HP Network Automation

Software Version 9.10

Horizontal Scalability User's Guide

Document Release Date: March 2011
Software Release Date: March 2011

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. 110410

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

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

Introduction

The *HP Network Automation 9.10 Horizontal Scalability User's Guide* describes the Network Automation System (NA) Horizontal Scalability functionality. It includes information on:

- What is NA Horizontal Scalability
- Configuring NA Horizontal Scalability on a two NA Core configuration
- Adding additional NA Cores to a NA Horizontal Scalability configuration

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 Sites.
- **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 9.10 User's Guide* for information on segmenting devices.
- **NA Mesh** — Multiple NA Cores connected via replication.

Supported Databases and System Requirements

The NA Horizontal Scalability solution is supported on the following databases:

- SQL Server 2005, Service Pack 2 (Standard Edition or Enterprise Edition)
- SQL Server 2008 (Standard Edition or Enterprise Edition)
- Oracle Server 10g (10.2.0.4) Standard Edition
- Oracle Server 11gR1 Standard or Enterprise Edition

The maximum number of NA Cores supported is five.

Horizontal Scalability system requirements are equivalent to those for standard NA installs, as are the required database settings. However, the Horizontal Scalability system is not designed to be used over WANs where NA Cores and the database are separated by large geographic distances. Refer to the *HP Network Automation 9.10 Support Matrix* for information on standard database server hardware requirements.

What is NA Horizontal Scalability?

NA 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.

The following figure shows a typical two-NA Core Horizontal Scalability installation.

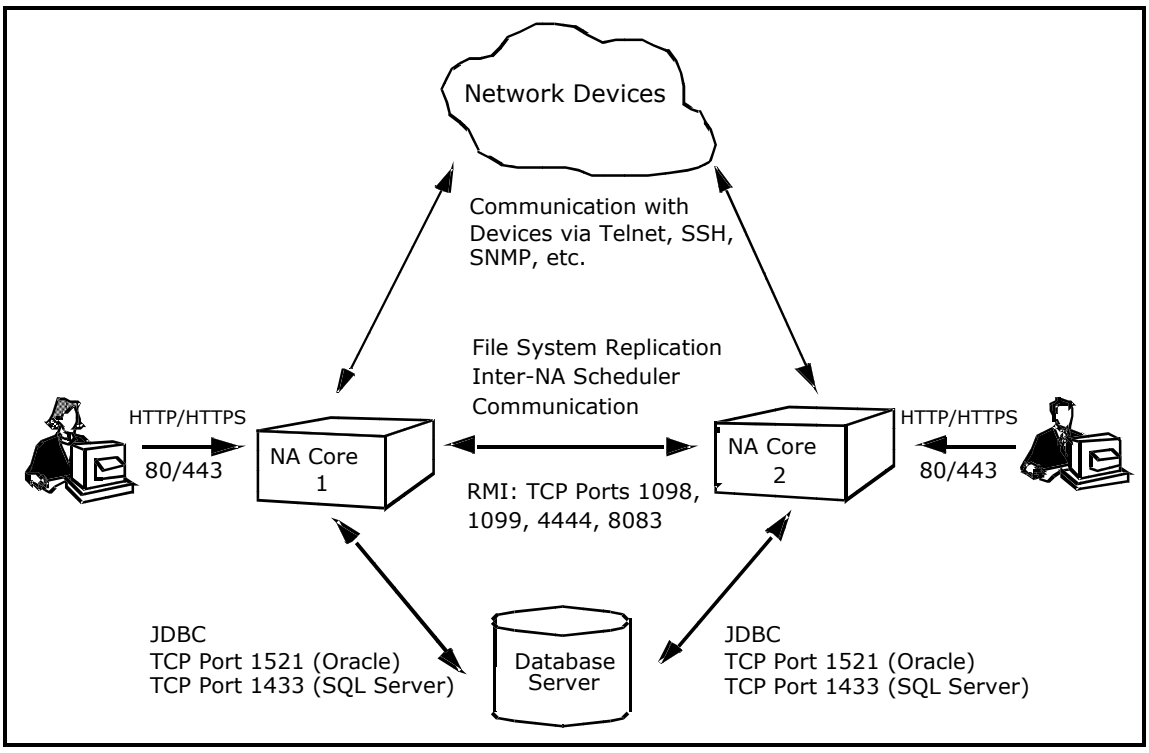

Keep in mind that some installations might have two separate network device collections that are inaccessible from the non-managing NA Core. These installations still provide scaling benefits, however.

