

HP Network Automation

Software Version: 9.0

Multimaster Distributed System on SQL Server User's Guide

Document Release Date: August 2010
Software Release Date: August 2010



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice. 081110

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notices

© Copyright 2010 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

Java™ is a US trademark of Sun Microsystems, Inc.

Acknowledgements

ANTLR, Apache, Bouncy Castle, GNU, Jaxen, Jython, Netaphor, MetaStuff, Radius, Sleepcat, TanukiSoftware

Documentation Updates

This guide's title page contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates, or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign-in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

You can visit the HP Software Support Online web site at:

<http://www.hp.com/go/hpsoftwaresupport>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software Support Online provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the HP Software Support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract.

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Table of Contents

Preface	7
Document Conventions	7
NA Documentation	8
Chapter 1: Getting Started	9
Terminology	9
Overview	10
Architecture	12
What is Horizontal Scalability?	13
NA Internal Monitoring: SQL Server Merge Replication	14
NA Internal Monitoring: Inter-NA Core Communication.....	15
Chapter 2: Installation, Setup, and Upgrading ...	17
System Requirements.....	18
Installation Planning.....	18
Protocols, Databases, and Ports	21
NA Server Installation Steps.....	22
Installation and Removal.....	22
Unpacking the Replication Script Installation Bundle	23
Creating a Two NA Core SQL Server Replication Environment	23
Preparation Requirements	23
Installation Steps.....	24
Deleting a Subscriber from a NA Mesh	27
System Requirements.....	27
Deletion Steps.....	27
Removing Replication	29
Upgrading from NAS 7.0/7.2 to NA 7.50	29
Upgrading from NA 7.5x to NA 7.60	32
Upgrading from NA 7.60 to NA 9.0	35
Chapter 3: System Administration	39
Getting Started.....	39
NA Generated Events	40
Distributed System - Uniqueness Conflict	40
Distributed System - Time Synchronization Warning	41
Distributed System - RMI Error.....	41
Distributed System - Stopped Merge Agent Job.....	42

Distributed System - Data Synchronization Delay Warning	42
Using the NA Distributed System Pages.....	43
Navigating to Distributed Systems Pages	43
Distributed Monitor Results Page.....	44
Distributed Conflict List.....	46
View Distributed Conflict Page	47
Site Reassignment Page	48
List Cores Page.....	48
Edit Core Page.....	49
Device Password Rule Priority Reset Page.....	50
Renew Configuration Options Page	50
Chapter 4: Failover and Recovery	51
Failover	51
Recovery	51
Loss of Network Connectivity.....	52
Loss of a NA Server.....	53
Loss of a Database Server.....	54
Chapter 5: Troubleshooting	57
SQL Server Replication Setup	57
Removing In-Memory and Database Information.....	58
Index	59

Preface

This document contains information on installing, configuring, and administering the HP Network Automation (NA) Distributed System on Microsoft SQL Server 2005 and SQL Server 2008.

Note: The NA Distributed System on SQL Server software requires SQL Server 2005, Service Pack 2 (Standard Edition or Enterprise Edition) or SQL Server 2008 (Standard Edition or Enterprise Edition.) Keep in mind that the NA Distributed System on SQL Server software can only be run on two NA Cores and no more than 6,500 devices can be managed.

Document Conventions

The following table explains the conventions used in this guide.

Convention	Description/Action
<i>Italic</i>	Used for system messages, paths, file names, and Web URLs. For example, <i>C:\HP\client\sdk\docs</i>
Link	Moves you from one location to another within the document, opens Web pages, or opens a new email message. In the guide, cross-references are contained within quotation marks and include a page number, while links to URLs and email addresses appear as underlined text.
Enter	Indicates that you should type the text or command that follows, then press the Enter key on the keyboard.
< >	Indicates variable information, such as a name or folder that you must supply. Do not include the angle brackets when replacing the placeholder.

NA Documentation

The core NA 9.0 Documentation Set includes:

- *HP Network Automation (NA) 9.0 User's Guide.*
- *HP Network Automation (NA) 9.0 Upgrade and Installation User's Guide.*
- *HP Network Automation (NA) 9.0 Release Notes.*
- Online HTML Help Files — To view the Online HTML Help files, after logging in to NA, click the Help icon at the top of any NA page.

Chapter 1: Getting Started

Use the following table to quickly locate information in this chapter.

Topic	Refer to:
Terminology	"Terminology" on page 9
Overview	"Overview" on page 10

Terminology

The following terms are used throughout this guide:

- **NA Core** — A single NA Management Engine, associated services (Syslog and TFTP), and a single database. A NA Core can manage multiple Partitions.
- **Partition** — A set of devices with unique IP addresses. A Partition is managed by one (and only one) NA Core. Multiple Partitions can be managed by a single NA Core. Refer to the *HP Network Automation 9.0 User's Guide* for information on segmenting devices.
- **NA Mesh** — Multiple NA Cores connected via replication.
- **Publisher** — A SQL Server database that defines what data is replicated and handles transferring changed data to and from the other databases in the NA Mesh.
- **Subscribers** — Databases in the NA Mesh that can send and receive replication data, but do not control the replication process.

Note: Both the Publisher and Subscribers will have the same data (subject only to replication time lags).

Overview

The NA Distributed System on SQL Server is a multimaster system where the data from each NA Core in a NA Mesh is accessible to all other NA Cores. This provides a comprehensive view of your data and allows for redundant data and failover in the event of a problem with a single NA Core.

The multimaster features include both database data and certain file system data, such as software images and device driver packages. Keep in mind that software images and device driver packages are also replicated across the NA Mesh.

The following comprises the NA Distributed System:

- A NA Core and a Partition:
 - A device is associated with a single Partition.
 - A Partition is associated with a single NA Core.
 - A task is associated with a specific NA Core.
- Merge replication on SQL Server:
 - Conflict resolution in the event that the same piece of data is modified almost simultaneously on two separate NA Cores. This conflict is typically resolved using the latest timestamp method.
 - Replication monitoring and conflict notification is built into NA. You can manage SQL Server replication conflicts and view merge agent job statuses from within the NA UI.
 - The NA scheduler is multi-core aware. You can schedule group tasks containing devices that are associated with different NA Cores. The system will run these tasks on the correct NA Core. You do not have to schedule tasks on the appropriate NA Core.

