (Optional)
- 5** Click **Next**. In the Run OS Sequence window, select an unprovisioned server or servers to provision. To add a server, click **Add**.
- 6** Click **Next**. In the Select OS Sequence pane, click **Add** to add an OS sequence.
- 7** Click **Next**, and in the Scheduling pane, choose if you want to run the OS sequence, immediately, or at a later date and time.
- 8** Click **Next** and in the Notifications pane, select an email notifier. Click **Add Notifier** and enter an email address.
- 9** (Optional) Specify if you want the email to be sent upon the success of the OS sequence job or failure of the OS sequence job.
- 10** You can also specify a Ticket Tracking ID in the Ticket ID field.
- 11** Click **Next**, and review the OS sequence information before you run the job.
- 12** Click **Start Job** to run the OS sequence. When the OS sequence has run, click **View Results** to view the results of the OS sequence job.
- 13** When the OS sequence job has been run, you can check the Devices ► All Managed Servers list to see the newly reprovisioned server.

# Appendix A: VAM Platform Support

## IN THIS APPENDIX

This section discusses the following topics:

- Supported Platforms in VAM

This appendix provides information about the operating system platforms and architecture that VAM supports scanning and displaying application (process families), server, and device information.



---

This list of operating system support for VAM is a subset of the supported platforms for the Opware Agent, since in order for VAM to be able to fully scan a server it must be under Opware management with an Opware agent. For more information on supported platforms for the Opware Agent, see the chapter on server asset tracking in the *Opware<sup>®</sup> SAS User's Guide: Server Automation*

---

## Supported Platforms in VAM

For non-Linux and non-VMware operating systems, VAM supports each operating systems kernel out of the box, and assumes no customizations have been made. For a list of non-Linux and non-VMware operating systems and kernels listed supported by VAM, see Table A-1.

For Linux and VMware ESX 3 operating systems, there are certain out of the box kernel versions that VAM supports. For information on Linux and VMware operating systems and kernels supported by VAM, see Table A-2.

Table A-1: VAM Supported Operating Systems - Non-Linux/VMware

VAM SUPPORTED OPERATING SYSTEMS	OS VERSIONS	ARCHITECTURE
<b>AIX</b>		
	AIX 4.3 AIX 5.1 AIX 5.2 AIX 5.3	POWER, Itanium
<b>HP-UX</b>		
	HP-UX 10.20 HP-UX 11.00 HP-UX 11.11	PA-RISC
	HP-UX 11i v2	PA-RISC and Itanium
<b>Sun Solaris</b>		
	Sun Solaris 6 Sun Solaris 7 Sun Solaris 8 Sun Solaris 9	Sun SPARC
	Solaris 10, through Update 4	Sun SPARC, 32 bit x86, 64 bit x86 and Niagara
<b>Fujitsu Solaris</b>		

Table A-1: VAM Supported Operating Systems - Non-Linux/VMware (continued)

VAM SUPPORTED OPERATING SYSTEMS	OS VERSIONS	ARCHITECTURE
	Fujitsu Solaris 8 Fujitsu Solaris 9 Fujitsu Solaris 10	Fujitsu SPARC
<b>Windows</b>		
	Windows NT 4.0 Windows 2000 Server Family Windows Server 2003	32 bit x86
	Windows Server 2003 x64	64 bit x86 (not Itanium)
	Windows XP Professional	32 bit x86
	Windows XP Professional x64	64 bit x86 (not Itanium)

Table A-2: VAM Supported Operating Systems - Linux and VMware

VAM SUPPORTED OPERATING SYSTEMS	VERSIONS	KERNEL	ARCHITECTURE
<b>Red Hat Linux</b>			
	Red Hat Linux 6.2	2.2.16	32 bit x86
	Red Hat Linux 7.1	2.4.20	32 bit x86
	Red Hat Linux 7.2	2.4.7-10	32 bit x86
	Red Hat Linux 7.3	2.4.18-3	32 bit x86
	Red Hat Linux 8.0	2.4.18-14 2.4.18-17.8.0	32 bit x86
	Red Hat Enterprise Linux 2.1 AS	2.4.9	32 bit x86
	Red Hat Enterprise Linux 3 AS Red Hat Enterprise Linux 3 ES Red Hat Enterprise Linux 3 WS	2.4.21-x.EL	32 bit x86 64 bit x86 and Itanium
	Red Hat Enterprise Linux 4 AS Red Hat Enterprise Linux 4 ES Red Hat Enterprise Linux 4 WS	2.6.9-x.EL	32 bit x86 64 bit x86  (VAM does not support Itanium for these Red Hat OS and kernel versions)
	Red Hat Enterprise Linux Desktop 5 Red Hat Enterprise Linux Server 5	2.6.18-8.el5xen	32 bit x86 64 bit x86
<b>SUSE Linux</b>			
	SUSE Linux Standard Server 8	2.4.18	32 bit x86
	SUSE Linux Enterprise Server 9	2.4.21, 2.6.5	32 bit x86 64 bit x86
	SUSE Linux Enterprise Server 10	2.6.16.13, 2.6.16.21	32 bit x86 64 bit x86
<b>VMware</b>			



Table A-2: VAM Supported Operating Systems - Linux and VMware

VAM SUPPORTED OPERATING SYSTEMS	VERSIONS	KERNEL	ARCHITECTURE
	VMware ESX Server 3 and 3.01	2.4.21- 37.0.2.ELvmlinix	32 bit x86 64 bit x86



## Appendix B: Glossary

### IN THIS APPENDIX

This appendix describes terminology and acronyms used in Opware 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 Opware SAS.

**administrator** See Opware administrator.

**Agent** See Opware Agent.

**Agent Installer** An application that installs the Opware Agent on a server.

**Agent Uninstaller** An application that uninstalls the Opware Agent on a server.

**application** A set of tiers and sub-tiers, such as a Web tier running Apache on Linux, an application tier running WebLogic on Windows, and a database tier running Oracle on Solaris.

**application component** An object that represents a process of running software, such as Apache, Oracle, BEA WebLogic, Microsoft® SQL Server, and so on.