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.

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

- RMI connectivity between each NA Core
- NA server timestamp differences between each NA Core
- Local NA Core identification

These monitors generate events during error conditions. The events can be emailed to the NA administrator using a standard event rule in NA. Refer to the *HP Network Automation 9.10 User's Guide* for information on configuring email notification.

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

In a scenario where the NA Core is the bottleneck in terms of task throughput or overall performance, NA Horizontal Scalability can ease the problem. For example, a NA Core could become a bottleneck if:

- The NA Core's memory or CPU usage is high due to running many concurrent tasks
- The NA Core's memory or CPU usage is high due to many concurrent users on the system

- NA is not able to run as many concurrent tasks as you would like due to memory restrictions

Note: Using NA Horizontal Scalability could place an addition load on your database, as certain queries for the task scheduler could be duplicated across different NA Cores.

Architecture

The NA Horizontal Scalability architecture has multiple components:

- A NA Mesh scheduler that is optimized for group task throughput. The scheduler enables group tasks to be load-balanced across multiple cores. The scheduler is not currently optimized for task failover, however.
- A file system synchronizer. Certain file system items are kept synchronized across multiple NA Cores. These items include NA Core configuration files, device drivers, and software images.
- A monitoring subsystem that periodically checks the health of inter-NA Core communication.

Topologies

The NA Mesh can be configured in multiple topologies, depending on your requirements. The following general topologies are supported:

- **Topology 1**: Dedicating each server to a fixed set of devices. You should use this topology if you have business units that need their own servers for performance reasons. This topology allows for global reporting on all devices.

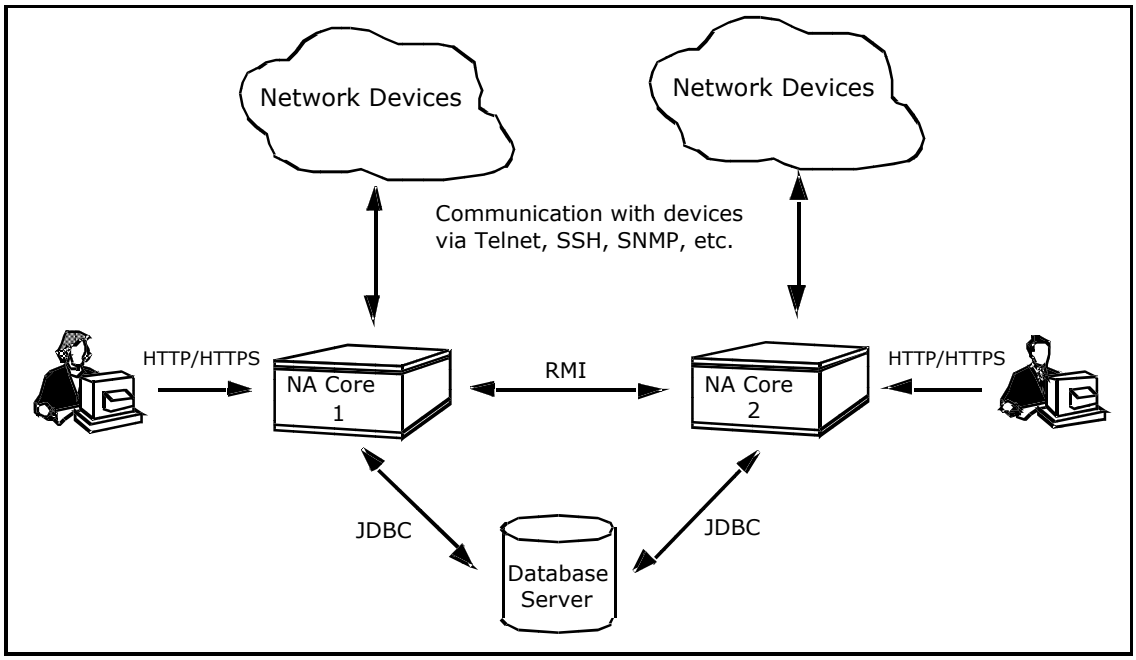

- **Topology 2**: Load-balancing all device group tasks across all servers. This topology is useful if you want to maximize task throughput across your entire NA Mesh. You should use this topology if you have performance requirements on task throughput, but do not need to limit access to certain devices for different servers.