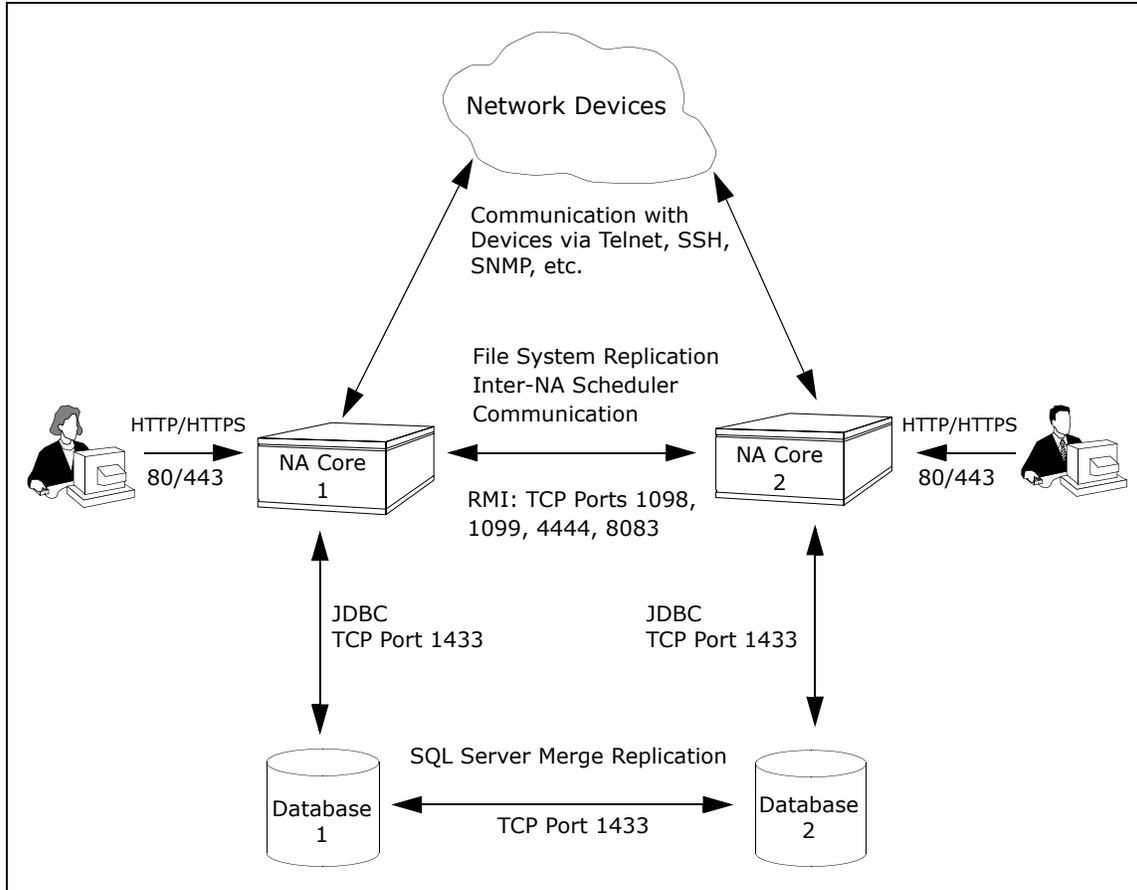
System setup requires a thorough understanding of SQL Server and NA. Installation includes number of steps that must be performed on the various servers.

In addition, certain network changes may be necessary to allow connections between the servers. Once setup is complete, you will need to partition your devices into Partitions to ensure proper NA Core access to devices. Refer to the *HP Network Automation 9.0 User's Guide* for information on segmenting devices.

Note: The NA Distributed System on SQL Server requires SQL Server 2005, Service Pack 2 (Standard Edition or Enterprise Edition) or SQL Server 2008 (Standard Edition or Enterprise Edition).

Architecture

The following figure shows a typical two-NA Core Multimaster Distributed System installation. The installation enables failover from one NA Core to another, while the remaining NA Core continues to manage all devices.



Keep in mind that some installations might have two separate network device collections that are inaccessible from the non-managing NA Core. These installations still provide data redundancy and scaling benefits, however.

Note that the database data transport uses standard SQL Server Merge Replication. For details on how replication works, refer to Chapter 3: System Administration.

The communication between the NA Cores is done using Java's Remote Method Invocation (RMI). RMI is used to ensure:

- Certain file system objects, such as software images and driver packages, are in sync.
- NA Tasks are scheduled and run on the correct NA Core.

NA does not replicate NA Core specific options, including:

- TFTP server information
- TACACS server information
- Scripting language settings
- Hop Box definitions
- Driver directory
- Install directory
- Local Gateway information
- SWIM proxy server
- Active Directory settings
- Scheduler (concurrent task limits)

What is Horizontal Scalability?

Horizontal Scalability is the ability to combine multiple NA Cores with a single database so that they work as a single logical unit to improve the performance of the overall system. NA Horizontal Scalability can be configured in different ways, depending on how you use your system and where bottlenecks are apt to occur. Refer to the *HP Network Automation 9.0 Horizontal Scalability User's Guide* for installation and configuration information.

NA Internal Monitoring: SQL Server Merge Replication

SQL Server Merge Replication can be configured to automatically handle certain types of conflicts. Where possible, NA uses SQL Server's built-in conflict handling software. For example, for most tables with *update/update* conflicts, NA installs the default SQL Server "latest timestamp wins" handler.

When it is not possible to use SQL Server's default handler, NA typically has additional code inside the application with logic that attempts to correct the error. If the error is uncorrected through that logic, NA alerts the NA administrator to the issue using NA's standard event notification emailing functionality. Refer to the *HP Network Automation 9.0 User's Guide* for information on configuring email notification.

Note: When a second NA Core is used as a standby for failover and not for normal operation, you should not encounter *update/update* conflicts.

Refer to "[Distributed Monitor Results Page](#)" on page 44 for detailed information on NA monitors.

In addition, in SQL Server, NA monitors on a regular schedule the following on each NA Core for the database transport (i.e., replication):

- Merge agent status — If the check fails and the merge agent is found to have stopped, NA will attempt to restart it.
- Merge agent delays — If the merge agent synchronization falls behind, NA will alert the NA Administrator via email notification.
- Unrecoverable uniqueness conflicts — NA alerts the NA Administrator in the event of unrecoverable uniqueness conflicts. Typically, there is no way to automatically correct these types of issues.

Note: Refer to the *HP Network Automation 9.0 User's Guide* for information on configuring email notification.

NA Internal Monitoring: Inter-NA Core Communication

In addition to replication monitoring, NA also monitors the following between each NA Core:

- RMI connectivity
- NA server timestamp differences

These monitors generate events during error conditions. The events can be emailed to the NA Administrator or SQL Server DBA using a standard event rule in NA. Refer to the *HP Network Automation 9.0 User's Guide* for information on configuring email notification. Refer to **"NA Generated Events" on page 40** for detailed information on NA system events.

Note: There is a standard example event rule shipped with NA. The event rule can be updated as necessary.

Chapter 2: Installation, Setup, and Upgrading

Use the following table to quickly locate information in this chapter.

Topic	Refer to:
System Requirements	"System Requirements" on page 18
Installation Planning	"Installation Planning" on page 18
Installation and Removal	"Installation and Removal" on page 22
Creating a Two NA Core SQL Server Replication Environment	"Creating a Two NA Core SQL Server Replication Environment" on page 23
Deleting a Subscriber from a NA Mesh	"Deleting a Subscriber from a NA Mesh" on page 27
Removing Replication	"Removing Replication" on page 29
Upgrading from NAS 7.0/7.2 to NA 7.50	"Upgrading from NAS 7.0/7.2 to NA 7.50" on page 29
Upgrading from NA 7.5x to NA 7.60	"Upgrading from NA 7.5x to NA 7.60" on page 32
Upgrading from NA 7.60 to NA 9.0	"Upgrading from NA 7.60 to NA 9.0" on page 35

System Requirements

The NA Distributed System on SQL Server software requires SQL Server 2005, Service Pack 2 (Standard Edition or Enterprise Edition) or SQL Server 2008 (Standard Edition or Enterprise Edition). Refer to the *HP Network Automation 9.0 Support Matrix* for information on database server hardware requirements.

Installation Planning

To properly install the Distributed System software, you must first complete:

- Device partitioning planning across NA Cores.
- Network configuration planning for connectivity between NA servers and SQL Servers in the NA Mesh.
- Network configuration planning for connectivity between NA servers and devices. For example, what network connectivity is required to support failover for device access? The ability of a NA Mesh to failover for device access depends in part on proper network setup to ensure access to devices. In some cases, you might not want to have failover work for complete device access, but instead have it ensure access to data while corrective action is taken to restore the network connectivity to the affected NA Core.
- Network configuration planning for connectivity and bandwidth between the different servers (NA and database) that comprise the NA Mesh and between the NA Management Engines and devices. The NA Cores in the NA Mesh will also need bandwidth between them equal to the bandwidth provided between a single NA server and its database in a single NA Core.
- Network configuration planning for bandwidth usage by the NA Mesh. The bandwidth required between a NA Management Engine and its database depends on the size of the deployment, including the number of devices, average configuration size, and the number of concurrent users. The bandwidth available between the SQL Server databases should be as large as between the NA Management Engine and the database. For a deployment of 14,000 devices with an average 10K configurations that change once a week, the bandwidth requirements

could conservatively be on the order of 1Mb/s. Peak bandwidth requirements could be higher, depending on usage.

- SQL Server setup planning. Access to SQL Server Management Studio and the *sqlcmd* command line tool are needed for initial database creation.

Keep in mind that during replication setup, a snapshot of the initial database is transferred to each database in the NA Mesh. This requires ample time, disk space, and bandwidth.

You can estimate the time it will take to copy data from server to server given the bandwidth between servers. You can also calculate the disk space requirements for the export (and subsequent import) operations by looking at the size of your database. If you want to export or import data from the same server as the database, the disk space requirement is twice the size of your database.

Note: Estimating time for the import and export operations could be difficult. You should allocate a lengthy time frame to complete this work. In addition, the NA server(s) must be off during the export and import steps. A sufficiently long outage window should be planned for.

You will also have to ensure:

- Time synchronization setup for the NA servers in the NA Mesh
- Users are instructed to login to their “closest” NA Core
- Access to a SQL Server DBA to support the NA Distributed System installation

To assist in planning, please note the following limitations and suggestions concerning the NA Distributed System:

- NA only supports two NA Cores in a Distributed System on SQL Server and can support no more than 6,500 nodes.
- NA currently does not support joining multiple NA installations into a NA Mesh. You can only create a NA Mesh from a single existing NA server, adding a new NA Core as appropriate. You can also create a NA Mesh from scratch.
- Users should not share logins. Due to the replication system used to share data across NA Cores, two users should never use the same login name to connect to two different NA Cores at the same time. If they do so, the system will likely require additional work by the system

administrator to ensure that the effected user's profile is properly synced up across the NA Mesh.

- NA currently assumes that all servers in the Mesh (NA and database) share a single timezone.
- Future NA upgrades will take longer and require more down-time due to the need to both update the replication setup and to update all servers in the NA Mesh. (NA does not support rolling upgrades where one part of the NA Mesh is running a version of NA while the rest of the NA Mesh is running a different version.)
- The *SQLServerReplicationScript.sql* script updates the RN_CORE table. When the replication script runs, the RN_CORE changes. There is no need for it to UPDATE and INSERT into the RN_CORE table on both Cores because replication is already running. The database changes will be pushed via replication to the second database. The contents of the RN_CORE tables on both databases should match whatever is setup in the *SQLServerReplicationScript.sql* script.
- To successfully recover from the loss of the Publisher and Subscriber server, you must reconfigure replication by using existing NA data. Refer to ["Loss of a Database Server" on page 54](#) for information on restoring databases.

Protocols, Databases, and Ports

NA communicates with devices using a combination of the following protocols, databases, and ports. If you use a given protocol, NA requires access to the corresponding port. Specifically, if NA communicates with devices protected by firewalls, these ports must be open. Refer to the figure on page 11 for information on the needed open ports.

Protocol/Database/Port	From/To
NA server (running the Mgmt Engine, Syslog, TFTP) and network devices	
Telnet (port 23)	From the NA server to network devices.
SSH (port 22)	From the NA server to network devices.
TFTP (port 69/udp)	From network devices to the NA server.
Syslog (port 514/udp)	From network devices to the NA server.
SNMP (port 161/udp)	From the NA server to network devices.
Between the database servers	
SQL Server (port 1433)	In a Distributed System configuration, the SQL Server processes connect to each other on port 1433.
Between the NA servers	
JNDI (ports 1098, 1099)	NA server to NA server. You can change this by editing the NA configuration files. Please contact Customer Support for assistance.
RMI (ports 4444 and 8083)	NA server to NA server. You can change this by editing the NA configuration files. Please contact Customer Support for assistance.
Between the NA server and the database server	
SQL Server (port 1433)	From the NA server to an SQL Server database server.
NA server and NA users	
HTTPS (port 443)	From the NA server to NA users. You can change this by editing the NA configuration files. Please contact Customer Support for assistance.

Protocol/Database/Port	From/To
Telnet (port 23 - Windows or 8023 - Solaris/Linux)	From the NA client to the NA server. This can be changed from the Administrative Settings option.
SSH (port 22 - Windows or 8022 - Solaris/Linux)	From the NA client to the NA server. This can be changed from the Administrative Settings option.

NA Server Installation Steps

For complete information on installing NA, refer to the *HP Network Automation 9.0 Upgrade and Installation Guide*. Keep in mind that you will have to:

- Verify time synchronization between Distributed Systems.
- Create a second database using the same database user account credentials. You do not need to use the same database name, however, for setting up replication.

Note: NA service accounts do not have to have different usernames for setting up replication.

Installation and Removal

This section includes information on:

- Creating a two NA Core SQL Server replication environment
- Deleting a Subscriber from a NA Mesh
- Removing replication

Note: Refer to ["Chapter 5: Troubleshooting" on page 57](#) for information on Troubleshooting the installation.

Unpacking the Replication Script Installation Bundle

The setup files for the NA Distributed System are the standard setup files for a normal NA installation, with the addition of the *SQLServerReplicationScriptToolBundle.zip* file. This file should be unzipped onto a standard NA server. The setup files include a Java program that will be run to generate replication setup SQL scripts. This program will need to use the Java run-time that is available on the NA server.

Creating a Two NA Core SQL Server Replication Environment

When creating a two NA Core SQL Server replication environment, NA Core 1 is the Publisher and NA Core 2 is the Subscriber.

Preparation Requirements

You must have the following items configured before creating a SQL Server replication environment:

- Both SQL Server Distributed Systems must be at SQL Server 2005, Service Pack 2, or SQL Server 2008.
- A current NA 7.60 database on NA Core 1.
- A NA server connected to the database on NA Core 1.
- The SQL Server agent service running on NA Core 1 (Publisher) database server.
- An empty database (with no data or NA schema setup) on NA Core 2.
- The empty database on NA Core 2 needs a 'dbowner' login account.
- A network connection from the NA Core 1 servers to the NA Core 2 servers (and vice-versa) that enable ports 1433 and 1099 (or appropriate variations) to be connected between the servers.
- sqlcmd access to NA Core 1 and NA Core 2. (**Note:** You will need to supply credentials for a login that is a member of the sysadmin role when running the scripts.)
- The SQLServerReplicationScriptTool application installed on a Java-capable system.

Installation Steps

Do the following to create a SQL Server replication environment:

1. Set up a shared directory that is accessible from both NA Core 1 and NA Core 2.
2. Collect the following information:
 - Login name and password of a SQL Server login that is a member of sysadmin on NA Core 1 and NA Core 2.
 - Login name and password of a Windows account under which SQLServer agents can run.
 - Database name, NA server hostname, NA server RMI listening port, database hostname, and database listening port for NA Core 1 and NA Core 2.
 - The time zone offset (integer from UTC) for the entire NA Mesh. This must be a constant across the NA Mesh. Do not consider daylight savings time when setting this value.
 - UNC path for the shared directory you setup in Installation Step 1.
3. Turn off the NA server that is accessing the database on NA Core 1.
4. Update the variables for NA Core 1 and NA Core 2 in the *SQLServerReplicationScriptTool.properties* file. In addition, update the timezone offset in that file. Make sure the "mode" property is set to "initial". These properties are described in detail in the file. Set *replication.data.dir* to be the directory you set up in Installation Step 1. Enter the "Windows username" value as "machine name\administrator" or "domain\domainUser". If you use "machine name\administrator", make sure the administrator can access the shared directory setup in Step 1 above.

5. Set the java path on NA Core 1 and run the script tool:
 - `C:\Program Files\Microsoft SQL Server\90> set JAVA_HOME = C:\Rendition\jre`
 - `C:\Program Files\Microsoft SQL Server\90>set PATH = %PATH% ; C:\Rendition\jre\bin`
 - Enter `java` at the prompt to verify that value returned is Java 1.6 or higher.
 - Run the script tool. This will output two files. Enter:
`java -classpath . SQLServerReplicationScriptTool`
6. Copy the *SQLServerPreSnapshotScript.sql* file to the shared directory that is accessible from both NA Core 1 and NA Core 2. The file defaults to *SQLServerPreSnapshotScript.sql*. The name is specified by the "pre-snapshot.file" property in the *SQLServerReplicationScriptTool.properties* file.
7. Run the *SQLServerReplicationScript.sql* file using the 'sqlcmd' command: `sqlcmd -S <PublisherServer> -U <USER> -P <PASSWORD> -i SQLServerReplicationScript.sql -o log.txt`

Where <PublisherServer> is the Publisher's server, and <USER> and <PASSWORD> are the credentials for a login that is a member of the sysadmin role. (Refer to the eighth bullet in "Preparation Requirements" on page 23.)

8. Using sqlcmd, enter the following t-sql query on both the NA Core 1 database and the NA Core 2 database.

NA Core 1:

```
> sqlcmd -S PublisherServer -U sa -P sa password
1> use <DatabaseName>
2> Select count(*) from Information_Schema.Tables where
Table_Type = 'Base Table' and Table_Name like 'RN_%'
3> go
```

NA Core 2:

```
> sqlcmd -S SubscriberServer -U sa -P sa password
1> use <Databasename>
2> Select count(*) from Information_Schema.Tables where
Table_Type = 'Base Table' and Table_Name like 'RN_%'
3> go
```

Once the results of the queries are the same on both NA Cores, the replication has been setup between the NA Cores.

9. Go to RN_CORE within SQL database on each NA Core and verify that the file sizes are the same. Note that <DatabaseName> is the appropriate database name for each server.)
10. Install a NA server. Be sure it points to an empty database on NA Core 2. When prompted, have NA connect to the database using the login from the sixth bullet in the "Preparation Requirements" on page 23 section.
11. Verify that the replication is working correctly. Make sure:
 - The log file generated by running the replication SQL script contains no errors.
 - The RN_CORE table on each database contains the appropriate list of servers in the NA Mesh. To check the RN_CORE table, using SQLPlus, enter the following command on each server:
"SELECT * FROM RN_CORE";
 - Verify that the replication of data is working correctly. To verify that replication is working between all NA Cores in the NA Mesh, enter the following commands:

On the Publisher, enter:

```
UPDATE RN_CORE SET COMMENTS = 'Test from db1 to db2' WHERE  
CoreID=1;
```

wait at least 60 seconds

On the Subscriber, enter:

```
SELECT Comments from RN_CORE where CoreID=1;
```

You should then reverse this, executing the update on the Subscriber with a new comment and the select on the Publisher. Be sure to execute this from every NA Core in the NA Mesh. Check the values each time on all of the other NA Cores to ensure that the replication between each database in the system are working correctly.

12. Add the *distributed.rcx* file to each NA server in the NA Mesh (in the same location as *appserver.rcx*).

13. Start the NA server on NA Core 1.

14. Start the NA server on NA Core 2.

Deleting a Subscriber from a NA Mesh

When deleting a subscriber from a NA Mesh, NA Core 1 is the Publisher and NA Core 2 is the Subscriber you want to remove.

System Requirements

You must have the following items configured before deleting a Subscriber from a NA Mesh:

- sqlcmd access to NA Core 1 and NA Core 2.
- The *SQLServerReplicationScriptTool* application installed on a Java-capable system.

Deletion Steps

To delete a Subscriber from a NA Mesh:

1. Collect the following information:
 - Login name and password of a SQL Server login that is a member of sysadmin on NA Core 1 and NA Core 2.
 - Database name, NA server hostname, NA server RMI listening port, database hostname, and database listening port for NA Core 1 and NA Core 2.
2. Ensure that all devices in NA belong to partitions on NA Cores that are not going to be removed.
3. Modify all partitions to point to a NA Core that is not being removed. Alternatively, remove those partitions.
4. Turn off the NA server from the NA Core that is being removed.

5. Delete the RN_CORE entry that was removed using the following commands using sqlcmd:

Note: To identify the Core ID of the NA Core to delete, enter:

```
SELECT * FROM RN_CORE;
```

```
UPDATE RN_PARTITION SET OwningCoreID = 1 WHERE OwningCoreID = <coreID>;  
UPDATE RN_PARTITION SET ManagingCoreID = 1 WHERE ManagingCoreID = <coreID>;  
UPDATE RN_SCHEDULE_TASK SET CoreID = 1 WHERE CoreID = <coreID>;  
DELETE FROM RN_CORE WHERE CoreID = <coreID>;
```

6. Update the variables for NA Core 2 in the *SQLServerReplicationScriptTool.properties* file. Make sure that the variables for NA Core 1 are correct. Make sure the "mode" property is set to "delete_server". These properties are described in detail in the file.
7. Run the script tool. Enter: `java -cp . SQLServerReplicationScriptTool`
8. Run the first output file using sqlcmd with a login that is a member of the sysadmin role on NA Core 1. Enter:

```
sqlcmd -S <PublisherServer> -U <USER> -P <PASSWORD> -i  
SQLServerReplicationScript.sql -o log.txt
```

Where <PublisherServer> is the Publisher's server, and <USER> and <PASSWORD> were collected in step 1.

9. Delete the *distributed.rcx* file on the NA server if it is the only NA Core in the NA Mesh. Restart the NA server on NA Core 1.

Removing Replication

To remove replication:

1. Remove each subscriber from the NA Mesh by following the steps above for each one.
2. Enter the following script on the publisher:

```
use [master]
exec sp_dropdistributor @no_checks = 1
GO
```

Note: After removing replication, edit the RN_CORE entry. Refer to ["Deletion Steps" on page 27](#) for information.

Upgrading from NAS 7.0/7.2 to NA 7.50

Note: To update a system to NAS 7.0, refer to NAS 7.0 documentation.

During the NA upgrade process, the NA Cores must be completely off-line. Keep in mind that the upgrade instructions only work on a current, up-to-date NAS 7.0 or NAS 7.2 database running on all servers.

You will be upgrading the:

- NA application servers
- NA database schema
- Replication between the database cores

Note: You will need to update each NA Core in your system.

You must have the following configured before upgrading a NA Mesh:

- sqlcmd access to NA Core 1 and NA Core 2
- The *SQLServerReplicationScriptTool* application installed on a Java-capable system

To upgrade a NA Mesh:

1. Turn off all NA servers in the NA Mesh.
2. Collect the following information:
 - a) Login name and password of a SQL Server login that is a member of sysadmin on NA Core 1 and NA Core 2.
 - b) Database name, NA server hostname, NA server RMI listening port, database hostname, and database listening port for NA Core 1 and NA Core 2.
3. Update the variables for NA Core 2 in the *SQLServerReplicationScriptTool.properties* file. Make sure that the variables for NA Core 1 are correct. Make sure the "mode" property is set to "upgrade_from_7_0" or "upgrade_from_7_2". These properties are described in detail in the file.

4. Run the script tool. Enter:

```
java -cp . SQLServerReplicationScriptTool
```

5. Run the first output file using sqlcmd with a login that is a member of the sysadmin role on NA Core 1. Enter:

```
sqlcmd -S <PublisherServer> -U <user> -P <password> -i  
SQLServerReplicationScript.sql -o log.txt
```

Where <PublisherServer> is the server of Publisher, <user> and <password> were collected in Step 2.

6. Verify that the replication is working correctly. Make sure:
 - The log file generated by running the replication SQL script contains no errors.
 - The RN_CORE table on each database contains the appropriate list of servers in the NA Mesh. To check the RN_CORE table, using SQLPlus, enter the following command on each server:

```
SELECT * FROM RN_CORE;
```

- Verify that the replication of data is working correctly. To verify that replication is working between all NA Cores in the NA Mesh, enter the following commands:

On the Publisher, enter:

```
UPDATE RN_CORE SET COMMENTS = 'Test from db1 to db2' WHERE  
CoreID=1;
```

wait at least 60 seconds

On the Subscriber, enter:

```
SELECT Comments from RN_CORE where CoreID=1;
```

You should then reverse this, executing the update on the Subscriber with a new comment and the select on the Publisher. Be sure to execute this from every NA Core in the NA Mesh. Check the values each time on all of the other NA Cores to ensure that the replication between each database in the system are working correctly.

7. Upgrade all NA servers using the NA 7.5 Service Pack Installer. For information on running the NA 7.5 Service Pack Installer, refer to the *HP Network Automation 7.5 Upgrade and Installation Guide*.
8. Restart the NA servers in the NA Mesh.

Upgrading from NA 7.5x to NA 7.60

To update a system to NA 7.60, you must first upgrade to NA 7.50 or 7.5.02.

During the NA upgrade process, the NA Cores must be completely off-line. Keep in mind that the upgrade instructions only work on a current, up-to-date NA 7.5x database running on all servers.

You will be upgrading the:

- NA application servers
- NA database schema
- Replication between the database cores

Note: You will need to update each NA Core in your system.

You must have the following configured before upgrading a NA Mesh:

- sqlcmd access to NA Core 1 and NA Core 2
- The *SQLServerReplicationScriptTool* application installed on a Java-capable system

To upgrade a NA Mesh:

1. Turn off all NA servers in the NA Mesh.
2. Collect the following information:
 - a) Login name and password of a SQL Server login that is a member of sysadmin on NA Core 1 and NA Core 2.
 - b) Database name, NA server hostname, NA server RMI listening port, database hostname, and database listening port for NA Core 1 and NA Core 2.
3. Update the variables for NA Core 2 in the *SQLServerReplicationScriptTool.properties* file. Make sure that the variables for NA Core 1 are correct. Make sure the "mode" property is set to "upgrade_from_7_5". These properties are described in detail in the file.

4. Run the script tool. Note that you must run this on a system that has Java 1.6.x installed. Enter:

```
java -cp . SQLServerReplicationScriptTool
```

5. Run the first output file using sqlcmd with a login that is a member of the sysadmin role on NA Core 1. Enter:

```
sqlcmd -S <PublisherServer> -U <user> -P <password> -i  
SQLServerReplicationScript.sql -o log.txt
```

Where <PublisherServer> is the server of Publisher, <user> and <password> were collected in Step 2.

6. Verify that the replication is working correctly. Make sure:
 - The log file generated by running the replication SQL script contains no errors.
 - The RN_CORE table on each database contains the appropriate list of servers in the NA Mesh. To check the RN_CORE table, using SQLPlus, enter the following command on each server:

```
SELECT * FROM RN_CORE;
```

- Verify that the replication of data is working correctly. To verify that replication is working between all NA Cores in the NA Mesh, enter the following commands:

On the Publisher, enter:

```
UPDATE RN_CORE SET COMMENTS = 'Test from db1 to db2' WHERE  
CoreID=1;
```

wait at least 60 seconds

On the Subscriber, enter:

```
SELECT Comments from RN_CORE where CoreID=1;
```

You should then reverse this, executing the update on the Subscriber with a new comment and the select on the Publisher. Be sure to execute this from every NA Core in the NA Mesh. Check the values each time on all of the other NA Cores to ensure that the replication between each database in the system are working correctly.

7. Upgrade all NA servers using the NA 7.60 Service Pack Installer. For information on running the NA 7.60 Service Pack Installer, refer to the *HP Network Automation 7.60 Upgrade and Installation Guide*.
8. Restart the NA servers in the NA Mesh.

Upgrading from NA 7.60 to NA 9.0

During the NA upgrade process, the NA Cores must be completely off-line. Keep in mind that the upgrade instructions only work on a current, up-to-date NA 7.60 database running on all servers.

You will be upgrading the:

- NA application servers
- NA database schema
- Replication between the database cores

Note: You will need to update each NA Core in your system.

You must have the following configured before upgrading a NA Mesh:

- sqlcmd access to NA Core 1 and NA Core 2
- The *SQLServerReplicationScriptTool* application installed on a Java-capable system

To upgrade a NA Mesh:

1. Turn off all NA servers in the NA Mesh.
2. Collect the following information:
 - a) Login name and password of a SQL Server login that is a member of sysadmin on NA Core 1 and NA Core 2.
 - b) Database name, NA server hostname, NA server RMI listening port, database hostname, and database listening port for NA Core 1 and NA Core 2.
3. Update the variables for NA Core 2 in the *SQLServerReplicationScriptTool.properties* file. Make sure that the variables for NA Core 1 are correct. Make sure the "mode" property is set to "upgrade_from_7_6". These properties are described in detail in the file.

4. Run the script tool. Note that you must run this on a system that has Java 1.6.x installed. Enter:

```
java -cp . SQLServerReplicationScriptTool
```

5. Run the first output file using sqlcmd with a login that is a member of the sysadmin role on NA Core 1. Enter:

```
sqlcmd -S <PublisherServer> -U <user> -P <password> -i  
SQLServerReplicationScript.sql -o log.txt
```

Where <PublisherServer> is the server of Publisher, <user> and <password> were collected in Step 2.

6. Verify that the replication is working correctly. Make sure:
 - The log file generated by running the replication SQL script contains no errors.
 - The RN_CORE table on each database contains the appropriate list of servers in the NA Mesh. To check the RN_CORE table, using SQLPlus, enter the following command on each server:

```
SELECT * FROM RN_CORE;
```

- Verify that the replication of data is working correctly. To verify that replication is working between all NA Cores in the NA Mesh, enter the following commands:

On the Publisher, enter:

```
UPDATE RN_CORE SET COMMENTS = 'Test from db1 to db2' WHERE  
CoreID=1;
```

wait at least 60 seconds

On the Subscriber, enter:

```
SELECT Comments from RN_CORE where CoreID=1;
```

You should then reverse this, executing the update on the Subscriber with a new comment and the select on the Publisher. Be sure to execute this from every NA Core in the NA Mesh. Check the values each time on all of the other NA Cores to ensure that the replication between each database in the system are working correctly.

7. Upgrade all NA servers using the NA 9.0 Service Pack Installer. For information on running the NA 9.0 Service Pack Installer, refer to the *HP Network Automation 9.0 Upgrade and Installation Guide*.
8. Restart the NA servers in the NA Mesh.

Chapter 3: System Administration

Use the following table to quickly locate information in this chapter.

Topic	Refer to:
Getting Started	"Getting Started" on page 39
NA Generated Events	"NA Generated Events" on page 40
Distributed Monitor Results	"Distributed Monitor Results Page" on page 44
Site Reassignment page	"Site Reassignment Page" on page 48
List Cores page	"List Cores Page" on page 48
Edit Core page	"Edit Core Page" on page 49
Device Password Rule Priority Reset page	"Device Password Rule Priority Reset Page" on page 50
Renew Configuration Options page	"Renew Configuration Options Page" on page 50

Getting Started

In general, a NA server that is part of a Distributed NA Mesh should be transparent to users. However, there are a number of operations that the system administrator may need to do to keep the Distributed NA Mesh functioning properly.

NA Generated Events

By default, NA generates system events. Event rules can alert you to certain error conditions requiring attention. Each event is listed below, along with an explanation and required action to be taken.

Distributed System - Uniqueness Conflict

Event format:

rowguid: <the guid of the database row that had the conflict>
origin_datasource: <database server>.<database name>
reason_text: <a description of why the conflict occurred>
conflict_type: <type of conflict according to SQLServer>
reason_code: <error message from SQLServer, depends on the type of conflict>
repl_create_time: <time the conflict was generated by SQLServer>

conflict_table: <where SQLServer stores the conflicting data>

dataTable: <NA table that contains the conflicting data>

SQLServerConflictID: <ID of the conflict recorded by NA>

status: <status>

Conflicting Data: <the columns that are conflicting>

Refer to the *SQL Server Replication* documentation for instructions on correcting this conflict.

This event is sent when NA detects a conflict in a uniqueness constraint. You will receive an event per NA Core, since the conflicts are local to each NA Core. To correct a naming conflict, go to one NA Core and update the names for the affected objects. Both the renamed <NAME>.<SID> and <NAME> should be edited to force an update on the other NA Cores.

To correct a rule priority conflict, go to the Device Password Rule Priority Reset page click the Reset Priority button. Refer to ["Device Password Rule Priority Reset Page" on page 50](#).

If this does not solve the problem, you will need to manually edit the rules on each NA Core, setting the priority order correctly and verifying existence of correct rules. When finished, return to the Device Password Rule Priority Reset page and click the Reset Priority button.

Distributed System - Time Synchronization Warning

Event format:

```
Time difference: <N> seconds  
Local Core: <hostname>  
Remote Core: <hostname>
```

NA replication conflict resolution depends on a latest timestamp method. To work correctly, this requires different NA servers' clocks to differ by only a small amount. To correct this problem, make sure that the time is synchronized on the NA server systems across the NA Mesh.

Distributed System - RMI Error

Event format:

```
Local Core: <hostname>  
Remote Core: <hostname>  
Error: <Exception text>
```

This error typically occurs when there are network problems between the NA servers. To troubleshoot this problem, make sure:

1. The host that the server cannot connect to is up and running.
2. The NA instance on that host is running.
3. From a command line, enter `ping <host>` to ensure that network connectivity exists between servers.
4. From a command line, enter `telnet <host>` to port 1099 (or whatever your RMI listen port is set to) to ensure that RMI connections are being accepted. If working correctly, you should get back some data that includes the text string "java.rmi.MarshalledObject".

Failures of any of these steps will point to corrective actions needed, such as updating the RMI port being used in the Edit NA Core page, or restarting NA to make sure that the RMI port has been bound correctly and is not being used by another application.

Distributed System - Stopped Merge Agent Job

Event format:

```
Merge Agent Job Stopped  
<Job Details>  
[Successfully attempted to restart the job.]
```

NA monitors the replication merge agent and sends the event if it determines that the agent has stopped. If the "Successfully attempted to restart the job" message is displayed, NA restarted the agent. The agent's status can be reviewed in SQLServer Management Studio.

Distributed System - Data Synchronization Delay Warning

Event format:

```
Publisher '<DatabaseName>' has been unable to push changes  
to subscriber '<DatabaseName2>' since <tim>. If you wish to  
avoid potential data loss, please rectify this as soon as  
possible.
```

```
Subscriber '<DatabaseName2>' has been unable to push changes  
to Publisher '<DatabaseName>' since <tim>. If you wish to  
avoid potential data loss, please rectify this as soon as  
possible.
```

NA has detected that the replication system is having trouble keeping the data and the databases in sync. The databases, merge agents, and network connections between the databases should be examined for issues.

Using the NA Distributed System Pages

When you install the Distributed System software, the NA user interface includes specific Distributed System pages to help you monitor and administer the system.

Navigating to Distributed Systems Pages

The screenshot shows the HP Network Automation user interface. The top navigation bar includes the HP logo, the text "HP Network Automation", and a "Logout" button. Below the navigation bar are several menu items: "Devices", "Tasks", "Policies", "Reports", "Admin", and "Help". The "Admin" menu is expanded, showing a list of options. A black arrow points from the "Admin" menu item to the expanded menu. Another black arrow points from the text "Distributed System Pages" to the "Distributed" section of the expanded menu.

HP Network Automation		Logout
Devices ▾	Tasks ▾	Policies ▾
Reports ▾	Admin ▾	Help ▾
<ul style="list-style-type: none"> Users User Groups New User New User Group Logged on Users 		
User Roles & Permissions		
<ul style="list-style-type: none"> Security Partitions Device Password Rules Event Notification & Response Rules 		
<ul style="list-style-type: none"> Custom Data Setup LDAP Setup Workflow Setup 		
Administrative Settings ▶		
Distributed ▶		
<ul style="list-style-type: none"> Monitor Results Error List Conflict List Site Reassignment Core List Device Password Rule Priority Reset Renew Configuration Options 		
<ul style="list-style-type: none"> Task Load System Status Start/Stop Services Troubleshooting Drivers 		
About HP Network Automation		
System Task ▶		

Distributed Monitor Results Page

The Distributed Monitor Results page displays the overall health of the Distributed System.

To open the Distributed Monitor Results page, on the menu bar under Admin select Distributed and click Monitor Results. The Distributed Monitor Results page opens.

NA monitor several properties necessary for proper functioning of the Distributed System, including:

- **RMI Connections** — RMI (Remote Method Invocation) is Java's remote procedure call protocol. The distributed system makes RMI calls between NA servers in the NA Mesh to transfer information about scheduled tasks, system settings, software images, and so on.
- **Uniqueness Conflicts** — Certain NA database constraints restrict columns to unique values. In a distributed environment, these constraints can be violated when updates are made on two different NA Cores where the unique column is set to the same value. These conditions are captured by the Replication Conflict Resolution System and logged. NA cannot automatically resolve these conflicts. They must be resolved manually.
- **Merge Agents** — Merge Agents are the processes at the Publisher that handle transferring replicated data. NA monitors the SQL Server jobs that schedule these processes. If for some reason the process stops, NA reports that here. Stopped Merge Agent jobs should be restarted as soon as possible.
- **Local NA Core Definition** — The local NA Core must be able to determine which entry in the RN_CORE table it is. If the "The *local core for this system is undefined.*" error message is displayed, the CoreHostname property needs to be updated for the NA Core. This can be done using the Edit Core page.

Note: When this condition occurs, the NA Management Engine's log file will contain the following text: "*Fatal error - could not assign local core.*"

The CoreHostname value can be either the DNS, *etc/hosts* value, or an IP address. If you are using a NA server with multiple IP addresses, you might need to tell NA which IP address to use. This is done by adding the following setting to the *distributed.rcx* file:

```
<option name="distributed/NA_server_local_ip">A.B.C.D</option>
```

Note: The *distributed.rcx* file is located in the same location as the *appserver.rcx* file.

The value A.B.C.D should be replaced with the appropriate NAT IP address for the NA server and should match the RN_CORE table's CoreHostname value for that NA Core.

Distributed Conflict List

The Distributed Conflict List page displays the uniqueness constraint conflict list. This provides information about uniqueness conflicts that will need to be manually corrected to ensure that the databases in the NA Mesh are in sync.

To open the Distributed Conflict List, on the menu bar under Admin select Distributed and click Conflict List. The Distributed Conflict List opens.

Note: Conflicts are currently viewable only at the Publisher. It is recommended that you set up an Event Rule to email you whenever a "Distributed System - Uniqueness Conflict" event is generated. Refer to the *HP Network Automation 9.0 User's Guide* for information on creating Event Rules.

Distributed Conflict List Page Fields

Field	Description
origin_datasource	The database on which the conflict occurred.
Table	The table on which the conflict occurred.
rowguid	The guid of the row on which the conflict occurred.
Status	Status is "event_generated" if the system has sent an alert that this conflict exists.
Actions	You can select the following options: <ul style="list-style-type: none">• Detail — Opens the View Distributed Conflict page, where you can view details on an individual uniqueness constraint. Refer to "View Distributed Conflict Page" on page 47.• Delete — Deletes the conflict from the database.

View Distributed Conflict Page

The View Distributed Conflict page provides details on a specific uniqueness constraint.

To open the View Distributed Conflict page:

1. On the menu bar under Admin select Distributed and click Conflict List. The Distributed Conflict List opens.
2. In the Actions column, click the Detail option. The View Distributed Conflict page opens.

View Distributed Conflict Page Fields

Field	Description
origin_datasource	The database on which the conflict occurred.
Table	The table on which the conflict occurred.
Conflicting Data	The columns and their values that are causing the conflict.
reason_text	A description of why the conflict occurred.
rowguid	The guid of the row on which the conflict occurred.
conflict type	The type of conflict according to the SQL Server.
reason_code	Depending on the type of conflict, the error message from the SQL Server.
MS_repl_create_time	The time the conflict was generated by the SQL Server.
conflict_table	The location where the SQL Server stores the conflicting rows.
Status	Status is "event_generated" if the system has sent an alert that this conflict exists.

Site Reassignment Page

The Site Reassignment page allows the Site-to-NA Core mapping to be modified. This is useful for failover of Sites from one NA Core to another and for restoring Sites back to their original NA Core.

To open the Site Reassignment page, on the menu bar under Admin select Distributed and click Site Reassignment. The Site Reassignment opens. You can select NA Cores from the drop-down menu.

List Cores Page

The List Cores page lists all NA Cores in the NA Mesh. This page provides information to properly manage the Distributed System.

To open the List Cores page, on the menu bar under Admin select Distributed and click Core List. The List Cores page opens.

List Cores Page Fields

Field	Description
Name	The NA Core's name.
Core Hostname	The hostname of the NA Core's NA server.
Is Master?	Is the NA Core the Publisher? (Yes or No)
Timezone Offset	The timezone offset of the actual NA Core server.
Status	Currently, there is only Normal status.
Realm	The default Realm for the NA Core.
Actions	You can select the following option: <ul style="list-style-type: none">• Edit — Open the Edit Core page. Refer to "Edit Core Page" on page 49.

Edit Core Page

The Edit Core page enables you to edit the NA Core definition.

To open the Edit Core page:

1. On the menu bar under Admin select Distributed and click Core List. The List Cores page opens.
2. In the Actions column, click the Edit option. The Edit Core page opens.

You can complete the following fields and click the Save Core button:

- Name — Enter the NA Core name.
- Database Identifier — Enter the database name. This information is needed to make connections to a particular SQL Server instance on a server.
- Core Hostname - Enter the hostname of this NA Core's server.
- RMI Port — Enter the RMI port. RMI (Remote Method Invocation) is Java's remote procedure call protocol. The distributed system makes RMI calls between NA servers in the NA Mesh to transfer information about scheduled tasks, system settings, software images, and so on.
- Database hostname — Enter the Database hostname.
- Database Port — Enter the port on the database server with which NA communicates with the database.
- Timezone Offset — Select a Timezone offset from the drop-down menu.
- Replication Admin User — Enter the name of the Replication Admin user. The Replication Admin user is created and used by the SQL Server database to manage replication.
- Replication Password — Not applicable for SQL Server.
- Confirm Replication Password — Not applicable for SQL Server.
- Comments — Add any comments about the NA Core.
- Realm Name — Enter the Realm in which the NA Core resides. For information on segmenting devices, refer to the *HP Network Automation 9.0 User's Guide*.

Device Password Rule Priority Reset Page

The Device Password Rule Priority Reset page enables you to reset device password rule priorities in the event that a uniqueness constraint conflict occurs for those objects.

To open the Reset Password Priority page, on the menu bar under Admin select Distributed and click Device Password Rule Priority Reset. The Device Password Rule Priority Reset page opens.

Click the Reset Priority button to reset the device password rule priorities.

Renew Configuration Options Page

The Renew Configuration Options page enables you to reset the configuration options when the configuration options on a NA Core become out-of-sync with other servers in the NA Mesh.

To open the Renew Configuration Options page, on the menu bar under Admin select Distributed and click Renew Configuration Options. The Renew Configurations Options page opens.

Click the Renew Config Options button to ensure that all options on the NA Core are in sync with the rest of the NA Mesh.

Chapter 4: Failover and Recovery

Use the following table to quickly locate information in this chapter.

Topic	Refer to:
Failover	"Failover" on page 51
Recovery	"Recovery" on page 51
Loss of Network Connectivity	"Loss of Network Connectivity" on page 52
Loss of a NA Server	"Loss of a NA Server" on page 53
Loss of a Database Server	"Loss of a Database Server" on page 54

Failover

When the network has been configured to failover, if a NA Core fails, users can continue to access all data in the system using a different NA Core. All Sites that were originally managed by the failed NA Core can be pointed to a new NA Core using the Site Reassignment page. Refer to ["Site Reassignment Page" on page 48](#) for information.

Note: Procedures for system recovery will vary depending on how the remote server failed.

Recovery

There are three basic recovery scenarios:

- Loss of network connectivity
- Loss of a NA server
- Loss of a Database server

Loss of Network Connectivity

In the case of lost network connectivity, failover occurred due to network issues. No problems occurred with the NA server or with the SQL Server database server. Recovery consists of the following steps:

1. Resolve the network issues.
2. Reset Sites that had been reassigned back to their original NA Core. This can be accomplished in NA. Refer to ["Site Reassignment Page" on page 48](#).
3. If any drivers have been added to the system during the outage, click the "Reload Drivers" button on the Start/Stop Services page. This action reloads the driver files and pushes them across to other NA Cores in the NA Mesh. This action should be performed on the NA server where the drivers were added. Refer to the *HP Network Automation 9.0 User's Guide* for information.
4. If any system settings have been modified during the outage, use the "Renew Config Options" page to make sure options are synced across the NA Mesh. Refer to ["Renew Configuration Options Page" on page 50](#).
5. If any NA Cores have lost connectivity for a long period of time, restart the NA Core server that lost connectivity after data sync so as to reload certain Site data and avoid exception errors due to obsolete data.

Once the network issues are resolved, the system should recover as replication syncs data between the databases.

Loss of a NA Server

In cases where the NA server suffers a failure that requires re-installation of the NA server, recovery consists of the following steps:

1. During NA installation, select the "Use existing database" option. In addition, the database selected should be the one the failed server was previously using.
2. Add the *distributed.rcx* file from the *ReplicationScriptToolBundle* to the directory where the *appserver.rcx* file resides.
3. Re-add any NA Core-specific special case options for patches and support issues.
4. Restart NA.
5. Reset Sites that have been re-assigned back to their original NA Core. This can be accomplished in NA. Refer to ["Site Reassignment Page" on page 48](#).
6. If any drivers have been added to the system during the outage, click the "Reload Drivers" button on the Start/Stop Services page. This action reloads the driver files and pushes them across to other NA Cores in the NA Mesh. This action should be performed on the NA server where the drivers were added. Refer to the *HP Network Automation 9.0 User's Guide* for information.
7. If any system settings have been modified during the outage, use the "Renew Config Options" page to make sure options are synced across the NA Mesh. Refer to ["Renew Configuration Options Page" on page 50](#).
8. Edit the original NA Core to modify any parameters that may be different (perhaps the installation happened on a new server with a different hostname).
9. Copy the software images repository from a good NA Core to the recovered NA Core.

Loss of a Database Server

In the case of a lost database server, the NA server is still running, but cannot access the database server. The database server will need to be rebuilt and replication must be reconfigured on the database server. Recovery consists of the following steps:

1. Pause or delete any tasks that appear to be waiting or pending, or not running because they are associated with the lost database server. You can perform this action on another working NA server in the NA Mesh.
2. If the lost database server is a Subscriber, replication needs to be removed from the NA Mesh if the lost Subscriber is the only Subscriber in the NA Mesh. Refer to ["Removing Replication" on page 29](#) for information. Keep in mind that Replication can be reconfigured after removal. Refer to ["Creating a Two NA Core SQL Server Replication Environment" on page 23](#) for information.
3. If the lost database server is the Publisher, replication needs to be reinstalled.
 - a) Export all RN_% tables from the Subscriber database.
 - b) Create a new database on the NA Core 1 database server with data exported from the Subscriber database.
 - c) Reinstall NA on the NA Core 1 Server. Be sure NA points to the existing database on NA Core 1.
 - d) Edit the RN_CORE entry. Refer ["Deletion Steps" on page 27](#) (step 5).
 - e) Install replication. Refer to ["Creating a Two NA Core SQL Server Replication Environment" on page 23](#) for information.
4. Make sure the SQL Server Agent is running on the restored database server. Open a command prompt on the restored database server and enter: **net start**. You should see `SQL Server Agent (MSSQLSERVER)` in the output. If you do not, enter: **net start SQL Server Agent <MSSQLSERVER>**.
5. Reset all sites that have been reassigned back to their original NA Core. Refer to ["Site Reassignment Page" on page 48](#) for information.
6. Login to NA and edit the NA Core that was recovered to make sure all information is correct for the new setup. Refer to the *HP Network Automation 9.0 User's Guide* for information.

Chapter 5: Troubleshooting

Use the following table to quickly locate information in this chapter.

Topic	Refer to:
SQL Server Replication Setup	"SQL Server Replication Setup" on page 57
Removing In-Memory and Database Information	"Removing In-Memory and Database Information" on page 58

SQL Server Replication Setup

If the replication setup process fails at any step, it is recommended that you do the following:

1. Delete the RN_CORE entry that was removed:

```
DELETE FROM RN_CORE WHERE CoreID = <id>;  
GO (using sqlcmd)
```

2. Update the variables for all masters in the *SQLServerReplicationScriptTool.properties* file.

Note: The NA Core being deleted must be the second NA Core entry and the publisher must be the first entry. Update the properties if needed. Make sure the "mode" property is set to "delete_server".

Removing In-Memory and Database Information

To ensure proper removal of all in-memory and database information, and to avoid replication conflicts due to matching timestamps, NA active tasks are automatically deleted on the NA Core with which they are associated. An active task is any task that does not have the "Succeeded," "Failed," "Duplicate", "Skipped", or "Warning" status.

Keep in mind that deleted tasks could be displayed in task lists for a few moments while the replication process pushes the deletes to other NA Cores in the system. In addition, if the NA Core from which the task originated is not accessible, the delete will fail.

Index

B

- Buttons
 - Renew Config Options 50
 - Reset Priority 50

C

- Conflict table 47

D

- Deleting a Subscriber 27
- Deleting a subscriber 27
- Device Password Rule Reset page 50
- Distributed Conflict List 46
- Distributed Monitor Results page 44
- Distributed systems
 - Failover 51
 - Installation 18
 - Overview 10
 - Recovery 52
 - Terminology 9
 - Troubleshooting 57

E

- Edit Core page 49
- Events
 - RMI error 41
 - Status 46, 47
 - Time synchronization warnings 41

F

- Failover 53
- Files
 - appserver.rcx 53
 - distributed.rcx 26

H

- Horizontal scalability 13

I

- Installation
 - Creating two NA Cores 23
 - Planning 18
 - SQL Server setup 18
 - Unpacking the Replication script 23

L

- List Cores page 48

M

- Monitor results 44

N

- NA
 - Core 9
 - Device Password Rule Reset page 50
 - Distributed Conflict List 46
 - Distributed Monitor Results page 44
 - Documentation 8
 - Edit Core page 49
 - Generated events 40
 - List Cores page 48
 - Mesh 9
 - Renew Configuration Options 50
 - Site Reassignment page 48
 - View Distributed Conflict page 47
- NA scheduler 10

P

- Ports
 - JNDI 21
 - SNMP 21
 - Telnet 21

R

- Recovery
 - Loss of network connectivity 51
 - Replication data 52
- Removing database information 58
- Removing replication 29
- Renew Configuration Options 50

Replication

- Monitoring 10
- Removing 29
- Script tool 25
- Two NA Cores 23
- Unpacking 23

RMI connections 44

Rowguid 46, 47

S

Site Reassignment page 48

SQL Server

- Deleting a subscriber 27
- Multimaster replication 10
- Publisher 23
- Subscriber 23

sqlcmd 25

Subscriber

- Definition 9
- Deleting 27

System administration

- List Cores 48
- Monitor results 44
- NA generated events 40
- NA UI pages 41, 43
- Site reassignment 48
- Uniqueness conflicts 46

T

Terminology 9

Timezone offset 48, 49

Troubleshooting

- Removing information 58
- Replication setup 57

U

Uniqueness conflicts 40, 44

V

View Distributed Conflict page 47