**application configuration** Contains application configuration templates associated with an application.

**Application Configuration Management (ACM)** An Opware feature that enables you to manage and modify configuration files for applications on managed servers.

**application definition** Specifies an application's logical construction in terms of tiers, subtiers and the application components contained in those tiers. Allows you to transform a data display that contains extraneous and hard-to-understand information into a focused and easy-to-understand view of the data that is relevant to the application of interest.

**application tier** A set of process families that you create to organize the Application View so that you can see a compelling diagram of multiple servers, the connections among them, clients connecting to them, and dependencies to which they connect.

**ASAS** See Opware Application Automation Storage System (ASAS).

**attribute** A single property of a configuration item, the value of which describes the behavior of the configuration item. Configuration items can have one or more attributes. For example, some of the attributes for the configuration item SAS Server includes hostname, life cycle, agent status, and management IP.

**audit** A set of rules that express the desired state of a managed server's configuration objects – for example, a server's file system directory structure or files, a server's Windows Registry, application configuration, and so on.

**audit job** The process that performs an audit.

**audit policy** A collection of rules that define a desired state of configuration for a server.

**audit result** The results of running an audit, which will show how a target server or groups of servers's configuration object values match or mismatch the values.

**Automated Configuration Tracking** An Opware feature that allows users to monitor critical configuration files and configuration databases. When Opware 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 Opware-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)** The 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.

**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.

**Bundle** Core Components are bundled as a unit for installation to improve performance and scalability. For example, the Infrastructure bundle, the OS Provisioning bundle and the Slice Component bundle.

**CDR** See Code Deployment & Rollback (CDR).

**change log** An audit trail of changes made to a node (read-only). The log identifies who has recently modified the node and who has added or removed software packages, operating systems, and servers. It also tracks who has created or removed subordinate nodes.

**Code Deployment & Rollback (CDR)** An Opware 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 Opware Code Deployment & Rollback feature.

**combined event history log** A record of events performed on servers and network devices in your environment. These events are recorded in detail as actions performed on a certain date, by a certain user, on a certain server, or on a certain network device.

**Command Engine** The Opware SAS component that enables distributed programs to run across many servers. The Command Engine handles the entry of scripts into the Opware Model Repository (the script storage location in Opware 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 Opware 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. To determine if an Opware Agent is reachable, the Communication Test runs the 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.

**compliance** The degree to which a server object conforms to a test.

**configuration item (CI)** An object which can be viewed and managed with Opware OMDB. For example SAS Server, NAS Device, SAS Patch Policy.

**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** The process by which new configuration tracking policies or changes to existing configuration tracking policies 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. This is used when performing a variety of Opsware functions, including network and server configuration, notifications, and CRON script configurations.

**custom extension** Developed and deployed by Opsware Inc. support representatives, a custom extension is a Command Engine script that extends Opsware SAS functionality. A custom extension is often tailored to meet the specific needs of a customer.

**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. The different files from the Update directory are automatically synchronized with the current Live directory.

**Data Access Engine** The XML-RPC interface to the Model Repository that simplifies interaction with various clients such as the SAS Web Client, system data collection, and monitoring agents on servers.

**data center** Legacy term. See facility.

**data link connection** A layer 2 connection that includes switches that are directly connected to a managed server and switches that are indirectly connected through other switches.

**DCML Exchange Utility (DET)** A command-line tool (cbt) for importing and exporting Opsware SAS content. The primary function of this tool is to inject a newly-installed Opsware core with content from an existing core.

**deactivated server** A 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** A managed server or network device.

**device group** In SAS, a device group can contain managed servers and network devices, or only managed servers. In NAS, a device group contains only network devices. A device group helps you categorize your devices (servers and network devices) in ways that make sense for your organization.

**Dormant Opsware Agent** An Opsware Agent that runs in the dormant mode after its installation when the 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.

**duplex mismatch** A configuration mismatch between the speed and duplex of a managed server and a connected network device.

**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** An the Automated Configuration Tracking feature that sends an email to those listed in an 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 policy, is used 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.

**folder** A hierarchal file-system type organization structure that allows you to organize and manage your software resources and allows you to define security boundaries to control access to its content across user groups.

**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.

**gateway** See Opsware Gateway.

**Global Shell** A Unix shell that accesses the Opware Global File System (OGFS) and managed servers in a secure environment.

**group** See device group.

**Health Check Monitor (HCM)** The command-line scripts that check the status of Opware SAS core components.

**hypervisor** The host operating system that allows multiple operating systems (virtual servers) to run on a computer at the same time. Example hypervisors are VMware ESX and Solaris 10.

**IDK** Intelligent Software Module (ISM) Development Kit. The command-line tools (primarily ismtool) that build and upload ISMs into Opware SAS.

**Import Media tool** A utility script included with Opware SAS that is used to import OS media from the Media Server to Opware 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.

**Infrastructure Component bundle** A group of Core Components that are bundled as a unit for installation and which include the primary Management Gateway, the primary Data Access Engine, the Model Repository Multimaster Component/Tibco, the Opware Command Engine, and the Software Repository.

**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 for applications that are to be installed on managed servers.

**ISM control** A script within an ISM package that can be run on a managed server.

**job** Any major process run by the SAS Web Client or the SAS Client such as a communication test or a software installation.



**layer 2 connection** See data link connection.

**Library** In the SAS Client, the Library provides a way to display features such as application configurations, software policies, patches, patch policies, packages, OS sequences, and OS profiles managed by Opware.

**Live directory** In CDR, the directory that stores the actual code and content required to run a live site.

**MAC** See Media Access Control Address (MAC).

**Machine ID (MID)** A unique identifier that Opware SAS uses to identify the server. Opware 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 Opware Agent installed on it and is under the control of a particular Opware core.

