

HP Network Automation Software

For the Windows[®], Linux, and Solaris operating systems

Software Version: 9.20

Multimaster Distributed System on SQL Server Guide

Document Release Date: May 2012
Software Release Date: May 2012



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.

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 2006–2012 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

AMD is a trademark of Advanced Micro Devices, Inc.

Intel and Intel Itanium are trademarks of Intel Corporation in the U.S. and other countries.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

Oracle and Java are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Oracle Technology — Notice of Restricted Rights

Programs delivered subject to the DOD FAR Supplement are 'commercial computer software' and use, duplication, and disclosure of the programs, including documentation, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, programs delivered subject to the Federal Acquisition Regulations are 'restricted computer software' and use, duplication, and disclosure of the programs, including documentation, shall be subject to the restrictions in FAR 52.227-19, Commercial Computer Software-Restricted Rights (June 1987). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

For the full Oracle license text, after product installation see the `$NA_HOME/server/license` directory (or the `%NA_HOME%\server\license` directory on Windows systems) on the NA application server.

Acknowledgements

This product includes software developed by the Apache Software Foundation.
(<http://www.apache.org>)

Parts of this software Copyright © 2003-2008 Enterprise Distributed Technologies Ltd. All Rights Reserved.
(<http://www.enterprisedt.com>)

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>

Contents

1	Getting Started	7
	Terminology	7
	Overview	7
	Architecture	9
	What is Horizontal Scalability?	10
	NA Internal Monitoring: SQL Server Merge Replication	10
	NA Internal Monitoring: Inter-NA Core Communication	11
2	Installation, Setup, and Upgrading	13
	System Requirements	13
	Prepare for Initial Replication	13
	Planning	15
	Protocols, Databases, and Ports	15
	Unpacking the Replication Script Installation Bundle	16
	Creating a Two NA Core SQL Server Replication Environment	16
	Recovering from a Failed Replication Configuration	22
	Removing Replication	22
	System Requirements	22
	Deletion Steps	23
	Upgrade the NA Mesh	24
	Upgrading from NAS 7.0/7.2 to NA 7.50	25
	Upgrading from NA 7.5x to NA 7.60	26
	Upgrading from NA 7.60 to NA 9.00	27
	Upgrading from NA 9.00 to NA 9.10	28
	Upgrading from NA 9.10 to NA 9.20	29
3	System Administration	33
	Getting Started	33
	How NA Removes In-Memory and Database Information	33
	Suspending Replication for Database Maintenance	34
	NA Generated Events	34
	Distributed System - Uniqueness Conflict	34
	Distributed System - Time Synchronization Warning	35
	Distributed System - RMI Error	35
	Distributed System - Stopped Merge Agent Job	36
	Distributed System - Data Synchronization Delay Warning	36
	Using the NA Distributed System Pages	36
	Distributed Monitor Results Page	36
	Distributed Conflict List	37

View Distributed Conflict Page	38
Partition Reassignment Page.....	38
List Cores Page	39
Edit Core Page	39
Device Password Rule Priority Reset Page	40
Renew Configuration Options Page.....	40
4 Failover and Recovery	41
Failover.....	41
Recovery	41
Loss of Network Connectivity	41
Loss of an NA Application Server	42
Loss of a Database Server	42
5 Troubleshooting	45
RMI Errors in the NA Console	45
Removing In-Memory and Database Information	46
Index.....	47

1 Getting Started

This chapter contains the following topics:

- [Terminology](#) on page 7
- [Overview](#) on page 7
- [Architecture](#) on page 9
- [What is Horizontal Scalability?](#) on page 10
- [NA Internal Monitoring: SQL Server Merge Replication](#) on page 10
- [NA Internal Monitoring: Inter-NA Core Communication](#) on page 11

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. An 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. For information about segmenting devices, see the *NA User Guide*.
- **NA Mesh** — Multiple NA Cores connected via multimaster 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.



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

Overview

The NA Distributed System on Microsoft® SQL Server is a multimaster system where the data from each NA Core in an 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.