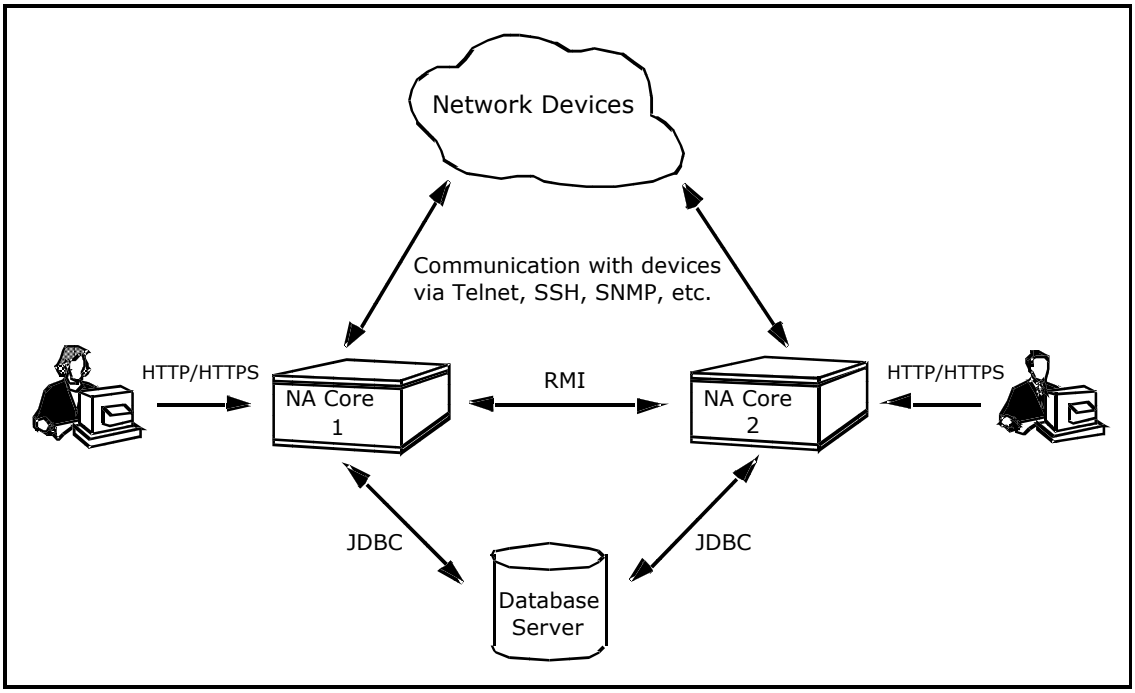

- **Topology 3**: Using one server to run automated regularly-scheduled group tasks, while another server is dedicated to running all user-initiated device tasks. This configuration is useful if you want to have user-initiated tasks avoid a scheduler bottleneck due to large automated tasks that are running. You should use this topology if it is critical that user-initiated tasks run as soon as possible.