**management IP** The IP address that Opware SAS uses to communicate with the Opware Agent on the server.

**manifest** Within CDR, a list of files that preview the results of an update to be performed. Each entry in the list specifies the file size, the last-modified date, the time stamp, 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 Opware 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 Opware database that stores information about managed server configurations within Opware SAS. It contains all information necessary to build, operate, and maintain an Opware-managed site. It includes data on all managed servers, the hardware associated with those servers, their memory, CPUs, storage capacity, configuration, 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** Opware SAS enables users to first model and simulate proposed operational changes to their environment before propagating these changes to production servers and applications. Using 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 Opware SAS control.

**MRL** See Media Resource Locator (MRL).

**multimaster core** An Opware core that belongs to a multimaster mesh.

**multimaster infrastructure component** See Model Repository Multimaster Component.

**multimaster mesh** A set of two or more Opware 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 Opware cores in a multimaster mesh can be managed through a single SAS Web Client.

**My Jobs** A page in the SAS Web Client that displays a list of jobs from the Model Repository such as software installation or server provisioning.

**name-value pairs** Legacy term. See custom attributes.

**NAS** See Network Automation System (NAS).

**NAS Integration** An Opware feature that enables you to closely examine detailed information about managed servers and the network devices connected to them so that you can determine how they are related and then, subsequently, coordinate and implement those changes.

**OCLI** See Opware command Line Interface (OCLI).

**OGFS** See Opware Global File System.

**OGFS Scripts** Scripts which can be executed in the Opware Global Shell from the SAS Client.

**OMDB See** Opware Operational Management Database.

**Opware administrator** The person responsible for overall administration, policy, and practices for accessing Opware SAS. The administrator can add users, define access to specific Opware SAS features, allow users to view site information, and deploy new code and content to their site.

**Opware Agent** Intelligent software on Opware-managed servers that is used to make changes to the servers. Depending on the request, the agent might use Global Opware

services. Some functions supported include software installation and removal, software and hardware configuration, server status reporting, and auditing.

**Opware API** An API that accesses the functionality and data model of Opware SAS. The Opware API supports Web Services, Java RMI, and other types of clients. The Opware API enables you to extend, integrate, and customize Opware SAS.

**Opware Application Storage Automation System (ASAS)** The Opware client that enables you to discover storage information in an Opware, Inc. data center. The information that ASAS discovers helps you optimize your storage environment by identifying storage resources that are under utilized or over provisioned.

**Opware Automation features** Opware SAS is made up of a set of Opware Automation features. Opware Automation features are the components that automate particular IT processes. The Opware Automation features include the following functions: Configuration Tracking, Code Deployment and Rollback, Script Execution, and Data Center Intelligence Reporting.

**Opware Automation Platform** The Opware API and runtime environment that facilitate the integration and extension of Opware SAS. The Opware APIs expose core services such as audit compliance, Windows patch management, and OS provisioning. The runtime environment (hub) executes scripts that can access the Opware Global File System (OGFS).

**Opware Discovery and Agent Deployment** A feature that helps deploy Opware Agents to a large number of servers through the SAS Client.

**Opware Operational Management Database (OMDB)** A configuration management database designed to create and maintain a record of the infrastructure data such as applications, servers, networks, and storage in your IT environment from various sources such as Opware SAS, Opware NAS, and other third party systems.

**OMDB Dashboard** The OMDB Dashboard displays either a default or user-configured collection of OMDB reports in an attempt to provide useful summary information on one window.

**Opware SAS Web Client** A web-based user interface for managing the Opware environment.

**Opware Command Line Interface (OCLI)** OCLI 1.0 has two commands, upload and download, which import and export files between the local file system and the Software Repository. OCLI 2.0 is intended for use only by Opware Inc. support representatives.

**Opware core** The server side of Opware SAS server-agent architecture. A core consists of the Opware 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. Opsware Agents communicate with the core through Opsware 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) structure that is derived from the Model Repository.

**Opsware Network Automation System (NAS)** The network device management application that tracks, validates, and automates network infrastructure changes. It enables compliance and enforces operational best practices across globally distributed, multi-vendor networks.

**Opsware Process Automation System (PAS)** The management application that automates incident resolution, maintenance tasks, and process integration in the data center. Opsware PAS is a run book automation platform.

**Opsware Satellite** Installed in a remote facility, an Opsware satellite provides a 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 Server Automation System (SAS)** The server management application that preserves the knowledge of system administrators, network engineers, and database administrators in a centralized knowledgebase. It automates previously manual tasks associated with the deployment, support, and the 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** The process of installing a basic set of software components, including an operating system and an Opsware Agent. After provisioning is complete, the server is ready to be managed by Opsware SAS.

**OS Provisioning bundle** Two Core components that are grouped together as a unit for installation and which include the Boot Server and the Media Server.

**Package Repository** Legacy term. See Software Repository.

**package** A collection of executables, configuration, or script files that are associated with an Opware-installable application or program. In Opware SAS, a package contains software package files registered in the Software Repository and 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.

**PAS** See Opware Process Automation System (PAS).

**patch** A piece of object code (binaries) that is inserted into (patched into) an executable program to temporarily fix a known defect.

**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 SAS Web Client.

**Patch Management** An Opware 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 Opware SAS features and resources. A resource is usually a set of managed servers. The set of managed servers corresponds to a facility, customer, or server group. A feature is an action, such as provisioning an OS.

**physical connection** The connection inferred from data link connections and represents direct connections (cables) between server and switches.

**physical server** TAn operating system running on a hardware device. A non-virtual server.

**platform** The name and version of an operating system.

**pluggable checks** Audit rules developed as scripts in a command-line environment and then imported into Opware SAS.

**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 remediate** Before Opsware SAS installs software on a server, it performs a preview remediate, and determines what will happen when the actual remediate 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.

**Pytwist** A set of Python libraries that provide access to the Opsware API from managed servers and custom extensions.

**realm** In Opsware SAS, a realm is a routable IP address space, which is serviced by one or more Opsware Gateways. The managed servers that connect to an Opsware core via a Gateway are identified as being in that Gateway's realm.

**remediate** The process of updating the actual software configuration of a server based on the specified configuration stored in the Model Repository.

**remediate output** After a remediate operation completes, Opsware SAS displays the remediate output for each server that was remediated. The remediate output aggregates output from the various installation, uninstallation, or post-installation scripts, messages from Opsware SAS, and messages from the system utilities that remediate uses to perform the installation and uninstallation of packages, operating systems, and patches.