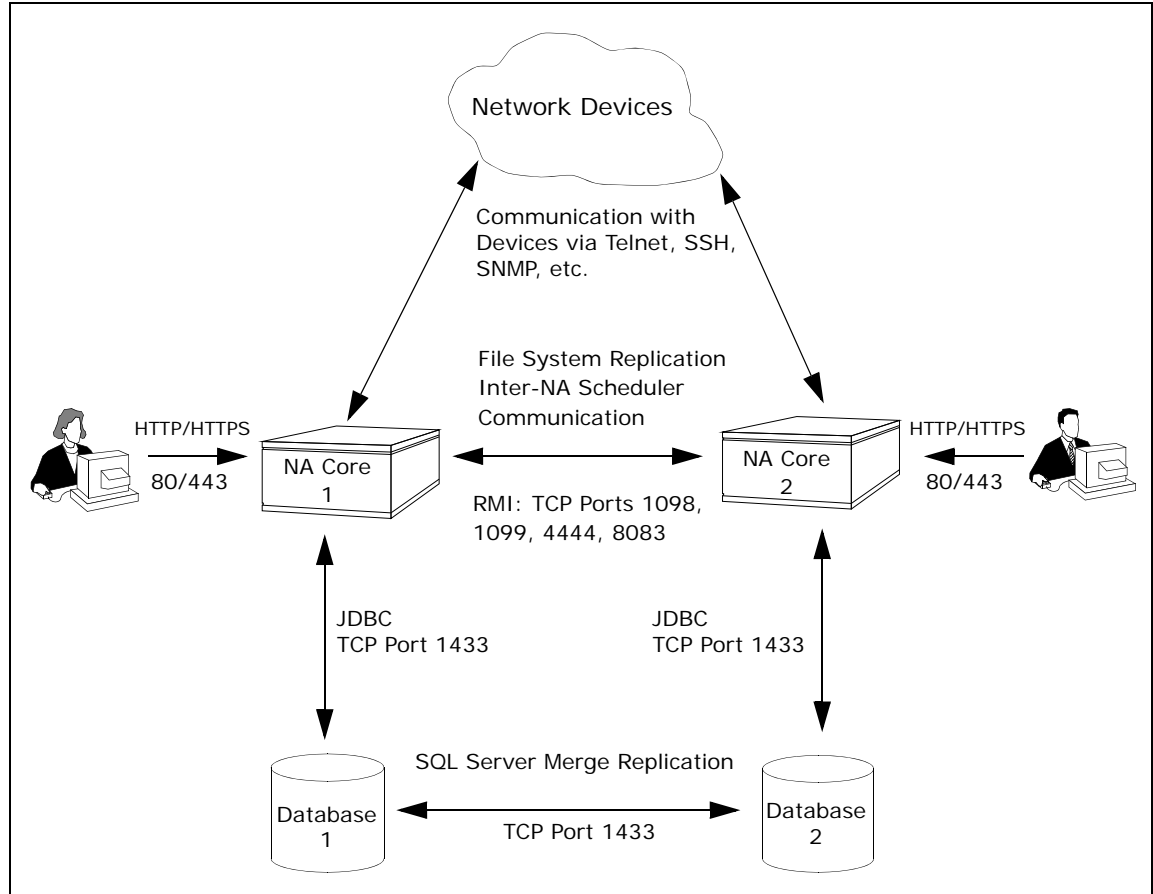
The following comprises the NA Distributed System:

- An 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 console.
 - 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 a 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. For information about segmenting devices, see the *NA User Guide*.

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, see [System Administration](#) on page 33.

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 synchronized.
- 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. For installation and configuration information, see the *NA Horizontal Scalability Guide*.

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. For information about configuring email notification, see the *NA User Guide*.



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

For detailed information on NA monitors, see [Distributed Monitor Results Page](#) on page 36.

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.



For information about configuring email notification, see the *NA User Guide*.

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. For information about configuring email notification, see the *NA User Guide*. For detailed information about NA system events, see [NA Generated Events](#) on page 34.



NA includes an example event rule that can be updated as necessary.

2 Installation, Setup, and Upgrading

The chapter contains the following topics:

- [System Requirements](#) on page 13
- [Prepare for Initial Replication](#) on page 13
- [Creating a Two NA Core SQL Server Replication Environment](#) on page 16
- [Removing Replication](#) on page 22
- [Upgrade the NA Mesh](#) on page 24

For information about troubleshooting an NA Mesh, see [Troubleshooting](#) on page 45.

System Requirements

For information about database version and server hardware requirements for the NA Multimaster Distributed System on SQL Server software, see the *NA Support Matrix*.

Prepare for Initial Replication

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 devices. For example, what network connectivity is required to support failover for device access? The ability of an 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. For more information, see [Architecture](#) on page 9.
- Network configuration planning for bandwidth usage by the NA Mesh. The bandwidth required between an 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 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 the servers. You can also calculate the disk space requirements for the export (and subsequent import) operations by knowing 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.



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 application servers must be offline during the export and import steps. A sufficiently long maintenance 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 existing NA installations into an NA Mesh. You can only create an NA Mesh from a single existing NA server, adding a new NA Core as appropriate. You can also create an NA Mesh from scratch.
- Users should not share logon accounts. Due to the replication system used to share data across NA Cores, two users should never use the same logon 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 affected user’s profile is properly synchronized across the NA Mesh.
- NA currently assumes that all application and database servers in the Mesh share a single time zone.
- Future NA upgrades will take longer and require more downtime 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` table 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 to the second database by using replication. The contents of the `RN_CORE` tables on both databases should match whatever is set up in the `SQLServerReplicationScript.sql` file.
- 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 42 for information on restoring databases.

Planning

This section includes information for initial replication setup on a Publisher and a Subscriber. NA will be offline during initial replication. It is recommended that you carefully read the following sections so as to streamline the replication process. A new database instance on the Subscriber is manually created when running initial replication setup.

When initially setting up replication, complete the following table before starting. Be sure to save this information. You will need it each time you upgrade in the future.

Table 1 Distributed System Server Information

Information to Collect	NA Core 1	NA Core 2
Time Zone Offset from UTC		Same for all servers in the NA mesh
Application Server Data	NA Application Server 1	NA Application Server 2
NA App Server FQDN		
NA App Server RMI Port		
Database Server Data	Database Server 1 (Publisher)	Database Server 2 (Subscriber)
Database Server FQDN		
SQL Server Listening Port		
Database Name		
Database User		Same as for publisher
Database Password		Same as for publisher

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. For more information, see [Architecture](#) on page 9.

Protocol/Database/Port	From/To
NA server (running the Management 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.

Protocol/Database/Port	From/To
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 a 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.
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.

Unpacking the Replication Script Installation Bundle

The setup files for the NA Multimaster 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 application server. The setup files include a Java program that you run to generate SQL scripts for replication setup. This program uses the Java runtime that is available on the NA application server.

Creating a Two NA Core SQL Server Replication Environment

NA can already be installed on one application server and database server (a single NA core). Alternatively, you can complete the first NA core installation as part of setting up initial replication.

This procedure uses the following identifiers:

- NA application servers: `App1` and `App2`
- SQL Server database names:
 - `NA1` on database server 1, the publisher
 - `NA2` on database server 2, the subscriber
- NA database user for both databases: `nauser`
- Password for the NA database user: `napass`

To use a different identifier, substitute the actual value for each instance of the example value within the procedure.