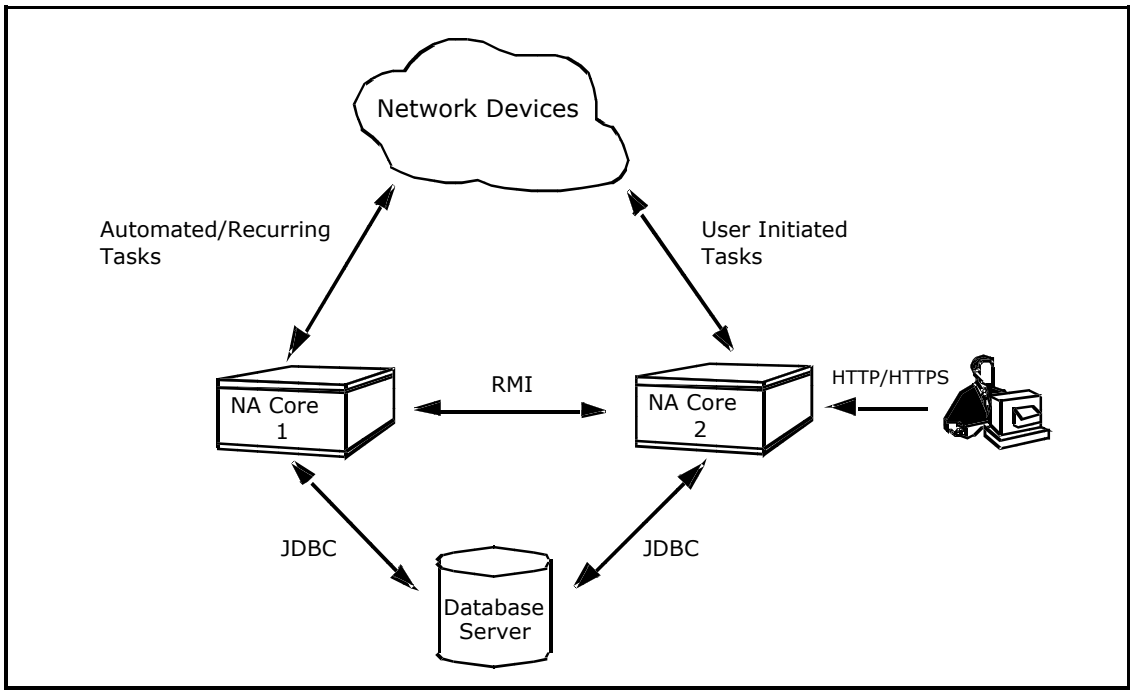

- **Topology 4**: Using one server for all user-interaction, while another server is used for all device tasks. This is useful if NA users experience UI performance limitations due to device tasks loads on a specific NA Core. You should use this topology if users run multiple search and reporting operations that impact overall system performance.

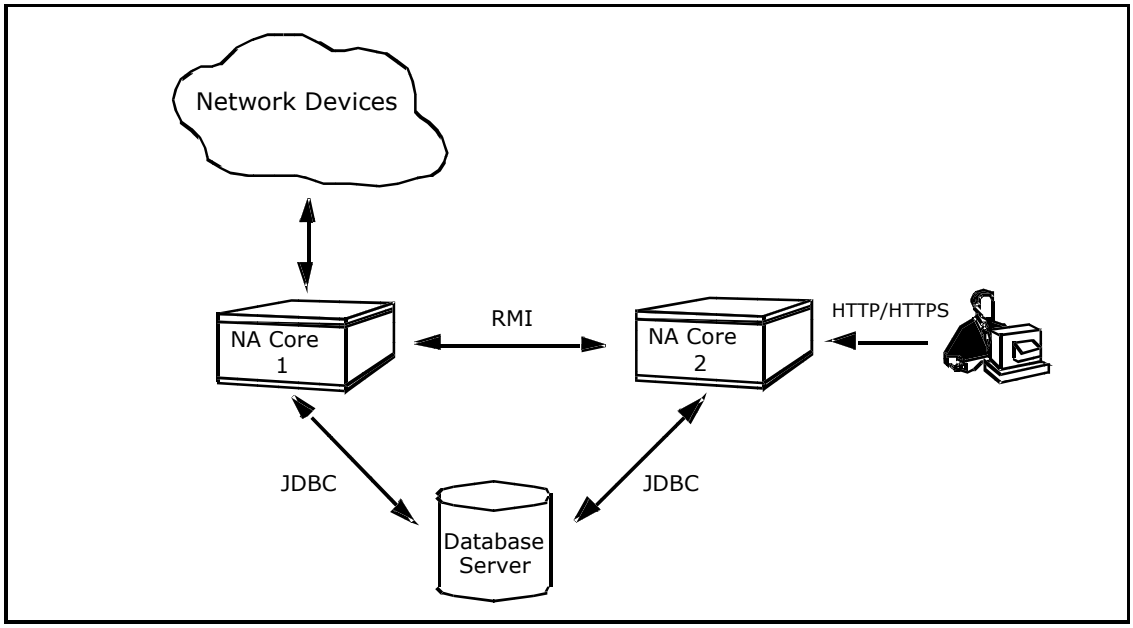

Note: Horizontal Scalability in conjunction with Multimaster Distributed Systems is not a supported configuration due to the replication load that Horizontal Scalability generates. Refer to the *HP Network Automation Scalability Architecture and Support Statement* for more information.

Configuration with NA Satellites

When configuring NA Satellites with Horizontal Scalability, each Satellite Remote Agent is associated with a NA Core. The NA Core associated with a Satellite receives snapshot requests when Syslog messages indicate a Snapshot is required.