**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.

**Reports** An Opsware feature that enables you to create reports that provide comprehensive, real-time information about managed servers, network devices, software, patches, customers, facilities, operating systems, compliance policies, and users and security in your environment.

**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** The queue in which the configuration files are placed before they are restored to a server.

**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 Opware shell is a command that makes a client connection enabling you to remotely run programs on managed servers.

**SAS** See Opware Server Automation System (SAS).

**Satellite** See Opware Satellite

**Saved Scripts** Public scripts that every Opware SAS user can access.

**selection criteria** Rules that instruct Opware SAS to collect information about server objects and how to collect that information. It optionally can also collect data for file comparison and inclusions/exclusions. Selection criteria is required for the snapshot and audit processes.

**sequence** The 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 used to create, modify, or delete a sequence definition.

**Sequence Performer (Production)** In CDR, a predefined User Group used to directly perform or request performance of a sequence action on production hosts.

**Sequence Performer (Staging)** In CDR, a predefined User Group used to directly perform or request performance of a sequence action on staging hosts.

**Sequence Requester (Production)** In CDR, a predefined User Group used to request performance of a sequence action on production hosts.

**Sequence Requester (Staging)** In CDR, a predefined User Group used 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** Opware 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 Opware Agents on the servers and registers them with the Model Repository.

**server baselines** The process of defining and provisioning servers with standard configurations. Opware templates can be used to automate the building of complete server baselines.

**Server Explorer** A feature of the SAS Client that allows you to browse and manage servers and server groups in your facility.

**Server ID** The primary key in the Opware Model Repository that represents a given server. The Server ID is used internally in Opware SAS.

**server lifecycle** The various server states assigned to a server by Opware SAS. Server states include Unprovisioned, Available, Installing OS, and Managed.

**server management** The process by which users can manage and track servers in an Opware-managed environment. Opware 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 Opware SAS, but do not have a full operating system installed.

**server provisioning** The process of installing a basic set of software components including the operating system, an Opware Agent, and other system utilities and debugging tools to manage the server. Configuration is defined in the Model Repository.

**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 used to define and modify or delete service definitions.

**Service Performer (Production)** In CDR, a predefined User Group used to directly perform or request performance of service operations on production hosts (servers).

**Service Performer (Staging)** In CDR, a predefined User Group used to directly perform or request performance of service operations on staging hosts.

**Service Requester (Production)** In CDR, a predefined User Group used to request performance of service operations on production hosts.

**Service Requester (Staging)** In CDR, a predefined User Group used 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 ontology.

**signature** A set of rules that you provide and that VAM uses to identify a process family.

**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.

**Slice Component bundle** A group of Core Components that are grouped together for installation and which includes the Core and Agent Gateways, the Opsware Command Center (OCC), the Opsware Global File System (OGFS), the Web Services Data Access Engine, the secondary Data Access Engine, and the Build Manager.

**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. Audit and Remediation records one snapshot per server.

**snapshot job** The process that created a snapshot of a server or server group.

**snapshot specification** 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 compliance** Software compliance indicates whether or not the all software policies attached to the selected server are compliant with the actual server configuration.

**software policy** A software policy contains packages, patches, application configurations, and other software policies and allows you to install software and configure applications on managed servers simultaneously.

**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.

**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.

**Storage Agent** An Opsware ASAS component that runs on a managed server to collect storage device information, such as SAN arrays, SAN switches, SAN fabrics, NAS filers, and Oracle databases. A Storage Agent must be installed in the Opsware core and then deployed on a managed server.

**Storage Host Agent Extension (SHA)** A Storage Host Agent Extension (SHA) runs on a managed server to discover storage related data on that managed server, such as Fibre Channel Adapter (FCA) hardware and drivers, file system mounts, and so on.

**synchronization** The 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 used to create, modify, or delete a synchronization definition.

**Synchronization Performer** In CDR, a predefined User Group used to directly perform or request performance of a synchronization action.

**Synchronization Requester** In (CDR), a predefined User Group used 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*.

**topology** A snapshot that represents the state of a set of network devices and managed servers, the process families running on those servers, the connections among those process families, and any external clients and dependencies.

**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).

**virtual server** The logical server or virtual machine that runs on a hypervisor. To an application, the virtual server appears to be a physical server. Example virtual servers are VMware virtual machines and Solaris zones.

**Visual Application Manager** An Opsware feature that is designed to help you understand and manage the operational architecture and behavior of distributed business applications in your IT environment.

**Web Services API** A SOAP-based interface that facilitates the integration of operations and business support systems with Opsware SAS. The Opsware Web Services API 2.2 is intended for Opsware SAS 5.x. A different Web Services API, introduced in Opsware SAS 6.1, accesses the Opsware API.

**Web Services Data Access Engine** A web services interface to the Model Repository that provides increased performance to other Opsware SAS components.