To set up initial replication between two NA cores, complete the following tasks:

- [Task 1: Verify Prerequisites](#)
- [Task 2: \(New Installations Only\) Create the NA Schema on Database Server 1](#)
- [Task 3: Prepare for Replication](#)
- [Task 4: Configure Replication](#)

Task 1: [Verify Prerequisites](#)

Prepare two database servers and two application servers. For best performance, each application server-database server pair should be co-located. Additionally, the application server-application server link and the database server-database server link should each be a high-speed connection.

Verify that the following prerequisites have been met:

- The two database servers are running a supported version of SQL Server. Both database servers must use the same version of the database software.
 - Identify the user name (for example, `sa`) and password of a SQL Server user account with the `sysadmin` role on both database servers.
 - Ensure `sqlcmd` access to both database servers.

You will need to supply credentials for a database user account that is a member of the `sysadmin` role when running the scripts.
- The SQL Server agent service is running on database server 1, the publisher.
 - Identify the user name and password of a Windows user account under which SQL Server agents can run.
- The two application servers that will run NA have working hostnames. Note the following:
 - Each application server should have a high-speed connection to one of the database servers.
 - For application server hardware and operating system requirements, see the *NA Support Matrix*.
- The host names of both database servers and both application servers are in the `hosts` file on each application server. This file is located as follows:
 - *Windows*: `<Drive>:\Windows\System32\drivers\etc\hosts`
 - *UNIX*[®]: `/etc/hosts`
- All database servers and application servers are set to use the same time and time zone.

Task 2: [\(New Installations Only\) Create the NA Schema on Database Server 1](#)

If NA has not yet been installed, create the NA schema on database server 1 by following these steps:

- 1 Install NA on the first application server (`App1`).

During installation, supply the following database credentials:

 - For the database IP address or hostname, enter the fully-qualified domain name of database server 1.


- For the database port, enter the port on database server 1 that NA will use to communicate with the database.
 - For the database name, enter the name of the NA database on the database server 1 (for example, NA1).
 - For the database user name and password, enter a name and password for accessing the NA database, for example, nauser and napass.
- 2 After NA installation is complete, log on to NA (as user nauser) to ensure that it works as expected.
If NA does not work correctly, follow the steps in [Recovering from a Failed Replication Configuration](#) on page 22.
 - 3 Stop the NA management engine on App1:

Task 3: Prepare for Replication

To configure replication between the publisher and the subscriber databases, follow these steps:

- 1 Stop the NA management engine on App1:
 - **Windows:** Open the **Services** control panel. In the list of services, right-click **TrueControl ManagementEngine**, and then click **Stop**.
 - **UNIX:** Run the following command:


```
/etc/init.d/truecontrol stop
```
- 2 If the NA database contains working data (is not a new installation), back up the database.
Alternatively, if you are recovering from a failed replication configuration and had working data before starting configuration, restore the database from backup at this time.
- 3 On database server 1, create a shared directory (for example, C:\NA_Share).
 - Set the sharing permissions so that both database servers have read-write access to this directory.
 - Also set the sharing permissions so that the administrative user on App1 (or other system) who will run the Java script in [step 2 of Task 4](#) has read-write access to this directory.
 - Ensure that the directory has at least 50MB available disk space. If the NA database contains working data, this requirement might increase.
- 4 On database server 2, create an empty database. Note the following:
 - Use a default database instance. (NA replication does not support a named instance.)
 - Use both Windows and SQL Server authentication.
 - Enable full-text indexing.
 - The name of this database must be *different* from the name of the first NA database, for example, NA2.
 - In all other aspects, such as privileges and owner, this database configuration must be the same as the configuration of the first database.
- 5 On App1, unpack the SQLServerReplicationScriptToolBundle.zip file to a known location with low security, for example, C:\tmp or /tmp.

- 6 From the known location of [step 5](#), customize the `SQLServerReplicationScriptTool.properties` file for your environment:
 - a In a text editor such as WordPad or vi, open the `SQLServerReplicationScriptTool.properties` file.
 - b Search for and replace the text `README` with the appropriate value for that variable.
 - Use short hostnames to identify the database servers.
-  Verify DNS resolution of the database server hostnames to IP addresses.
- Use fully-qualified hostnames to identify the application servers and the replication directory.
 - Set `windows.username.0` to either `<machineName>\\<administrator>` or `<domain>\\<domainUser>`. If you use the first format, ensure that the administrator can access the shared directory created in [step 3](#) of this task.
 - For example:


```
db.server.0=db1
db.name.0=NA1
db.username.0=nauser
db.password.0=napass
windows.username.0=db1\\Administrator
windows.password.0=adminpass
db.tableowner.0=dbo
db.port.0=1433
app.server.0=App1.com
app.rmiport.0=1099
app.corename.0=Core1

replication.data.dir=\\\\db1.example.hp.com\\NA_Share

db.server.1=db2
db.name.1=NA2
db.port.1=1433
app.server.1=App2.example.com
app.rmiport.1=1099
app.corename.1=Core2
```
- c Set `timezone_offset` to the value that matches the time zone setting for all servers in the distributed system.
 - d Verify that `mode` is set to `initial`.

Task 4: Configure Replication

- 1 From the known location of [step 5](#) of [Task 3](#), run `java -version` to verify that Java is version 1.6 or later.
- 2 From the known location of [step 5](#) of [Task 3](#), run the `ReplicationScriptTool`:


```
java -classpath . SQLServerReplicationScriptTool
```

This command creates two scripts for starting replication. By default, the scripts are named `SQLServerPreSnapshotScript.sql` and `SQLServerReplicationTool.sql`.