When using Topology 1, the NA Core associated with a Satellite should be the same NA Core associated with the Partition for the Realm of the Satellite. For example,

- Partition Bellevue is associated with the Bellevue Realm and NA Core A.
- Satellite Gateway Bellevue1 is installed in the Bellevue Realm.
- The Remote Agent deployed to Gateway Bellevue1 should also be NA Core A.

When using Topology 2 and 3, the NA Cores associated with the Satellites should be distributed evenly between the NA Cores.

Load-Balancing

NA Horizontal Scalability enables multiple NA Cores to communicate with a single database. This configuration uses several methods to boost NA performance:

- You can configure the NA Cores to execute tasks in a round-robin fashion across the NA Cores. NA tasks are not bound to the NA Core that owns the device. Note that the round-robin algorithm is used only for group tasks, and occurs at group sub-task creation time.
- You can configure one or more NA Cores to run all tasks that were locally created on it.
- NA can be configured to only look at in-memory running tasks when searching for duplicate tasks scheduled to run at the same time on the same device.

Note: Load-balancing assumes that the database is not the limiting factor for performance. If the database performance is the bottleneck, NA Horizontal Scalability will not help. In fact, it could exacerbate the problem. Be sure to check the CPU and memory usage on all of your servers. If your NA Core is using a lot of CPU and memory, the NA Core is the bottleneck. Otherwise, the bottleneck could by the database.

Configuring NA Horizontal Scalability on Two NA Cores

The following steps describe how to configure two NA Cores communicating with one database. This configuration can be started using an existing NA Core or on two new NA Cores.

Note: SQL scripts are provided in the appropriate bundle described below. The SQL scripts should be run on your database using the appropriate tool, for example SQLPlus for Oracle or SQLServer Management Studio for SQL Server. *Appendix A: Running Scripts on Oracle* on page 24 includes information for running scripts on Oracle. *Appendix B: Running Scripts on SQL Server* on page 25 includes information on running scripts on SQL Server.

1. [optional] If you do not have an existing SQL Server or Oracle NA database server installed, perform a clean NA install (NA1) for either the SQL Server or Oracle database server. Refer to the *HP Network Automation 9.10 Upgrade and Installation User's Guide* in the Doc folder on the installation DVD to review the NA installation steps.
2. Record the IP addresses and the database instance name and credentials for the NA Cores and Database server.
3. Locate the *OracleHorizontalScalabilityBundle.zip* file or the *SQLServerHorizontalScalabilityBundle.zip file* in the oracle_horizontal_scalability folder or the sql_server_horizontal_scalability folder on the NA 9.10 Multimaster and Horizontal Scalability DVD. Be sure to use the scripts for the appropriate database type.
4. Stop the NA Cores/daemons (NA1).
5. Extract either the *OracleInitialSetup.sql* file or *SQLServerInitialSetup.sql* file from the appropriate replication script bundle. You will need to edit the file to replace certain variables. The variables to be replaced are located in the script between comment lines `BEGIN VARIABLE REPLACEMENT SECTION` and `END VARIABLE REPLACEMENT SECTION`. Be sure to completely replace the variables, including replacing the angle brackets (<>), as follows:

   `<REPLACEME_DATABASE_NAME>` Use the SQL Server database name or Oracle SID.
   `<REPLACEME_DATABASE_SERVER_NAME_OR_IP>` Use the DNS hostname or static IP address of your database server.
   `<REPLACEME_CORE_SERVER_NAME_OR_IP_1>` Use the DNS hostname or static IP address of the first NA Core.
   `<REPLACEME_CORE_SERVER_NAME_OR_IP_2>` Use the DNS hostname or static IP address of the second NA Core.

   Note: The CoreRMIPort in the script could be different from 1099 if your installation has been customized. The default DatabasePort is 1521 for Oracle and 1433 for SQL Server. If you are using a port other than the standard one for the database, you will need to modify the script accordingly.

Note: If you have previously configured a second Horizontal Scalability Core and then removed it, you will need to comment out the following line by putting two hyphens (--) at the beginning of the line:

```
-- INSERT INTO RN_KEY_INCREMENTOR SELECT Name,Value,2 FROM
RN_KEY_INCREMENTOR WHERE CoreID = 1;
```

Keep in mind that the `RN_KEY_INCREMENTOR` table will already be updated for the second Horizontal Scalability Core from before.

6. Run the appropriate setup script, for example *SQLServerInitialSetup.sql* or *OracleInitialSetup.sql*, that you edited in Step 5. (Refer to *Appendix A: Running Scripts on Oracle* on page 24 or *Appendix B: Running Scripts on SQL Server* on page 25 for information.)
7. On the second NA Core, install NA (NA2). When prompted, choose "Use Existing Database". The database name/SID is the same one you used in Step 5.
8. Stop the second NA Core Management Engine (NA2). Refer to the *HP Network Automation 9.10 User's Guide* for information on how to stop the NA Management Engine.
9. On all of the NA Cores, edit and copy the *distributed.rcx* file from the appropriate *HorizontalScalabilityBundle.zip* to C:\NA\jre\ (Windows) or equivalent Unix directory, for example /opt/NAfiles/jre/. Make sure the following options are set accordingly for Oracle.

```
<options>
…
   <option name="distributed/enabled">true</option>
   <option name="distributed/replication_enabled">false</option>
   <option name="distributed/type">Oracle</option>
   <option name="distributed/cache_delay">30</option> <!-- in seconds -->
   <option name="distributed/bind_tasks_to_core">true</option>
…
</options>
```

The following options are for SQL Server:

```
<options>
…
   <option name="distributed/enabled">true</option>
   <option name="distributed/replication_enabled">false</option>
   <option name="distributed/type">SQLServer</option>
   <option name="distributed/cache_delay">120</option> <!-- in seconds-->
   <option name="distributed/bind_tasks_to_core">true</option>
…
</options>
```

Note: <option name="distributed/bind_tasks_to_core">false</option> enables the round-robin algorithm. If you are configuring your NA Mesh topology for round-robin, this should be set to false.

10. Restart all the NA Cores/daemons.

11. Verify your installation is working correctly. Refer to “Verifying Installation and Setup” on page 15).

Adding Additional NA Cores

To add a third, or more NA Cores to an existing NA Mesh, do the following:

1. Record the IP addresses and the database instance name and credentials for the NA Core you are adding.
2. Locate the *HP Network Automation 9.10 User's Guide* in the Doc folder on the installation DVD to review the NA installation steps.
3. Locate the *OracleHorizontalScalabilityBundle.zip* file or the *SQLServerHorizontalScalabilityBundle.zip file* in the oracle_horizontal_scalability folder or the sql_server_horizontal_scalability folder on the NA 9.10 Multimaster and Horizontal Scalability DVD. Be sure to use the scripts for the appropriate database type.
4. Stop the NA Core Management Engines on all the other servers.
5. Extract either the *OracleAddServer.sql* file or *SQLServerAddServer.sql* file from the appropriate replication script bundle. You will need to edit the file to replace certain variables. The variables to be replaced are located in the script between comment lines `BEGIN VARIABLE REPLACEMENT SECTION` and `END VARIABLE REPLACEMENT SECTION`. Completely replace the variables, including replacing the angle brackets (<>), as follows:

   `<REPLACEME_DATABASE_NAME>` Use the SQL Server database name or Oracle SID.
   `<REPLACEME_DATABASE_SERVER_NAME_OR_IP>` Use the DNS hostname or static IP address of your database server.
   `<REPLACEME_ADDED_CORE_SERVER_NAME_OR_IP>` Use the DNS hostname or static IP address of the NA Core you are adding.

   Note: The CoreRMIPort in the script could be different from 1099 if your installation has been customized. The default DatabasePort is 1521 for Oracle and 1433 for SQL Server. If you are using a port other than the standard one for the database, you will need to modify the script accordingly.
6. Run the appropriate add server script, for example *OracleAddServer.sql* or *SQLServerAddServer.sql* that you edited in Step 5. (Refer to *Appendix A: Running Scripts on Oracle* on page 24 or *Appendix B: Running Scripts on SQL Server* on page 25 for information.)
7. On the new NA Core, install NA (NA_*n*). When prompted, choose "Use Existing Database". The database name/SID is the same one you used in step 5 (i.e. dbname1).
8. Stop the new NA Core Management Engine (NA_*n*).
9. On the new NA Core, edit and copy the *distributed.rcx* file from the appropriate *HorizontalScalabilityBundle.zip* to *C:\NA\jre* (Windows) or equivalent Unix directory, for example /opt/NAfiles/jre/. Make sure the following options are set accordingly for Oracle.

```
<options>
…
   <option name="distributed/enabled">true</option>
   <option name="distributed/replication_enabled">false</option>
   <option name="distributed/type">Oracle</option>
   <option name="distributed/cache_delay">30</option> <!-- in seconds -->
   <option name="distributed/bind_tasks_to_core">true</option>
…
</options>
```