**zone** A virtual server that is a Solaris non-global zone.

# Index

## A

- ad hoc audit, running .....214
- adding
  - application configuration, server or group to ..526
  - configuration template .....518
  - rule exceptions to an audit .....207
- Ad-Hoc Scripts
  - definition of .....579
- AIX
  - APARs
    - about .....411, 412
    - uploading .....411, 412
    - LPPs, about .....411
- APARs. See AIX APARs.
- application and storage signatures (VAM)
  - about .....106
- application configuration
  - Compliance Dashboard .....302
- Application Configuration Management
  - adding, configuration template .....518
  - adding, server or group to .....526
  - application configuration, definition of .....579
  - application configuration, overview .....499
  - applying and inheriting values .....503
  - CML .....503
  - comparing, configuration templates .....535
  - configuration out of synch with server .....537
  - configuration template, overview .....499
  - configuring audit rules .....168
  - creating
    - application configuration .....513
    - configuration template .....515
  - default values, inheritance .....504
  - definition of .....579
  - deleting, configuration template .....518
  - editing, default values .....523
  - instance values, inheritance .....505
  - loading
    - template .....519
    - values .....530
  - main components .....498
  - preserve values .....502
  - process .....496
  - pushing changes, servers or groups to .....532
  - restore to previous state .....540
  - scheduling
    - push .....533
    - scan for compliance .....537
  - sequence merging
    - append mode .....509
    - overview .....508
    - prepend mode .....511
    - primary key option .....510
    - replace mode .....509
  - setting values, on .....528
  - setting, templates to run as scripts .....521
  - show inherited values .....502
  - specifying, template order .....522
  - value set editor
    - definition of .....596
    - overview .....499
- application or storage signature (VAM)
  - creating .....109
- application signature (VAM)
  - defined .....44
  - evaluation order .....107
- attaching
  - software policy to server .....449
- audit
  - audit job, definition of .....580
  - audit results
    - viewing and remediating .....229
  - auditing process .....145
  - audits and the Compliance Dashboard .....140
  - configuring, overview .....152
  - elements of .....146
  - performing
    - audit results, from .....216
    - re-running from audit results .....216
    - results, comparison based remediation .....221
    - results, value based remediation .....223
    - running from the library .....213
    - running on a server .....214
    - saving as audit policy .....212

schedule configuration compliance scan . . . . .	538	definition of . . . . .	592
scheduling . . . . .	217	inclusions/exclusions . . . . .	198
scheduling recurring . . . . .	217	inclusions/exclusions, definition of . . . . .	584
searching for . . . . .	235, 253	snapshot	
selection criteria		definition of . . . . .	594
inclusions/exclusions . . . . .	198	job, definition of . . . . .	594
snapshot used in . . . . .	236	template, definition of . . . . .	594
sources, server or snapshot . . . . .	154	source, definition of . . . . .	595
viewing completed audit job . . . . .	219	target, definition of . . . . .	595
ways to create . . . . .	148	terms and concepts . . . . .	140
creating from a server . . . . .	148	viewing	
from a group of servers . . . . .	149	and remediating audit results . . . . .	229
from a snapshot . . . . .	149	ways to create an audit . . . . .	148
from an audit policy . . . . .	150	audit policy	
from the library . . . . .	149	creating . . . . .	209
Audit and Remediation		linking and importing . . . . .	211
audit policies . . . . .	208	overview . . . . .	208
audit process overview . . . . .	145	saving . . . . .	212
audit results . . . . .	220	audit template	
capturing golden server configuration . . . . .	138	definition of . . . . .	580
creating an audit policy . . . . .	209	available	
deleting		patches, definition of . . . . .	580
snapshot . . . . .	592	servers, definition of . . . . .	580
snapshot specification . . . . .	254	<b>B</b>	
enforcing security policies . . . . .	139	backups	
examples (use cases) . . . . .	138	backup event, definition of . . . . .	580
exceptions . . . . .	206	full, definition of . . . . .	583
adding to an audit . . . . .	207	incremental backup, definition of . . . . .	584
editing . . . . .	208	booting	
rules that cannot have exceptions . . . . .	206	Solaris servers, over network . . . . .	556
linking and importing audit policies . . . . .	211	Build Manager	
overview . . . . .	138	definition of . . . . .	581
performing audit		OS Build Agents, locating . . . . .	557
audit results, from . . . . .	216	<b>C</b>	
rules . . . . .	157	change logs, definition of . . . . .	581
configuring, application configuration . . . . .	168	code deployment	
configuring, COM+ . . . . .	174	definition of . . . . .	580, 581
configuring, custom script . . . . .	175	deployment, definition of . . . . .	582
configuring, file system . . . . .	178	live directory, definition of . . . . .	585
configuring, hardware . . . . .	180	restore, definition of . . . . .	592
configuring, IIS Metabase . . . . .	181	sequences	
configuring, operating system . . . . .	185	definition of . . . . .	592
configuring, software . . . . .	188	sequence editor, definition of . . . . .	592
configuring, TON . . . . .	195	sequence performer, definition of . . . . .	592
configuring, users and groups . . . . .	189	sequence requester, definition of . . . . .	592
configuring, Windows Registry . . . . .	191	site backup directory, definition of . . . . .	594
configuring, Windows services . . . . .	193	site previous directory, definition of . . . . .	594
example of configuring . . . . .	164		
server objects . . . . .	160		
scheduling audits . . . . .	217		
selection criteria			

- user roles, definition of .....581
  - COM+ object
    - configuring Audit and Remediation rule .....174
  - Command Engine, definition of .....581
  - command line interface, definition of .....587
  - Communication Test
    - definition of .....581
  - comparing scan results (VAM)
    - heuristics used .....134
  - comparing snapshots (VAM)
    - "significant" object attribute differences .....124
    - how to .....125
    - object attribute differences .....122
    - object existence comparison .....121
    - source and comparison .....120
  - comparing, configuration templates .....535
  - compliance
    - performing software compliance scan .....468
    - software .....467
    - viewing server compliance in VAM .....87
  - Compliance Dashboard
    - application configuration .....302
    - audit .....289
    - compliance statuses .....262
    - general categories .....285
    - overview .....257
    - patch .....298
    - refreshing .....283
    - remediation overview .....285
    - software compliance .....293
    - terms and concepts .....268
    - usage, proactive or reactive .....259
    - viewing .....268
  - compliance status for Compliance Dashboard .....262
  - components
    - Application Configuration Management, of .....498
  - Configuration Markup Language .....503
  - configuration template
    - adding .....518
    - adding, scripts .....521
    - comparing .....535
    - creating .....515
    - definition of .....582
    - deleting .....518
    - loading
      - CML file .....519
      - values into .....530
    - overview .....499
    - specifying, order .....522
  - configuration tracking
    - backups, definition of .....580
    - definition of .....580
    - email notification lists
      - definition of .....583
    - policies
      - definition of .....582
    - reconcile, definition of .....582
    - restore queue, definition of .....592
    - restore, definition of .....591
    - server-based policies, definition of .....593
  - configuring
    - snapshot specification .....242
  - conventions used in the guide .....27
  - copying
    - objects to a server from a snapshot .....256
  - creating
    - application configuration .....513
    - application or storage signature (VAM) .....109
    - application tier (VAM) .....105
    - audit policy .....209
    - business application contacts in VAM .....103
    - configuration template .....515
    - OS sequence .....560
    - scripts .....474
    - snapshot in VAM .....119
    - snapshot specification .....241
    - snapshot specification from library .....241
  - custom attributes
    - definition of .....582
  - custom extensions
    - definition of .....582
  - custom script
    - configuring Audit and Remediation rule .....175
  - customers
    - definition of .....582
  - cutting over
    - definition of .....582
- ## D
- Data Access Engine, definition of .....582
  - data center. See facility.
  - deactivating
    - servers, definition of .....582
  - deleting
    - configuration template .....518
    - scripts .....482
    - snapshot .....234, 254, 592
    - snapshot job schedule .....249
    - snapshot specification .....254
  - depots
    - patch management .....410