The following error message can mean that Java does not have permission to run the replication tool:

```
Could not find the main class: ReplicationScriptTool.
```

Move all files from the replication bundle to a directory with lower security, and then run the java command from that directory.

- 3 Copy the `SQLServerPreSnapshotScript.sql` script to the shared directory created in [step 3 of Task 3](#).
- 4 Copy the `SQLServerReplicationTool.sql` script to a known location on database server 1, for example, `C:\tmp`.
 - Set the sharing permissions so that the SQL Server sysadmin user account has read-write access to this directory.
 - Verify that `sqlcmd` is accessible from this directory.
- 5 From the known location on database server 1 ([step 5 of this task](#)), run the `SQLServerReplicationTool.sql` script by using the `sqlcmd` command:

```
sqlcmd -S <Publisher> -U <User> -P <Password>
-i SQLServerReplicationScript.sql -o log.txt
```

For `<Publisher>`, use the short hostname of database server 1, for example `db1`.

For `<User>` and `<Password>`, use a SQL Server user account with the sysadmin role, for example, `sa`.

This command creates a directory named `UNC` in the shared directory created in [step 3 of Task 3](#). It also populates the `UNC` directory with a snapshot of the NA database on database server 1.

- 6 Verify that the new `UNC` directory contains the NA schema.

If the `UNC` directory does not contain the NA schema, follow the steps in [Recovering from a Failed Replication Configuration](#) on page 22.
- 7 Finish configuring replication by entering the following Transact-SQL queries on the appropriate database servers:



If you copy the select statement from this PDF and paste it to the `sqlcmd` line, re-type the single quotation mark characters (`'`) on the command line.

- Publisher:

```
sqlcmd -S <Publisher> -U <User> -P <Password>
use <DatabaseName>
Select count(*) from Information_Schema.Tables where Table_Type =
'Base Table' and Table_Name like 'RN_%'
go
```

For `<Publisher>`, use the short hostname of database server 1, for example `db1`.

For `<User>` and `<Password>`, use a SQL Server user account with the sysadmin role, for example, `sa`.

`<DatabaseName>` is the name of the NA database on database server 1, for example `NA1`.

- **Subscriber:**

```
sqlcmd -S <Subscriber> -U <User> -P <Password>
use <DatabaseName>
Select count(*) from Information_Schema.Tables where Table_Type =
'Base Table' and Table_Name like 'RN_%'
go
```

For <Subscriber>, use the short hostname of database server 2, for example db2.

For <User> and <Password>, use a SQL Server user account with the sysadmin role, for example, sa.

<DatabaseName> is the name of the NA database on database server 2, for example NA2.

- 8 Verify that the previous step completed correctly.

The query results should be the same on both database servers. For example:

```
Changed database context to 'NA'.
-----
              79
(1 rows affected)
```

If the query results for both database servers are not identical, follow the steps in [Recovering from a Failed Replication Configuration](#) on page 22.

- 9 Verify that the replication is working correctly. Make sure:

- The log file generated by running the replication SQL script contains no errors.
- In Microsoft SQL Management Studio, examine the NA databases:

Under Databases, look for the database name.

- Verify that the table count is the same for each database.
- Verify that the RN_CORE table on each database contains the appropriate list of servers in the NA mesh.
- Verify that the replication of data is working correctly by adding a comment (for example, Test from db1 to db2) to the CoreID=1 row of the RN_CORE table for the publisher database. After a minute, this comment should be visible for the same table row of the subscriber database.

If replication is not working correctly, follow the steps in [Recovering from a Failed Replication Configuration](#) on page 22.

- 10 Install NA on the second application server (App2).


During installation, connect NA to the database on Master 2. When prompted for the database credentials, choose: use existing database.

- 11 Stop the NA management engine on App2:

- **Windows:** Open the **Services** control panel. In the list of services, right-click **TrueControl ManagementEngine**, and then click **Stop**.
- **UNIX:** Run the following command:

```
/etc/init.d/truecontrol stop
```

- 12 Copy the `distributed.rcx` file from the known location of [step 5](#) to the following directory on both application servers, App1 and App2:
 - *Windows*: `<NA_HOME>\jre`
 - *UNIX*: `<NA_HOME>/jre`
- 13 Start the NA management engine on both application servers, App1 and App2:
 - *Windows*: Open the **Services** control panel. In the list of services, right-click **TrueControl ManagementEngine**, and then click **Start**.
 - *UNIX*: Run the following command:


```
/etc/init.d/truecontrol restart
```
- 14 Connect to the NA console on both application servers, App1 and App2, (as user `nauser`), and ensure that there are no issues.
 -  An RMI error message at this point can mean that NA is unable to identify the localhost. For more information, see [RMI Errors in the NA Console](#) on page 45.
- 15 Verify the NA installation:
 - a Edit an object on one NA server (for example, a Comments field for a device).
 - b Wait a minute and then verify that the updated comment exists on the second server.
 - c Check the status of the Oracle Distributed Monitor in the NA console to ensure that no problems are being reported. (For information, see [Distributed Monitor Results Page](#) on page 36.) It could take up to five minutes for this monitor to initially run.
- 16 *Optional*. Delete the shared directory created in [Task 3: Prepare for Replication](#) on page 18 (for example, `C:\NA_data`).

Recovering from a Failed Replication Configuration

If you see errors during replication setup, follow these steps:

- 1 Uninstall NA from both application servers.
- 2 Delete the NA databases from both database servers.
- 3 Redo the replication configuration process starting with [Task 2: \(New Installations Only\) Create the NA Schema on Database Server 1](#) on page 17.