The following options are for SQL Server:

```
<options>
…
   <option name="distributed/enabled">true</option>
   <option name="distributed/replication_enabled">false</option>
   <option name="distributed/type">SQLServer</option>
   <option name="distributed/cache_delay">120</option> <!-- in seconds-->
   <option name="distributed/bind_tasks_to_core">true</option>
…
</options>
```

Note: <option name="distributed/bind_tasks_to_core">false</option> enables the round-robin algorithm. If you are configuring your NA Mesh topology for round-robin, this should be set to false.

10. Restart all the NA Cores/daemons.

Verifying Installation and Setup

To verify installation and setup, do the following:

1. Verify that the RN_CORE table database contains the appropriate list of servers in the NA Mesh. To check the RN_CORE table, using the query tool appropriate for your database, enter:

   ```
   SELECT * FROM RN_CORE;
   ```

   You can also view the Edit Core page in NA to review each NA Core's settings. Refer to page 23 for information.

2. Edit an object on one NA server (for example, a Comments field for a device).
3. Verify that the updated comment exists on the second NA server.
4. Check the status of the Distributed Monitor in the NA UI to ensure that no problems are being reported. It could take up to five minutes for this monitor to initially run.

Configuring Topology After Setup

The following general topologies are supported:

- **Topology 1**: Dedicating each server to a fixed set of devices.
- **Topology 2**: Load-balancing all device group tasks across all servers.
- **Topology 3**: Using one server to run automated regularly-scheduled group tasks, while another server is dedicated to running all user-initiated device tasks.
- **Topology 4**: Using one server for all user-interaction, while another server is used for all device tasks.

Option	Topology 1 Value	Topology 2 Value	Topology 3 Value	Topology 4 Value
"distributed/bind_tasks_to_core" (in *distributed.rcx*)	true	false	true	true
Allow local task running: Admin Settings→ Server→ Performance Tuning **Note**: In a Distributed System, this option notifies the entire NA Mesh that there can be one or more NA Cores in the NA Mesh that locally run all tasks scheduled on those NA Cores. This is a NA Mesh-wide global option. If you need to change this option, you must restart all of the NA Management Engines after making the change.	Unchecked (i.e., false)	Unchecked (i.e., false)	Checked (i.e., true)	Unchecked (i.e., false)
Run tasks locally: Admin Settings→ Server→ Performance Tuning **Note**: In a Distributed System, this option notifies the NA Core to locally run all tasks created on it. This option will only take effect if the NA Mesh-wide option is set and is NA Core-specific.	Unchecked	Unchecked	Checked (on the server where tasks should run locally) Unchecked (on the other servers)	Unchecked

Post Installation Setup

Once you have a functioning replication system, there are additional steps you can take to complete setup:

1. Add new Sites — This will enable you to partition your devices across the different NA Cores in the NA Mesh.
2. Add new Realm definitions — A Realm is a network segment. A Partition is not required to be in the same Realm as its managing NA Core. Keep in mind that a Realm is a large area that can include many Partitions. However, a Realm does not have to include any NA Cores. Typically, a Realm is identified by a set of unique IP addresses. For example, a Realm cannot contain two devices numbered as 10.255.111.128. Instead, the devices must be broken out into separate Realms. (Refer to the *HP Network Automation 9.10 User's Guide* for information.)

Uninstall Procedures

If you want to remove NA Horizontal Scalability from two NA Cores and return to a single NA Core configuration, do the following:

Note: If you are using more than two NA Cores, the following steps need to be applied to each of the NA Cores you are removing.

1. Stop and disable as appropriate the NA Cores/daemons on the NA Core that is being removed.
2. On the database server, run the following script as appropriate for your database type:

   Oracle Script

```
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>;
COMMIT;
```

   SQL Server Script

```
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>;
```

   Note: Change <coreID> as appropriate. The script assumes you do not want to remove NA Core 1.

3. Remove the *distributed.rcx* file from NA Core 1 (assuming you want to leave only NA Core 1).
4. Restart NA on NA Core 1.

Managing Your Horizontal Scalability NA Mesh

In addition to replication monitoring, NA also monitors the following for each of the NA Cores:

- RMI connectivity between each NA Core
- NA server timestamp differences between each NA Core
- Local NA Core identification

These monitors will generate events during error conditions. The events can be emailed to the NA administrator using a standard event rule in NA.

NA Generated Events

By default, NA generates system events. Event rules can alert you to certain error conditions requiring attention. 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.

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.

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. The following options are included under the Admin→ Distributed option:

- Monitor Results
- Site Reassignment
- Core List
- Renew Configuration Options

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 monitors 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.
- 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. Refer to "Edit Core Page" on page 22 for information. (**Note**: This condition could be displayed in the error logs with 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>
  ```

  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.

Refer to the *HP Network Automation 9.10 Multimaster Distributed System on Oracle User's Guide* or the *HP Network Automation 9.10 Multimaster Distributed System on SQL Server User's Guide* for detailed information on the Distributed Monitor Results page.

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. Refer to the *HP Network Automation 9.10 Multimaster Distributed System on Oracle User's Guide* or the *HP Network Automation 9.10 Multimaster Distributed System on SQL Server User's Guide* for detailed information on the Site Reassignment page.

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. Refer to the *HP Network Automation 9.10 Multimaster Distributed System on Oracle User's Guide* or the *HP Network Automation 9.10 Multimaster Distributed System on SQL Server User's Guide* for detailed information on the List Cores page.

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.

Complete the following fields and click the Save Core button.

Refer to the *HP Network Automation 9.10 Multimaster Distributed System on Oracle User's Guide* or the *HP Network Automation 9.10 Multimaster Distributed System on SQL Server User's Guide* for detailed information on the Edit Core page.

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 Cores are in sync with the rest of the NA Mesh.

Refer to the HP *Network Automation 9.10 Multimaster Distributed System on Oracle User's Guide* or the *HP Network Automation 9.10 Multimaster Distributed System on SQL Server User's Guide* for detailed information on the Renew Configuration Options page.

Known Issues and Limitations

The network connectivity between different NA Cores and Database servers must be adequate for the NA Horizontal Scalability to function properly. In networking terms, the latency must be low.

If there is a firewall in the network configuration, make sure all the appropriate TCP ports are opened for traffic to tunnel through the firewalls. For example, ports 1099 must be open between the NA Cores. Refer to the "Protocols, Databases, and Ports" section in the *HP Network Automation 9.10 Upgrade and Installation User's Guide* for detailed information on port usage.

If you want to round-robin device tasks, all NA Cores must have network connectivity to all devices. Using the round-robin algorithm requires that each server access the database to determine the full list of running tasks in the NA Mesh. This can incur additional performance demands on the database. In addition, the maximum concurrent tasks settings should be the same on all NA Cores in the NA Mesh.

In the event of a NA Core failure while running tasks, some tasks that were scheduled to run on the failed NA Core can remain in the running task state. You can delete these tasks and resubmit them to run again on a running NA Core.

Appendix A: Running Scripts on Oracle

To run your edited setup scripts on Oracle using SQLPlus, do the following:

1. Execute the following command from the appropriate shell prompt:
   `sqlplus <USER>/<PASSWORD>@<DATABASE_SID>`
2. At the "SQL>" prompt, either paste the script you have edited or use `@<ScriptName>.sql` to run the script. Substitute the appropriate script you want to run for `<ScriptName>`.

Note: The `<USER>`, `<PASSWORD>`, and `<DATABASE_SID>` are the appropriate user, password, and database SID that the NA Core uses to access the database.

Appendix B: Running Scripts on SQL Server

You can run your edited setup script using either SQLServer Management Studio or the sqlcmd command line utility.

To run your script using the sqlcmd command line utility, use the following command:

```
sqlcmd -S <SERVER_NAME> -U <USER> -P <PASSWORD> -d <DATABASE_NAME> -i
<SCRIPT_NAME>.sql
```

Where the variables should be substituted as follows:

`<SERVER_NAME>` The name of your database server.
`<USER>` The user that the NA Core uses to access the database.
`<PASSWORD>` The user's password.
`<DATABASE_NAME>` The database name NA uses.
`<SCRIPT_NAME>` The name of the script you want to run.

Index