deprecated packages, definition of .....583

detaching  
     software policy to server .....460

device tree (VAM) ..... 46

DHCP  
     Linux servers, requirements for using ... .551, 554  
     servers, booting with ..... 546  
     Solaris servers, usage of ..... .549, 550

directories  
     live, definition of .....585  
     site backup directory, definition of .....594  
     Update directory, definition of .....596

## E

editing  
     audit rule exceptions .....208  
     audit schedule .....218  
     default values, application configuration for ... 523  
     snapshot job schedule .....248

email  
     notification lists for configuration tracking, definition  
         of .....583

e-mail notifications ..... 360, 361, 388, 397, 426, 434

encoding  
     choose for application configuration source ... 502  
     uploading template file .....519

Environment Tree  
     definition of .....583

error messages (VAM) .....131

evaluation order, VAM application signatures ... 107

Exceptions, Audit and Remediation  
     about .....206  
     adding to an audit .....207  
     considerations .....206  
     rules that cannot have exceptions .....206

executing  
     OGFS script .....490  
     scripts .....482  
     server script .....484

executing scripts .....482

exporting  
     scripts .....481  
     VAM properties to .csv ..... 81

exporting a VAM map ..... 79

## F

facilities, definition of .....583

File System  
     configuration Audit and Remediation rule ..... 178

font ..... 380

## G

Global Shell  
     definition of ..... 584

global shell  
     opening from VAM ..... 101

groups, definition of .....584

## H

hardware preparation, overview ..... 548

hardware,configuring Audit & Remediation rule .. 180

hotfix chaining ..... 332

HP-UX  
     depots  
         patch management ..... 410

hypervisor  
     definition of ..... 584

## I

icons, toolbar icons (VAM) ..... 55

IIS Metabase  
     configuring Audit and Remediation rule ..... 181

import media tool, definition of ..... 584

importing audit policy .....212

Install Patch Wizard .....421, 428

Install Patch wizard .....381, 420

install scripts,specifying .....358, 432

installation  
     flags, overview ..... 381, 392, 420

installing  
     Install Patch Wizard .....421, 428  
     Opware Agents  
         definition of ..... 592  
     OS Build Agents  
         verification .....557  
         software using software policy .....449, 460

Intelligent Software Module (ISM)  
     definition of ..... 584

IP range groups, definition of ..... 584

IP ranges  
     definition of ..... 584

ISM Control  
     running ..... 462

ISM control, definition of ..... 584



**J**

Japanese	378
jobs	
definition of	584
snapshot, definition of	594

**K**

Korean	378
--------	-----

**L**

layer 2 connections displayed in VAM	69
life cycle	
definition of	593
linking an audit policy	211
loading	
template	519
values, configuration template into	530
local attachment, definition of	585
locales	377

**M**

Machine ID (MID), definition of	585
management IP	
definition of	585
manifests, definition of	585
mbsacl.exe	328, 332
Media Access Control Address, definition of	585
media resource locators (MRLs), definition of	585
media server, definition of	585
Microsoft Baseline Security Analyzer	328
Microsoft patch management prerequisites	328
Model Repository	330, 351, 352
Model Repository Multimaster Component, definition of	585
Model Repository, definition of	585
modeling and change simulation engine	
definition of	586
msiexec.exe	332
multimaster	
mesh, definition of	586
multimaster core, definition of	586
My Jobs	
definition of	586

**N**

network	
Solaris servers, booting over	556

**O**

opening	
Device Explorer from VAM	101
global shell from VAM	101
remote terminal from VAM	101
scripts	478
opening a Business Application	112
operating systems	
configuring Audit and Remediation rule	185
patch management, supported for	408
provisioning	541
Opware administrator, definition of	586
Opware Agent	330
definition of	586
dormant, definition of	583
registration	404
Opware Agent Installer	
definition of	579
Opware Agent Uninstaller, definition of	579
Opware Discovery and Agent Deployment	
definition of	587
Opware Gateway, definition of	588
Opware Global File System	
definition of	588
Opware guides	
contents of	25
conventions used	27
documentation set	28
icons in guide, explained	27
Opware model space, definition of	588
Opware SAS	
core, definition of	587
documentation set	28
features, definition of	587
installation, definition of	588
related documentation	28
Opware Satellite, definition of	588
OS Build Agents	
Build Manager, locating	557
definition of	588
failure to install, recovering from	558
overview	551
verifying installation	557
OS media, definition of	588
OS provisioning	
definition of	588
hardware preparation	548
import media tool, definition of	584
managed server values	545
media resource locators (MRLs), definition of	585
media server, definition of	585

OS Build Agents	
definition of	588
overview	551
using	551
SAS Client	
creating an OS sequence	560
overview	559
reprovisioning a managed server	571
running an OS sequence	570
select unprovisioned servers	565
Server Pool values	545
Solaris servers	547
Windows servers	547
OS sequence	
attach device group	561
attach patch policy	561
attach software policy	560
creating	560
set remediate policy	561
overview	
script execution	471
<b>P</b>	
package types	
AIX APAR	411, 412
HP-UX depots	410, 412
LPP	411
RPM	408
Windows Hotfix	331, 382
packages	
definition of	589
deprecated, definition of	583
packaging server	
definition of	589
patch	
Compliance Dashboard	298
patch compliance	351, 352, 354
patch compliance scan	329
patch installation, previewing	389, 427
patch installation, scheduling	359, 388, 426
patch management	
definition of	589
Microsoft patch releases	329
operating systems, supported	408
patch administrators	
definition of	589
patch information from Agent	404
patch testing, support for	404
roles	407
supported Unix versions	408
uploading automatically	344
patch policy	350
patch policy exception	352
patch reboot options	357, 385, 423
patch uninstallation, previewing	398, 434
patch uninstallation, scheduling	397, 433
patches	
installation flags	381, 420
types supported	408
uninstallation flags	392
performing	
audit	
audit results, from	216
physical server	
definition of	589
policy setter	334
populate-opsware-update-library	344
post-install script, definition of	589
post-uninstall script, definition of	589
pre-install script, definition of	589
pre-uninstall script, definition of	590
preview remediate	
definition of	591
primary IP, definition of	591
printing a VAM map	79
private group, definition of	591
process families (VAM)	
defined	45
properties of	82
properties	
application signature (VAM)	84
storage signature (VAM)	84
properties (VAM)	
links	91
network device	91
network interface	92
port group	92
process family	82
server compliance	87
tiers	83
virtual server	90
virtual switch	92
property page (VAM)	
port group	92
public group, definition of	591
push	
application configuration	
onto a server or group	532
scheduling	533

**Q**

qchain.exe .....332  
 QNumber .....332, 361

**R**

realm, definition of .....591  
 reference server, definition of .....591  
 refreshing  
   Compliance Dashboard .....283  
 remediate  
   definition of .....591  
   output  
     definition of .....591  
     overview .....453  
   preview  
     definition of .....591  
     software policy .....454  
 remote Opsware shell  
   definition of .....592  
 remote terminal  
   definition of .....591  
   opening from VAM .....101  
 renaming  
   scripts .....482  
 reports  
   software policy .....468  
 reprovisioning a managed server, SAS Client ....571  
 rolling back  
   definition of .....592  
 RPM  
   patching .....408  
 rules  
   Audit and Remediation .....157  
 running  
   ad hoc audit .....214  
   audit from server .....214  
   audit from the library .....213  
   ISM Controls .....462  
   OS sequence, SAS Client .....570  
   snapshot specification .....245

**S**

SAS Client  
   OS installation with .....559  
 SAS Web Client  
   My Jobs  
     definition of .....586  
     patch administration in .....418  
 saving

  snapshot specification as policy .....245  
 scan  
   configuration compliance .....537  
   patch compliance .....368  
   software compliance .....468  
 scan timeout (VAM) .....99  
 scheduling  
   audit .....217  
   audit, recurring .....217  
   snapshot in VAM .....119  
   snapshot job .....247  
 script execution .....482  
   creating .....474  
   deleting .....482  
   editing scripts  
     editing  
       **scripts** 479  
       executing OGFS script .....490  
       executing server script .....484  
       exporting .....481  
       opening a script .....478  
       overview .....471  
       process .....472  
       renaming .....482  
       types .....472  
       viewing script history .....480  
 scripts  
   Distributed Scripts  
     overview .....471  
   My Scripts, definition of .....586  
   post-install script, definition of .....589  
   post-uninstall script, definition of .....589  
   pre-install script, definition of .....589  
   pre-uninstall script, definition of .....590  
   running on devices in VAM .....113  
   setting, templates to run as script .....521  
   Shared Scripts, definition of .....594  
 search  
   audit .....235, 253  
 searching  
   server search, definition of .....593  
 sending email to Business Application contacts .104  
 sequence merging, application configuration  
   append mode .....509  
   prepend mode .....511  
   primary key option .....510  
   replace mode .....509  
 sequences  
   editor for CDR, definition of .....592  
 server  
   attaching software policy .....449

detaching software policy . . . . .	460	editor for CDR, definition of . . . . .	593
server baselines, definition of . . . . .	593	performer (production) for CDR, definition of . . . . .	593
Server Explorer		performer (staging) for CDR, definition of . . . . .	594
definition of . . . . .	593	requester (production) for CDR, definition of . . . . .	594
server groups		requester (staging) for CDR, definition of . . . . .	594
adding, application configuration . . . . .	526	service-instance for CDR, definition of . . . . .	594
definition of . . . . .	593	setting	
dynamic, definition of . . . . .	583	application configuration values . . . . .	528
private, definition of . . . . .	591	configuration templates to run as scripts . . . . .	521
public, definition of . . . . .	591	Shared Scripts	
pushing changes, application configuration of . . . . .	532	definition of . . . . .	594
setting values on . . . . .	528	showing IPC service names in VAM maps . . . . .	80
static, definition of . . . . .	595	signatures (VAM)	
Server ID, definition of . . . . .	593	cutting and copying . . . . .	111
server management		deleting . . . . .	111
definition of . . . . .	593	editing . . . . .	110
groups, definition of . . . . .	584	pasting . . . . .	111
IP range groups, definition of . . . . .	584	snapshot	
IP ranges, definition of . . . . .	584	copying objects to server from . . . . .	255
Server Map (VAM) . . . . .	65	definition of . . . . .	594
server objects		deleting . . . . .	234, 254, 592
Audit and Remediation . . . . .	160	template . . . . .	254
Server Pool		deleting job schedule . . . . .	249
definition of . . . . .	593	difference between snapshot specification . . . . .	236
servers		editing job schedule . . . . .	248
adding, application configuration . . . . .	526	job, definition of . . . . .	594
booting		locating . . . . .	250
over network . . . . .	556	locating in SAS Client . . . . .	249
configuration policies for, definition of . . . . .	593	process . . . . .	240
deactivating, definition of . . . . .	582	saving as audit policy . . . . .	212
definition of . . . . .	592	scheduling . . . . .	247
life cycle		template, definition of . . . . .	594
definition of . . . . .	593	used in an audit . . . . .	236
Machine ID, definition of . . . . .	585	used with audit policies . . . . .	237
managed, definition of . . . . .	585	viewing contents of . . . . .	251
management IP, definition of . . . . .	585	snapshot specification . . . . .	241
media access control address, definition of . . . . .	585	and audit policies . . . . .	237
primary IP, definition of . . . . .	591	configuring . . . . .	242
provisioning, definition of . . . . .	593	configuring rules for . . . . .	244
pushing changes, application configuration of . . . . .	532	creating from library . . . . .	241
reference server, definition of . . . . .	591	creating from server . . . . .	241
reprovisioning, SAS Client . . . . .	571	deleting . . . . .	254
Server ID, definition of . . . . .	593	elements of . . . . .	237
Service Levels, definition of . . . . .	594	relationship to snapshots . . . . .	236
setting values on . . . . .	528	running . . . . .	245
Solaris servers, booting . . . . .	556	selection criteria	
source, definition of . . . . .	595	inclusions/exclusions . . . . .	198
status, definition of . . . . .	593	software	
Service Levels, definition of . . . . .	594	configuring Audit and Remediation rule . . . . .	188
services		software compliance	
definition of . . . . .	593	Compliance Dashboard . . . . .	293

- compliance remediate options .....294
  - software installation
    - attaching software policy .....449
    - detaching software policy .....460
    - installing using software policy .....449, 460
    - overview .....437
    - process .....438
    - remediate
      - overview .....453
      - remediate software policy .....454
      - software policy template overview .....461
      - ways to install .....440
  - software management
    - installation, overview .....437
  - software policy
    - attaching to server .....449
    - compliance overview .....467
    - detaching to server .....460
    - installing software .....449, 460
    - performing software compliance scan .....468
    - remediate .....454
    - remediate overview .....453
    - reports .....468
    - running ISM controls .....462
    - software policy template overview .....461
  - software policy template
    - overview .....461
  - Software Repository .....330
  - Software Repository Cache, definition of .....595
  - Software Repository, definition of .....595
  - Solaris
    - booting servers over network .....556
    - OS provisioning .....547
  - specifying
    - application configuration template order .....522
    - application configuration to run as script .....521
  - standalone core, definition of .....595
  - static group, definition of .....595
  - Storage Host Agent Extension, definition of .....595
  - storage signature (VAM) .....45
  - Summary Review .....361, 390, 398, 427, 434, 435
  - synchronizations
    - definition of .....595
    - editor for CDR, definition of .....595
    - performer for CDR, definition of .....595
    - requester for CDR, definition of .....595
- T**
- target, definition of .....595
  - templates
    - local attachment, definition of .....585
    - Tiers properties (VAM) .....83
    - TON rules
      - configuring Audit and Remediation rule .....195
    - troubleshooting
      - OS Build Agents
        - installation failure .....558
        - verifying installation .....557
    - tunnel, definition of .....596
    - types of scripts .....472
- U**
- Uninstall Patch wizard .....392, 429
  - uninstallation
    - flags, overview .....381, 392, 420
  - unzip.exe .....332
  - user roles, definition of .....596
  - users and groups, configuring Audit and Remediation rules .....189
  - users, definition of .....596
- V**
- value set editor
    - definition of .....596
  - value set editor, application configuration
    - editing default values .....523
    - overview .....499
    - preserve values .....502
    - show inherited values .....502
  - verifying
    - installation of OS Build Agents .....557
  - viewing
    - audit results .....229
    - audit server usage .....150
    - completed audit job .....219
    - Compliance Dashboard .....268
    - script history .....480
    - snapshot contents .....251
  - virtual server
    - definition of .....596
  - Virtualization Director
    - scan time out preference .....99
    - virtualization settings .....99
  - Visual Application Manager (VAM)
    - .business application
      - overview .....36
    - accessing servers
      - opening a remote terminal .....101
      - opening Device Explorer .....101

- opening global shell .....101
- Business Application
  - opening .....112
  - saving .....113
  - saving as template .....113
- business application
  - application signatures evaluation order ...107
  - copying and cutting a tier .....105
  - creating a tier .....105
  - deleting a tier .....105
  - pasting a tier .....106
  - templates .....103
  - tiers .....104
- comparing snapshots
  - "significant" object attribute differences ...124
  - comparison types .....120
  - how to .....125
  - object attribute difference .....122
  - object existence comparison .....121
  - source and comparison .....120
- creating a snapshot .....119
- data collection and display ..... 41
- error messages .....131
- filtering data .....127
  - criteria to use .....129
  - using regular expressions .....130
- icons in toolbar .....55
- menus .....60
- opening a snapshot .....119
- overview .....32
- prerequisites to run .....35
- process and link connection symbols .....77
- properties
  - application signature .....84
  - connection links .....91
  - network devices .....91
  - network interfaces .....92
  - port group .....92
  - process families .....82
  - server compliance .....87
  - storage signature .....84
  - tiers .....83
  - virtual servers .....90
  - virtual switches .....92
- property pages .....80
  - port groups .....92
  - server .....81
- scheduling a snapshot .....119
- supported operating systems .....35
- symbols used in maps .....77
- user interface .....54

- VAM application .....42
  - application and storage signatures explained .  
106
  - application signature .....44
  - creating definition .....113
  - devices tree .....46
  - process families .....45
  - storage signatures .....45
  - tiers tree .....42
- VAM maps
  - exporting .....79
  - infrastructure pane .....75
  - network map .....68
  - printing .....79
  - SAN map .....74
  - server map .....65
  - showing IPC service names .....80
  - storage map .....72
- virtualization settings .....99
- Visual Packager
  - definition of .....589

## W

- Web services APIs, definition of .....596
- Web Services Data Access Engine, definition of .596
- Windows Hotfix
  - installation flags .....382
  - uploading .....382
- Windows Registry
  - configuring Audit and Remediation rule .....191
- Windows servers
  - OS provisioning .....547
- Windows services
  - configuring Audit and Remediation rule .....193
- Windows Update Agent .....332
- wizards
  - Distributed Script Execution .....471
  - Install Patch .....421, 428

## Z

- zone, definition of .....596