Removing Replication

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

System Requirements

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

- `sqlcmd` access to both database servers in the mesh.
- The `SQLServerReplicationScriptTool` application installed on a Java-capable system.

Deletion Steps

To delete the Subscriber from an NA Mesh:

- 1 Collect the following information:
 - Login name and password of a SQL Server login that is a member of sysadmin on both database servers in the mesh.
 - 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 the publisher. Alternatively, remove those partitions.
- 4 Stop the NA management engine on the subscriber application server:
 - **Windows:** Open the **Services** control panel. In the list of services, right-click **TrueControl ManagementEngine**, and then click **Stop**.
 - **UNIX:** Run the following command:


```
/etc/init.d/truecontrol stop
```
- 5 Delete the RN_CORE entry for the subscriber using the following commands using sqlcmd:


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



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

```
SELECT * FROM RN_CORE;
```

- 6 Update the variables for NA Core 2 in the `SQLServerReplicationScriptTool.properties` file.
 - a Make sure that the variables for NA Core 1 are correct.
 - b Set mode to `delete_server`.

These properties are described in detail in the file.
- 7 From the replication bundle directory, run the `ReplicationScriptTool`:


```
java -classpath . SQLServerReplicationScriptTool
```
- 8 Copy the `SQLServerReplicationTool.sql` script to a known location on the publisher database server, for example, `C:\tmp`.
 - Set the sharing permissions so that the SQL Server sysadmin user account has read-write access to this directory.
 - Verify that `sqlcmd` is accessible from this directory.
- 9 From the known location on database server 1 (step 8 of this task), run the `SQLServerReplicationTool.sql` script by using the `sqlcmd` command:

```
sqlcmd -S <Publisher> -U <User> -P <Password>
-i SQLServerReplicationScript.sql -o log.txt
```

For `<Publisher>`, use the short hostname of database server 1, for example `db1`.

For `<User>` and `<Password>`, use a SQL Server user account with the sysadmin role, for example, sa.

- 10 Remove the `distributed.rcx` file from each NA application server.
- 11 Start the NA management engine on NA Core 1:
 - *Windows*: Open the **Services** control panel. In the list of services, right-click **TrueControl ManagementEngine**, and then click **Start**.
 - *UNIX*: Run the following command:

```
/etc/init.d/truecontrol restart
```

- 12 On the publisher, create a script with the following contents:

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

Upgrade the NA Mesh

During the NA upgrade process, the NA Cores must be completely offline. You will be upgrading the:

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

Note the following:

- The NA mesh upgrade path is as follows:
 - (NAS 7.0 or NAS 7.2) to NA 7.50 to NA 7.60 to NA 9.00 to NA 9.10 to NA 9.20
 Enter this path at your current version of NA.
- At each step on the NA mesh upgrade path, both NA cores in the NA mesh before moving to the next step in the upgrade path.
- For instructions for upgrading from NAS 6.2 to NAS 7.0, see the *NAS 7.0 Multimaster Distributed System on SQL Server User's Guide*.


For upgrade instructions, see the following topics:

- [Upgrading from NAS 7.0/7.2 to NA 7.50](#) on page 25
- [Upgrading from NA 7.5x to NA 7.60](#) on page 26
- [Upgrading from NA 7.60 to NA 9.00](#) on page 27
- [Upgrading from NA 9.00 to NA 9.10](#) on page 28
- [Upgrading from NA 9.10 to NA 9.20](#) on page 29

Upgrading from NAS 7.0/7.2 to NA 7.50

To upgrade from NAS 7.0/7.2 to NA 7.50, be sure that:

- You have `sqlcmd` access to NA Core 1 and NA Core 2.
 - The *ReplicationScriptTool* bundle files for NA 7.50 are installed on a Java-capable system (on one of the NA systems).
- 1 Shut down both NA Management Engines in the NA Mesh.
 - 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.
 - 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 the publisher and subscriber in the `SQLServerReplicationScriptTool.properties` file. Ensure that you modify all “REPLACEME” text entries in the `SQLServerReplicationScriptTool.properties` file with the appropriate values for your environment.

 All application and database servers should be listed in the `SQLServerReplicationScriptTool.properties` file. Make sure the `mode` property is set to the appropriate value for the upgrade you want to perform. For example, if you are upgrading from NAS 7.0 to NA 7.50, this value should be set to “`upgrade_from_7_0`”. If you are upgrading from NAS 7.2 to NA 7.50, this value should be set to “`upgrade_from_7_2`”.

- 4 Run the `ReplicationScriptTool`. Enter:

```
java -classpath . 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.
 - In Microsoft SQL Management Studio, examine the NA databases:
 - Under Databases, look for the database name.
 - Verify that the table count is the same for each database.
 - Verify that the `RN_CORE` table on each database contains the appropriate list of servers in the NA mesh.
 - Verify that the replication of data is working correctly by adding a comment (for example, `Test` from `db1` to `db2`) to the `CoreID=1` row of the `RN_CORE` table for the publisher database. After a minute, this comment should be visible for the same table row of the subscriber database.

- 7 Upgrade each NA server using the NA 7.50 Service Pack Installer. For information about running the NA 7.50 Service Pack Installer, see the *HP Network Automation 7.50 Upgrade and Installation Guide*.
- 8 Restart the NA servers in the NA Mesh.

Upgrading from NA 7.5x to NA 7.60

To upgrade from 7.5x to NA 7.60, be sure that:

- You have `sqlcmd` access to NA Core 1 and NA Core 2.
 - The *ReplicationScriptTool* bundle files for NA 7.60 are installed on a Java-capable system (on one of the NA systems).
- 1 Shut down both NA Management Engines in the NA Mesh.
 - 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.
 - 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 the publisher and subscriber in the `SQLServerReplicationScriptTool.properties` file. Ensure that you modify all “REPLACEME” text entries in the `SQLServerReplicationScriptTool.properties` file with the appropriate values for your environment.
 -  All servers should be listed in the `SQLServerReplicationScriptTool.properties` file. Set the mode property to “`upgrade_from_7_5`”.
 - 4 Run the `ReplicationScriptTool`. Note that you must run this on a system that has Java 1.6.x installed. Enter:


```
java -classpath . 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.
 - In Microsoft SQL Management Studio, examine the NA databases:
 - Under Databases, look for the database name.
 - Verify that the table count is the same for each database.
 - Verify that the `RN_CORE` table on each database contains the appropriate list of servers in the NA mesh.

- Verify that the replication of data is working correctly by adding a comment (for example, `Test` from `db1` to `db2`) to the `CoreID=1` row of the `RN_CORE` table for the publisher database. After a minute, this comment should be visible for the same table row of the subscriber database.
- 7 Upgrade each NA server using the NA 7.60 Service Pack Installer. For information about running the NA 7.60 Service Pack Installer, see 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 9.00

To upgrade from 7.60 to NA 9.00, be sure that:

- You have `sqlcmd` access to NA Core 1 and NA Core 2.
 - The *ReplicationScriptTool* bundle files for NA 9.00 are installed on a Java-capable system (on one of the NA systems).
- 1 Shut down both NA Management Engines in the NA Mesh.
 - 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.
 - 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 the publisher and subscriber in the `SQLServerReplicationScriptTool.properties` file. Ensure that you modify all “REPLACEME” text entries in the `SQLServerReplicationScriptTool.properties` file with the appropriate values for your environment.
 -  All servers should be listed in the `SQLServerReplicationScriptTool.properties` file. Set the `mode` property to “`upgrade_from_7_6`”.
 - 4 Run the `ReplicationScriptTool`. Note that you must run this on a system that has Java 1.6.x installed. Enter:


```
java -classpath . 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.
 - In Microsoft SQL Management Studio, examine the NA databases:
 - Under Databases, look for the database name.
 - Verify that the table count is the same for each database.

- Verify that the RN_CORE table on each database contains the appropriate list of servers in the NA mesh.
 - Verify that the replication of data is working correctly by adding a comment (for example, Test from db1 to db2) to the CoreID=1 row of the RN_CORE table for the publisher database. After a minute, this comment should be visible for the same table row of the subscriber database.
- 7 Upgrade each NA server using the NA 9.00 Service Pack Installer. For information about running the NA 9.00 Service Pack Installer, see the *HP Network Automation 9.00 Upgrade and Installation Guide*.
 - 8 Restart the NA servers in the NA Mesh.

Upgrading from NA 9.00 to NA 9.10

To upgrade from NA 9.00 to NA 9.10, be sure that:

- You have sqlcmd access to NA Core 1 and NA Core 2.
 - The *ReplicationScriptTool* bundle files for NA 9.10 are installed on a Java-capable system (on one of the NA systems).
- 1 Shut down both NA Management Engines in the NA Mesh.
 - 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.
 - 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 the publisher and subscriber in the `SQLServerReplicationScriptTool.properties` file. Ensure that you modify all “REPLACEME” text entries in the `SQLServerReplicationScriptTool.properties` file with the appropriate values for your environment.
 -  All servers should be listed in the `SQLServerReplicationScriptTool.properties` file. Set the mode property to “upgrade_from_9_0”.
 - 4 Run the `ReplicationScriptTool`. Note that you must run this on a system that has Java 1.6.x installed. Enter:


```
java -classpath . 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.
 - In Microsoft SQL Management Studio, examine the NA databases:
 - Under Databases, look for the database name.
 - Verify that the table count is the same for each database.
 - Verify that the RN_CORE table on each database contains the appropriate list of servers in the NA mesh.
 - Verify that the replication of data is working correctly by adding a comment (for example, `Test from db1 to db2`) to the `CoreID=1` row of the RN_CORE table for the publisher database. After a minute, this comment should be visible for the same table row of the subscriber database.
- 7 Upgrade each NA server using the NA 9.10 Service Pack Installer. For information about running the NA 9.10 Service Pack Installer, see the *HP Network Automation 9.10 Upgrade and Installation Guide*.
- 8 Restart the NA servers in the NA Mesh.

Upgrading from NA 9.10 to NA 9.20

To upgrade from NA 9.10 to NA 9.20, be sure that:

- You have `sqlcmd` access to NA Core 1 and NA Core 2.
 - The *ReplicationScriptTool* bundle files for NA 9.20 are installed on a Java-capable system (on one of the NA systems).
- 1 Shut down both NA Management Engines in the NA Mesh.
 - 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.
 - 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 the publisher and subscriber in the `SQLServerReplicationScriptTool.properties` file. Ensure that you modify all “REPLACEME” text entries in the `SQLServerReplicationScriptTool.properties` file with the appropriate values for your environment.
 -  All servers should be listed in the `SQLServerReplicationScriptTool.properties` file. Set the `mode` property to “`upgrade_from_9_1`”.
 - 4 Run the `ReplicationScriptTool`. Note that you must run this on a system that has Java 1.6.x installed. Enter:


```
java -classpath . 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.
 - In Microsoft SQL Management Studio, examine the NA databases:
 - Under Databases, look for the database name.
 - Verify that the table count is the same for each database.
 - Verify that the `RN_CORE` table on each database contains the appropriate list of servers in the NA mesh.
 - Verify that the replication of data is working correctly by adding a comment (for example, `Test` from `db1` to `db2`) to the `CoreID=1` row of the `RN_CORE` table for the publisher database. After a minute, this comment should be visible for the same table row of the subscriber database.
- 7 Upgrade each NA application server as described in the “Upgrading to NA 9.20 from a Different System” chapter or the “Upgrading to NA 9.20 on the Same System” chapter of the *NA Installation and Upgrade Guide*, as appropriate.

Run the upgrade procedures in parallel. That is, complete step 1 on each NA application server before initiating step 2 on any NA application server, and so forth. Follow the referenced procedure through the step to run the NA 9.20 Service Pack Installer.

Note the following:

- To prevent database access, on each NA application server, after the NA 9.20 Service Pack Installer runs, stop all NA services.
 - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Stop**:
 - TrueControl ManagementEngine**
 - TrueControl FTP Server**
 - TrueControl SWIM Server**
 - TrueControl Syslog Server**
 - TrueControl TFTP Server**
 - *UNIX*: Run the following command:


```
/etc/init.d/truecontrol stop
```
- The NA database upgrade has already been handled by running the replication script in [step 4](#).
- Do *not* upgrade SQL Server while upgrading NA.
- If you are moving NA to different application servers, copy the NA-specific files (the NA directory and, on UNIX, the startup files) from one existing NA application server to only one new NA application server.

For example, consider a two-core Distributed System environment. The existing NA application servers are Server X1 and Server X2. The new NA application servers are Server Y1 and Server Y2. Do the following:

- Copy the NA directory from Server X1 to Server Y1.

Retain the permissions on the files.

On UNIX systems, also copy the `/etc/truecontrol` and `/etc/init.d/truecontrol` files from Server X1 to Server Y1.

- Copy the NA directory from Server X2 to Server Y2.

Retain the permissions on the files.

On UNIX systems, also copy the `/etc/truecontrol` and `/etc/init.d/truecontrol` files from Server X2 to Server Y2.

- 8 If you moved NA to different application servers in [step 7](#), in Microsoft SQL Management Studio update the `RN_CORE` table on one database with the fully-qualified domain names of the new NA application servers.
- 9 Restart the NA servers in the NA Mesh.
- 10 Complete the upgrade procedure from the step to run the Data Pruning task.

3 System Administration

This chapter contains the following topics:

- [Getting Started](#) on page 33
- [How NA Removes In-Memory and Database Information](#) on page 33
- [Suspending Replication for Database Maintenance](#) on page 34
- [NA Generated Events](#) on page 34
- [Using the NA Distributed System Pages](#) on page 36

Getting Started

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

How NA Removes In-Memory and Database Information

There could be a delay in removing or cancelling a task because NA must communicate with the appropriate owning NA Core to perform this operation. 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.

Suspending Replication for Database Maintenance

To suspend replication because you are bringing down one or both database servers for maintenance, complete the following procedure:

- 1 Shut down the NA management engine on both NA application servers:
 - *Windows*: Open the **Services** control panel. In the list of services, right-click **TrueControl ManagementEngine**, and then click **Stop**.
 - *UNIX*: Run the following command:


```
/etc/init.d/truecontrol stop
```
- 2 Disable database replication by using the Disable Publishing and Distribution Wizard in SQL Server Management Studio. For information, see the Microsoft TechNet Library:
 - SQL Server 2005: *How to Disable SQL Server Database Replication* at <http://technet.microsoft.com/en-us/library/bb693954.aspx>
 - SQL Server 2008: *How to: Disable Publishing and Distribution (Replication Transact-SQL Programming)* at <http://technet.microsoft.com/en-us/library/ms147921.aspx>

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. You should examine the default “Distributed System” event rule to ensure all of the events are included in the event rule and that the event rule is configured to send the email notification to the appropriate administrator.



For information about configuring email notification, see the *NA User Guide*.

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. See [Device Password Rule Priority Reset Page](#) on page 40.

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 SQL Server 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. Examine the databases, merge agents, and network connections between the databases 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.

Distributed Monitor Results Page

The Distributed Monitor Results page displays the overall health of the Distributed System. By default, the Distributed monitor runs every five minutes.

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. For information, see [Edit Core Page](#) on page 39.

➤ 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 an 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>
```

➤ 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.

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.

Field	Description
Actions	<p>You can select the following options:</p> <ul style="list-style-type: none"> • Detail — Opens the View Distributed Conflict page, where you can view details on an individual uniqueness constraint. See View Distributed Conflict Page on page 38. • 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.

Partition Reassignment Page

The Partition Reassignment page enables the partition-to-NA Core mapping to be modified. This is useful for failover of partitions from one NA Core to another and for restoring partitions back to their original NA Core.

To open the Partition Reassignment page, on the menu bar under Admin select Distributed and click Partition Reassignment. The Partition Reassignment page 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. See Edit Core Page on page 39.

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:

- 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 Oracle 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 about segmenting devices, see the *NA User 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 an 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.

4 Failover and Recovery

This chapter contains the following topics:

- [Failover](#) on page 41
- [Recovery](#) on page 41

Failover

When the network has been configured to failover, if an NA Core fails, users can continue to access all data in the system using a different NA Core. All partitions that were originally managed by the failed NA Core can be pointed to a new NA Core using the [Partition Reassignment](#) page. For information, see [Partition Reassignment Page](#) on page 38. 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 an NA application 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 application server or with the SQL Server database server. Recovery consists of the following steps:

- 1 Resolve the network issues.
- 2 Reset partitions that had been reassigned back to their original NA Core. This can be accomplished in NA. See [Partition Reassignment Page](#) on page 38.
- 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. For information, see the *NA User Guide*.
- 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. See [Renew Configuration Options Page](#) on page 40.

- 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 partition data and avoid exception errors due to obsolete data.

After the network issues are resolved, the system should recover as replication synchronizes data between the databases.

Loss of an NA Application Server

In cases where the NA application server suffers a failure that requires re-installation of the NA application, 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 partitions that had been re-assigned back to their original NA Core. This can be accomplished in NA. See [Partition Reassignment Page](#) on page 38.
- 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. For information, see the *NA User Guide*.
- 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. See [Renew Configuration Options Page](#) on page 40.
- 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 application server is still running but cannot access the database. The database will need to be rebuilt and replication set up again on the database. Recovery consists of the following steps:

- 1 Pause or delete any tasks that appear to be waiting, pending, or not running because they are associated with the lost database server. You can perform this action on the other NA application server in the NA Mesh.
- 2 If the lost database server is the Subscriber, replication needs to be removed from the NA Mesh. For information, see [Removing Replication](#) on page 22. Keep in mind that Replication can be reconfigured after removal. For information, see [Creating a Two NA Core SQL Server Replication Environment](#) on page 16.
- 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. See [Deletion Steps](#) on page 23 (step 5).
 - e Install replication. For information, see [Creating a Two NA Core SQL Server Replication Environment](#) on page 16.
- 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 partitions that had been reassigned back to their original NA Core. This can be accomplished in NA. See [Partition Reassignment Page](#) on page 38.
 - 6 In NA, edit the NA Core that was recovered to make sure all information is correct for the new setup. For information, see the *NA User Guide*.

5 Troubleshooting

This chapter contains the following topic:

- [RMI Errors in the NA Console](#) on page 45
- [Removing In-Memory and Database Information](#) on page 46

RMI Errors in the NA Console

An RMI error that prevents NA console logon or inhibits use of the NA console can mean that NA is unable to identify the localhost.

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.

To correct the problem, update the localhost section of the `hosts` file on each NA application server as follows:



This solution is for static IP environments only.

- 1 In a text editor such as WordPad or `vi`, open the following file:
 - *Windows:* `<Drive>:\Windows\System32\drivers\etc\hosts`
 - *UNIX:* `/etc/hosts`
- 2 Set the localhost line to read:


```
127.0.0.1 localhost
```

- 3 For each NA application server, add a line in the following format:

```
<xx.xx.xx.xx> <NA.example.com> <NA>
```

 - Replace `<xx.xx.xx.xx>` with the IP address of the NA application server.
 - Replace `<NA.example.com>` with the fully-qualified domain name of the NA application server.
 - Replace `<NA>` with the short hostname of the NA application server.
- 4 Repeat [step 3](#) until the `hosts` file includes all NA application servers in the distributed system environment.
- 5 To use the updated hosts information, restart the NA application 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 deletions to other NA Core in the system. In addition, if the NA Core from which the task originated is not accessible, the task deletion fails.

Index

B

- Buttons
 - Renew Config Options 40
 - Reset Priority 40

C

- Conflict table 38

D

- Deleting a Subscriber 22
- Deleting a subscriber 22
- Device Password Rule Reset page 40
- Distributed Conflict List 37
- Distributed Monitor Results page 36
- Distributed systems
 - Failover 41
 - Installation 13
 - Overview 7
 - Recovery 41
 - Terminology 7
 - Troubleshooting 45

E

- Edit Core page 39
- Events
 - RMI error 35
 - Status 37, 38
 - Time synchronization warnings 35

F

- Failover
 - Loss of a database server 42
 - Loss of a NA server 42
- Files
 - distributed.rcx 22
 - ReplicationScriptTool 19, 42

H

- Horizontal scalability 10

I

- Installation
 - Creating two NA Cores 16
 - Planning 13
 - SQL Server setup 13
 - Unpacking the Replication script 16

L

- List Cores page 39

M

- Monitor results 36

N

- NA
 - Core 7
 - Device Password Rule Reset page 40
 - Distributed Conflict List 37
 - Distributed Monitor Results page 36
 - Edit Core page 39
 - List Cores page 39
 - Mesh 7
 - Renew Configuration Options 40
 - Site Reassignment page 38
 - View Distributed Conflict page 38

- NA scheduler 8

P

- Partition Reassignment page 38

Ports

- JNDI 16
- SNMP 15
- Telnet 15

R

- Recovery
 - Loss of network connectivity 41
 - Replication data 41
- Removing database information 33, 46
- Removing replication 22
- Renew Configuration Options 40

Replication

- Monitoring 8
- Script tool 20
- Two NA Cores 16
- Unpacking 16

RMI connections 36

Rowguid 37, 38

S**SQL Server**

- Deleting a subscriber 22
- Multi-master replication 8
- Publisher 16
- Server setup 13
- Subscriber 16

sqlcmd 20

Subscriber

- Definition 7
- Deleting 22

System administration

- List Cores 39
- Monitor results 36
- NA generated events 34
- NA UI pages 35, 36
- Partition reassignment 38
- Uniqueness conflicts 37

T

Terminology 7

Timezone offset 39

Troubleshooting

- Removing information 46

U

Uniqueness conflicts 34, 37

V

View Distributed Conflict page 38

We appreciate your feedback!

If an email client is configured on this system, by default an email window opens when you click *here*.

If no email client is available, copy the information below to a new message in a web mail client, and then send this message to **ovdoc-nsm@hp.com**.

Product name and version: NA 9.20

Document title: *NA Multimaster Distributed System on SQL Server Guide, May 2012*

Feedback: