

HP Network Automation

Software Version: 9.11

Horizontal Scalability User's Guide

Document Release Date: February 2013
Software Release Date: October 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 2008–2013 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

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.

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>

Contents

1	NA Horizontal Scalability Concepts	7
	Horizontal Scalability Architecture	7
	Horizontal Scalability Topologies	8
	Distributed Round Robin	9
	Distributed Round Robin Variation	9
	Distributed Round Robin Illustrated	10
	Core Binding	12
	Core Binding Variations	13
	Core Binding Illustrated	13
2	Configuring Horizontal Scalability	17
	Ports for Horizontal Scalability	17
	Configuring a Two NA Core Horizontal Scalability Environment	18
	Running the Horizontal Scalability Configuration Scripts	23
	Running Scripts on Oracle	24
	Running Scripts on SQL Server	24
	Verifying Installation and Setup	25
	Adding Additional NA Cores to the Horizontal Scalability Environment	25
	Configuring a Horizontal Scalability Topology Variation	28
	Configure the Distributed Round Robin Variation for User Interaction	28
	Configure the Core Binding Variation for User Interaction	29
	Configure the Core Binding Variation for Local Device Tasks	29
	Configuring NA Satellites in a Horizontal Scalability Environment	30
	Configuring Dynamic Device Group Calculation in a Horizontal Scalability Environment	32
	Uninstall Procedures	33
3	NA Core Failover in a Horizontal Scalability Environment	35
	Enabling Failover	36
	Disabling Failover	36
	Re-Enabling Failover	36
	NA Core Failover Behavior	37
	Delay Before Initiating Failover	38
	Task Status	38
	Movement of Tasks	39
	NA Failover Events	42
	Configuring Failover Order (Core Binding Only)	42
	Optional Actions after NA Core Failover Occurs	44
	Notify Users	44
	Configure Integrations to a Different NA Core	44

- Pause Non-Critical Tasks 44
- Locate Tasks Impacted by NA Core Failover 44
- Ensure Continuity of Dynamic Device Group Calculation. 45
- Stop the Core Gateway for the Stopped NA Core 45
- Failing Back to the Original NA Core. 46
- 4 System Administration 47**
 - NA-Generated Events for Horizontal Scalability. 47
 - Distributed System – Time Synchronization Warning. 47
 - Distributed System – RMI Error 47
 - Using the NA Distributed System Pages 48
 - Distributed Monitor Results Page 48
 - Distributed Error List 49
 - Distributed Conflict List. 49
 - Site Reassignment Page 49
 - Distributed Core List Page. 50
 - Edit Core Page 50
 - Device Password Rule Priority Reset Page 51
 - Renew Configuration Options Page 51
- We appreciate your feedback! 53**

1 NA Horizontal Scalability Concepts

With HP Network Automation Software (NA) Horizontal Scalability, multiple NA cores connect to one NA database. In this configuration, the NA cores work as a logical unit. Any NA console user sees the same information on each of the NA cores. NA distributes device tasks across the NA cores according to the Horizontal Scalability topology. Users do not need to know which NA core accesses a particular device.

NA Horizontal Scalability provides for scaling NA beyond the limits of a single NA core. Two to five NA cores can connect to one NA database. For scale information, see “Maximum Supported Managed Environment” in the *NA Support Matrix*.

As of NA 9.11, NA Horizontal Scalability also provides for high availability of NA with automated failover when an NA core becomes unavailable. For more information see [Chapter 3, NA Core Failover in a Horizontal Scalability Environment](#).



This document is updated as new information becomes available. 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>

For more information, see [Documentation Updates](#) on page 2.

This chapter includes the following topics:

- [Horizontal Scalability Architecture](#) on page 7
- [Horizontal Scalability Topologies](#) on page 8

Horizontal Scalability Architecture

In an NA Horizontal Scalability environment, the NA database is the primary source of the information shared across the NA cores. The database record for each task in the waiting queue includes the NA core on which that task will run.

Additionally, a lightweight channel exists between each pair of NA cores in the Horizontal Scalability environment. This channel uses Java Remote Method Invocation (RMI) for synchronization across the NA cores of certain file system objects, such as software images, driver packages, and the some of NA configuration.

The maximum number of NA cores in a Horizontal Scalability environment is five.

NA Horizontal Scalability is not designed for use over a WAN in which NA cores and the database are separated by WAN links. For best performance, it is recommended that each NA core be located in the same data center as the NA database server.

NA Horizontal Scalability is supported with an Oracle or Microsoft SQL Server database only. For information about supported versions, see the *NA Support Matrix*.

Horizontal Scalability Topologies

NA Horizontal Scalability can be configured for either of the following topologies:

- [Distributed Round Robin](#) on page 9
- [Core Binding](#) on page 12

The distinguishing feature of these topologies is the way NA assigns device tasks to the NA cores. A device task is an NA task that runs against a device.

Each of the supported topologies includes at least one variation for specialization of the NA cores in the Horizontal Scalability environment.

Core binding is the default NA Horizontal Scalability topology. However, if either topology could be used in your environment, distributed round robin is the recommended NA Horizontal Scalability topology.

The Horizontal Scalability environment can include one or more NA satellites for communicating with devices behind firewalls. For more information, see [Configuring NA Satellites in a Horizontal Scalability Environment](#) on page 30.

[Table 1](#) presents a high-level comparison of the supported Horizontal Scalability topologies.

Table 1 Comparing the Horizontal Scalability Topologies

Category	Distributed Round Robin	Core Binding
Network accessibility	All NA cores must be able to communicate with all managed devices.	Some or all NA cores cannot communicate with all managed devices.
Load sharing	Each NA core manages approximately the same number of devices. NA determines which NA core manages a given device.	Partitions determine which NA core manages which devices. The NA administrator assigns each partition to an NA core to distribute the device task load among the NA cores.
Failover	Automated with balanced load sharing. NA distributes responsibility for the devices that were managed by the unavailable NA core among the remaining NA cores.	Automated with uneven load sharing. All partitions from the unavailable NA core are assigned to one NA core. Device accessibility limitations might prevent successful failover.
Failback	NA automatically gives the recovered NA core responsibility for managed devices.	Manually assign partitions to the recovered NA core.

Distributed Round Robin

In the distributed round robin topology, each NA core can communicate with every device in the managed network, as shown in [Figure 1](#) on page 10. This configuration does not rely on the use of partitions, though partitions might be in use for other reasons. Within the configured access limits, users can log on to any of the NA cores and can create tasks to run against any device in the NA inventory. NA distributes the tasks across the NA cores. Additionally, users can see the results of any task on any NA core.

For a task that targets multiple devices with the same action, NA runs a lightweight group task that creates one child device task for each of the targeted devices. In the distributed round robin topology, the group task runs on the NA core where it was created, while each child device task could run on any NA core.

In the distributed round robin topology, NA distributes responsibility for the device tasks in the managed network across the available NA cores to provide approximate balance of the device task load. This distribution is based solely on the number of device tasks. It does not consider the NA core configuration or the size of device configurations. After a task for a given device is assigned to an NA core, all tasks for that device run on the same NA core until the number of running NA cores in the Horizontal Scalability environment changes. At that point, for new device tasks, NA recalculates the device distribution across the NA cores in the Horizontal Scalability environment.

Benefits of the distributed round robin topology include the following:

- Automated load sharing of device tasks across the NA cores.
- Automated failover and failback of tasks.
- By default, NA runs only one task on a device at any given time.



This behavior can be overridden for the Run Command Script task by clearing the **Wait option** check box.

Limitations of the distributed round robin topology include the following:

- Every device must be accessible to all NA cores or managed by an NA satellite.
- The NA administrator cannot control which NA core runs tasks on a given device.



Prior to NA 9.11, the distributed round robin topology was called “topology 2.”

Distributed Round Robin Variation

NA supports one variation of the standard distributed round robin topology: distributed round robin variation for user interaction.

This variation separates all device tasks from all user interaction with NA. User interaction includes all operational NA console sessions, NA proxy sessions, tools or programs that use the NA API, and connections from integrating products. See [Figure 2](#) on page 11.

The key benefit of this variation is that the resources of any NA core reserved for user interaction are fully devoted to responding to user requests. Additionally, this variation has the same benefits and limitations as for the standard distributed round robin topology.



This variation of the distributed round robin topology is new with NA 9.11.

Distributed Round Robin Illustrated

This section presents sample architecture drawings for a three-NA core Horizontal Scalability environment. It includes:

- Figure 1: Example Architecture: Standard Distributed Round Robin
- Figure 2: Example Architecture: Distributed Round Robin Variation for User Interaction

Figure 1 Example Architecture: Standard Distributed Round Robin





Figure 2 Example Architecture: Distributed Round Robin Variation for User Interaction



Core Binding

In the core binding topology, each NA core communicates with a fixed, non-overlapping set of devices, as shown in [Figure 3](#) on page 14. In this configuration, each device is associated with a partition, and each partition is associated with one NA core. Within the configured access limits, users can log on to any of the NA cores and can create tasks to run against any device in the NA inventory. NA assigns each task to the appropriate NA core for the targeted device. Additionally, users can see the results of any task on any NA core.

For a task that targets multiple devices with the same action, NA runs a lightweight group task that creates one child device task for each of the targeted devices. In the core binding topology, the group task runs on the NA core where it was created, while each child device task runs on the appropriate NA core for that device.

Benefits of the core binding topology include the following:

- Automated failover of partitions and tasks for topologies in which devices are reachable from multiple NA cores.
- Devices that are not accessible from all NA cores can be assigned to a specific NA core.
- The NA administrator controls which NA core runs tasks on a given device. This control can be important for heterogeneous environments. For example:
 - When the NA servers have different capacities, more devices can be assigned to the larger NA cores and fewer devices can be assigned to the smaller NA cores.
 - When the network includes some devices with very large configurations, the NA administrator can spread responsibility for these devices across the NA cores.

Limitations of the core binding topology include the following:

- No automated load sharing.
- The NA administrator must do extra NA configuration work to define the device partitions.
- Failback of partitions and tasks is manual.



Prior to NA 9.11, the core binding topology was called “topology 1.”

Core Binding Variations

NA supports the following variations of the standard core binding topology:

- Core binding variation for user interaction

This variation separates all device tasks from all user interaction with NA. User interaction includes all operational NA console sessions, NA proxy sessions, tools or programs that use the NA API, and connections from integrating products. One or more partitions are associated with each NA core that runs device tasks. No partitions are associated with an NA core that is designated for user interaction only. This variation maintains the strict relationship between each NA core and the devices associated with it through partitions. See [Figure 4](#) on page 15.

The key benefit of this variation is that the resources of any NA core designated for user interaction only are fully devoted to responding to user requests. Additionally, this variation has the same benefits and limitations as for the standard core binding topology.



Prior to NA 9.11, this variation of the core binding topology was called “topology 4.”

- Core binding variation for local device tasks

This variation separates the pre-planned device tasks (for example, daily snapshots) from the local device tasks that are created at the point of need (for example, an immediate update to an access control list). One or more partitions are associated with each NA core that runs pre-planned device tasks. No partitions are associated with an NA core that runs only local device tasks. The NA core for local device tasks bypasses the partition association to directly connect to any device targeted by a device task created on this NA core. See [Figure 5](#) on page 16.

The key benefit of this variation is that local device tasks run independently of the pre-planned tasks. This independence generally results in faster results for the local device tasks. Limitations of this variation include the following:

- Multiple NA cores might attempt to access one device at the same time.
- Pre-planned device tasks must be created on an NA core that is *not* reserved for local tasks.

Additionally, this variation has the same benefits and limitations as for the standard core binding topology.



Prior to NA 9.11, this variation of the core binding topology was called “topology 3.”

Core Binding Illustrated

This section presents sample architecture drawings for a three-NA core Horizontal Scalability environment. It includes:

- [Figure 3: Example Standard Core Binding Architecture](#)
- [Figure 4: Example Architecture: Core Binding Variation for User Interaction](#)
- [Figure 5: Example Architecture: Core Binding Variation for NA Local Device Tasks](#)

Figure 3 Example Standard Core Binding Architecture




Figure 4 Example Architecture: Core Binding Variation for User Interaction





Figure 5 Example Architecture: Core Binding Variation for NA Local Device Tasks



2 Configuring Horizontal Scalability

This chapter describes how to configure an HP Network Automation Software (NA) Horizontal Scalability environment. It includes the following topics:

- [Ports for Horizontal Scalability](#) on page 17
- [Configuring a Two NA Core Horizontal Scalability Environment](#) on page 18
- [Running the Horizontal Scalability Configuration Scripts](#) on page 23
- [Verifying Installation and Setup](#) on page 25
- [Adding Additional NA Cores to the Horizontal Scalability Environment](#) on page 25
- [Configuring a Horizontal Scalability Topology Variation](#) on page 28
- [Configuring NA Satellites in a Horizontal Scalability Environment](#) on page 30
- [Configuring Dynamic Device Group Calculation in a Horizontal Scalability Environment](#) on page 32
- [Uninstall Procedures](#) on page 33

Ports for Horizontal Scalability

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.

Table 2 Ports Used in an NA Horizontal Scalability Environment

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 NA servers	
JNDI (ports 1098, 1099)	NA server to NA server. You can change this by editing the NA configuration files. Contact your Support representative for assistance.

Table 2 Ports Used in an NA Horizontal Scalability Environment (cont'd)

Protocol/Database/Port	From/To
RMI (ports 4444 and 8083)	NA server to NA server. You can change this by editing the NA configuration files. Contact your Support representative for assistance.
Between the NA server and the database server	
Oracle (port 1521)	From the NA server to an Oracle database server.
Microsoft 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. Contact your Support representative for assistance.
Telnet (port 23 - Windows or 8023 - Solaris/Linux)	From the NA client to the NA server. For information about changing the port number, see "Telnet/SSH Page Fields" in the NA help.
SSH (port 22 - Windows or 8022 - Solaris/Linux)	From the NA client to the NA server. For information about changing the port number, see "Telnet/SSH Page Fields" in the NA help.

Configuring a Two NA Core Horizontal Scalability Environment

First, set up and verify a two NA core Horizontal Scalability environment. After this environment works correctly, add additional NA cores if necessary.

NA can be already installed on one NA server and one database server. Alternatively, you can complete the first NA core installation as part of setting up NA Horizontal Scalability.

This procedure identifies the NA servers as NA1 and NA2. To use a different identifier, substitute the actual value for each instance of the example value within the procedure.

To set up a two NA core Horizontal Scalability environment, complete the following tasks:

- [Task 1: Verify Prerequisites for All Servers](#) on page 19
- [Task 2: \(New Installations Only\) Install the First NA Core](#) on page 20
- [Task 3: Configure Horizontal Scalability for the First Two NA Cores](#) on page 20
- [Task 4: Configure the Standard Horizontal Scalability Topology](#) on page 22
- [Task 5: Optional. Configure Additional syslog Destinations for Continuity after NA Core Failover](#) on page 23

After verifying the two NA core Horizontal Scalability configuration, optionally complete the following tasks:

- [Adding Additional NA Cores to the Horizontal Scalability Environment](#) on page 25
- [Configuring a Horizontal Scalability Topology Variation](#) on page 28
- [Configuring NA Satellites in a Horizontal Scalability Environment](#) on page 30
- [Configuring Dynamic Device Group Calculation in a Horizontal Scalability Environment](#) on page 32

The setup files for NA Horizontal Scalability are the standard setup files for a normal NA installation, with the addition of an installation bundle for configuring Horizontal Scalability. The name and location of this bundle depends on the database type:

- *Oracle*: The `OracleHorizontalScalabilityBundle.zip` file is in the `oracle_horizontal_scalability` folder on the NA Multimaster and Horizontal Scalability DVD.
- *Microsoft SQL Server*: The `SQLServerHorizontalScalabilityBundle.zip` file is in the `sql_server_horizontal_scalability` folder on the NA Multimaster and Horizontal Scalability DVD.

The setup files include SQL scripts to be customized and run on the database server.

Task 1: [Verify Prerequisites for All Servers](#)

Prepare one database server and two NA servers. For best performance, all NA servers should be co-located with the database server.

- 1 Verify that the following prerequisites have been met:
 - All servers that will run NA have working hostnames. Note the following:
 - Each NA server should have a high-speed connection to the database server.
 - For NA server hardware and operating system requirements, see the *NA Support Matrix*.
 - It is recommended that the host names of the database server and both NA servers are in the `hosts` file on each NA server. This file is located as follows:
 - *Windows*: `<Drive>:\Windows\System32\drivers\etc\hosts`
 - *UNIX*[®]: `/etc/hosts`
 - The database server and all NA servers are set to use the same time and time zone. It is recommended to synchronize the servers with an external time service.
 - The ports listed in [Table 2](#) on page 17 are open.
- 2 Record the following information:
 - Database server:
 - IP address
 - Administrator or root credentials
 - Database instance name
 - Each NA server:
 - IP address
 - Administrator or root credentials

Task 2: (New Installations Only) Install the First NA Core

If NA has not yet been installed, install NA on NA1 by following these steps:

- 1 If NA is not currently installed, install NA on the first NA server (NA1) as described in the *NA Upgrade and Installation Guide*.
- 2 After NA installation is complete, log on to NA to ensure that it works as expected.

Task 3: Configure Horizontal Scalability for the First Two NA Cores

To connect the NA servers and the database server for Horizontal Scalability, follow these steps:

- 1 On NA1, unpack the Horizontal Scalability bundle to a known location.
- 2 In a text editor, such as WordPad or vi, open the initial setup script in the known location of [step 1](#) on page 20.
 - *Oracle*: OracleInitialSetup.sql
 - *SQL Server*: SQLServerInitialSetup.sql
- 3 Edit the initial setup script to completely replace the variables, including the angle brackets (<>), with information about the NA environment.
 - Replace <REPLACEME_DATABASE_NAME> with the Oracle SID, the service name of the Oracle RAC cluster, or the SQL Server database name.
 - Replace <REPLACEME_DATABASE_SERVER_NAME_OR_IP> with the DNS hostname or the static IP address of the database server.
 - Replace <REPLACEME_CORE_SERVER_NAME_OR_IP_1> with the DNS hostname or the static IP address of NA1.
 - Replace <REPLACEME_CORE_SERVER_NAME_OR_IP_2> with the DNS hostname or the static IP address of NA2.

For example:

- *Oracle*:

```
database_name := 'NA_SID';
database_server_name_or_ip := 'nadb.example.com';
core_server_name_or_ip_1 := 'na1.example.com';
core_server_name_or_ip_2 := 'na2.example.com';
```

- *SQL Server*:

```
SET @database_name = 'NA_DB';
SET @database_server_name_or_ip = 'nadb.example.com';
SET @core_server_name_or_ip_1 = 'na1.example.com';
SET @core_server_name_or_ip_2 = 'na2.example.com';
```

- 4 In the initial setup script, also do the following:
 - If NA is configured to use a non-default port for RMI communication, modify the value of CoreRMIPort.

Also, in the VALUES block, replace 1099 with the port in use.
 - If the database server is configured to use a non-default port for communicating with NA, modify the value of DatabasePort.

Also, in the VALUES block, replace 1520 (Oracle) or 1433 (SQL Server) with the port in use.

- Set `TimezoneOffset` to the value that matches the time zone setting for all NA servers and the database server in the Horizontal Scalability environment.

Also, in the `VALUES` block, replace `-8` with the value that matches the time zone setting for all NA servers and the database server in the Horizontal Scalability environment.

- If you have previously configured a second Horizontal Scalability NA core and then removed it, the `RN_KEY_INCREMENTOR` table will already be updated for the second Horizontal Scalability NA core. 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;
```

- 5 Copy the customized initial setup script to the database server.
- 6 Stop all NA services on NA1.
 - *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
```
- 7 If NA contains production data, back up the NA file system and the NA database.
- 8 On the database server, run the initial setup script as described in the appropriate procedure for the database type:
 - [Running Scripts on Oracle](#) on page 24
 - [Running Scripts on SQL Server](#) on page 24
- 9 On the second NA server (NA2), install NA. When prompted, choose **an existing HP Network Automation database**, and then enter the database name from the initial setup script of [step 3](#) on page 20.
- 10 Stop all NA services on NA2.
 - *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
```

Task 4: Configure the Standard Horizontal Scalability Topology

On one NA core, customize the `distributed.rcx` file to configure the Horizontal Scalability environment for a standard topology. (The topology variations are configured in other locations, as described in [Configuring a Horizontal Scalability Topology Variation](#) on page 28.)

- 1 In a text editor, such as WordPad or vi, open the `distributed.rcx` file in the known location of [Task 3, step 1](#) on page 20.

- 2 Add the `distributed/horizontalscalability` option.

```
<option name="distributed/horizontalscalability">true</option>
```



As of NA 9.11, the `distributed/horizontalscalability` option is required.

- 3 Add the `distributed/bind_tasks_to_core` option for the Horizontal Scalability topology.

- *Distributed Round Robin*: Add the following line:

```
<option name="distributed/bind_tasks_to_core">false</option>
```

- *Core Binding*: Add the following line:

```
<option name="distributed/bind_tasks_to_core">true</option>
```



Omitting the `distributed/bind_tasks_to_core` option results in the core binding topology.

- 4 Save the `distributed.rcx` file, and then copy this file to the following directory on both NA cores (NA1 and NA2):

- *Windows*: `<NA_HOME>\jre`
- *UNIX*: `<NA_HOME>/jre`

- 5 Add options to the `site_options.rcx` and `appserver.rcx` files as described in the additional information for fix QCCR1D95740 in the NA 9.11 (or later) patch `readme.txt` file.

- 6 Start all NA services on both NA cores (NA1 and NA2).

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Start**:

- **TrueControl ManagementEngine**
- **TrueControl FTP Server**
- **TrueControl SWIM Server**
- **TrueControl Syslog Server**
- **TrueControl TFTP Server**

- *UNIX*: Run the following command:

```
/etc/init.d/truecontrol start
```

- 7 *Core Binding only*. Create partitions of the managed devices, and then assign one or more partitions to each NA core that will run pre-planned device tasks.

- a Log on to the NA console for any NA core as an NA administrator.
- b Open the Partitions page (**Admin > Security Partitions**), and then click **New Partition**.

- c On the New Partition page, do the following:
 - Enter the partition name and description.
 - Select the NA core to manage this partition.



If you plan to configure a core binding topology variation, do not assign any partitions to an NA core that will be designated for user interaction only or for local device tasks.

- Use the device selector to specify the devices to include in the partition.



To simplify device selection, use the Search for Device page to locate the devices that match specific criteria, and then create a device group from the search results.

- d Repeat [step b](#) and [step c](#) until each managed device is included in a partition.
For more information, see “Partitions” in the NA help.

- 8 Verify that the installation is working correctly as described in [Verifying Installation and Setup](#) on page 25.
- 9 Choose your next step:
 - To add one or more additional NA cores to the Horizontal Scalability environment, follow the procedure in [Adding Additional NA Cores to the Horizontal Scalability Environment](#) on page 25.
 - If the Horizontal Scalability environment will contain only two NA cores, optionally configure the environment with a topology variation as described in [Configuring a Horizontal Scalability Topology Variation](#) on page 28.
 - If the Horizontal Scalability environment will contain only two NA cores and uses a standard topology, optionally add one or more NA satellites to the environment as described in [Configuring NA Satellites in a Horizontal Scalability Environment](#) on page 30.

Task 5: [Optional. Configure Additional syslog Destinations for Continuity after NA Core Failover](#)

To maintain immediate processing of syslog messages after NA core failover, configure the managed devices (or the syslog relay) to send syslog messages to two or more NA cores in the Horizontal Scalability environment.



Such configuration requires that each managed device can reach the selected NA cores, which might not be possible in an environment that uses a core binding topology.

Without this redundant configuration, after NA core failover NA detects configuration changes only at the next scheduled snapshot.

Running the Horizontal Scalability Configuration Scripts

After customizing each SQL script, run it as described in the appropriate procedure for the database type:

- [Running Scripts on Oracle](#) on page 24
- [Running Scripts on SQL Server](#) on page 24

Running Scripts on Oracle

To run a customized setup script on an Oracle database server using SQLPlus, follow these steps:

- 1 Copy the customized SQL script to the database server.
 - *Windows*: Place the file in C:\.
 - *UNIX*: Place the file in the \$ORACLE_HOME/bin directory, for example /u01/app/oracle/product/11.2.0/dbhome_1/bin.

- 2 Log on to a SQLPlus window as the NA database user. For example:

```
sqlplus <USER>/<PASSWORD>@<SID>
```

- 3 In the SQLPlus window, run the customized script. For example:

```
do @OracleInitialSetup.sql
```

For <USER> and <PASSWORD>, use the Oracle user account for the NA database user.

For <SID>, use the Oracle SID of the NA database.

Running Scripts on SQL Server

You can run a customized setup script using SQL Server Management Studio or the sqlcmd command.

To run a customized setup script on a SQL Server database server using the sqlcmd command, follow these steps:

- 1 Copy the customized SQL script to a known location on the 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.

- 2 From the known location on the database server, run the customized script by using the sqlcmd command. For example:

```
sqlcmd -S <Server> -U <User> -P <Password> -d <Database_Name>  
-i SQLServerInitialSetup.sql
```

For <Server>, use the short hostname of database server, for example nadb.

For <User> and <Password>, use the SQL Server user account for the NA database user.

For <Database_Name>, use the name of the NA database.

If necessary, replace SQLServerInitialSetup.sql with the name of the script to run.

Verifying Installation and Setup

To verify installation and setup, on *each* NA core, follow these steps:

- 1 Log on to the NA console as an NA administrator.
- 2 Open the List Cores page (**Admin > Distributed > List Cores**).
- 3 Verify that the list includes all NA cores with the expected status for each NA core.
- 4 Verify that the information on the List Cores page is identical in each NA console.

Adding Additional NA Cores to the Horizontal Scalability Environment

This procedure identifies the new NA server as NA3.

To add an NA core to an existing NA Horizontal Scalability environment, follow these steps:

- 1 On NA3, unpack the Horizontal Scalability bundle to a known location.
Alternatively, located the unpacked bundle on NA1.
- 2 In a text editor, such as WordPad or vi, open the add server script in the known location of [step 1](#) of this task.
 - *Oracle*: OracleAddServer.sql
 - *SQL Server*: SQLServerAddServer.sql
- 3 Edit the add server script to completely replace the variables, including the angle brackets (<>), with information about the NA environment.
 - Replace <REPLACEME_DATABASE_NAME> with the Oracle SID, the service name of the Oracle RAC cluster, or the SQL Server database name.
 - Replace <REPLACEME_DATABASE_SERVER_NAME_OR_IP> with the DNS hostname or the static IP address of the database server.
 - Replace <REPLACEME_ADDED_CORE_SERVER_NAME_OR_IP> with the DNS hostname or the static IP address of NA3.

For example:

- *Oracle*:


```
database_name := 'NA_SID';
database_server_name_or_ip := 'nadb.example.com';
added_core_server_name_or_ip := 'na3.example.com';
```
- *SQL Server*:


```
SET @database_name = 'NA_DB';
SET @database_server_name_or_ip = 'nadb.example.com';
SET @added_core_server_name_or_ip = 'na3.example.com';
```

- 4 In the VALUES block of the add server script, also do the following:
 - If NA is configured to use a non-default port for RMI communication, replace *1099* with the port in use.
 - If the database server is configured to use a non-default port for communicating with NA, replace *1520* (Oracle) or *1433* (SQL Server) with the port in use.
 - Replace *-8* with the value that matches the time zone setting for all NA servers and the database server in the Horizontal Scalability environment.
 - If you have previously configured a third Horizontal Scalability NA core and then removed it, the `RN_KEY_INCREMENTOR` table will already be updated for the third Horizontal Scalability NA core. Comment out the following line by putting two hyphens (`--`) at the beginning of the line:

```
-- INSERT INTO RN_KEY_INCREMENTOR SELECT Name, Value, core_number FROM
RN_KEY_INCREMENTOR WHERE CoreID = 1;
```
- 5 Copy the customized add server script to the database server.
- 6 Stop all NA services on all NA cores.
 - *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
```
- 7 If NA contains production data, back up the NA file systems and the NA database.
- 8 On the database server, run the add server script as described in the appropriate procedure for the database type:
 - [Running Scripts on Oracle](#) on page 24
 - [Running Scripts on SQL Server](#) on page 24
- 9 On the new NA server (NA3), install NA. When prompted, choose **an existing HP Network Automation database**, and then enter the database name from the add server script of [step 3](#) on page 25.

- 10 Stop all NA services on NA3.
 - *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
```
- 11 To add a fourth or fifth NA core to the Horizontal Scalability environment, repeat [step 2](#) on page 25 through [step 10](#) on page 27 for each additional NA core.
- 12 Copy the customized `distributed.rcx` file from an existing NA core to all new NA cores. This file is located in the following directory:
 - *Windows*: `<NA_HOME>\jre`
 - *UNIX*: `<NA_HOME>/jre`
- 13 Start all NA services on all NA cores.
 - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Start**:
 - **TrueControl ManagementEngine**
 - **TrueControl FTP Server**
 - **TrueControl SWIM Server**
 - **TrueControl Syslog Server**
 - **TrueControl TFTP Server**
 - *UNIX*: Run the following command:

```
/etc/init.d/truecontrol start
```
- 14 *Core Binding only*. On the Partitions page (**Admin > Security Partitions**) in the NA console, assign one or more partitions to each NA core that will run pre-planned device tasks. For more information, see “Partitions” in the NA help.
- 15 Verify that the installation is working correctly as described in [Verifying Installation and Setup](#) on page 25.
- 16 Choose your next step:
 - Optionally configure the environment with a topology variation as described in [Configuring a Horizontal Scalability Topology Variation](#) on page 28.
 - If the Horizontal Scalability environment uses a standard topology, optionally add one or more NA satellites to the environment as described in [Configuring NA Satellites in a Horizontal Scalability Environment](#) on page 30.

Configuring a Horizontal Scalability Topology Variation

The procedure for configuring Horizontal Scalability sets NA Horizontal Scalability for either the standard distributed round robin topology or the standard core binding topology. To configure the NA Horizontal Scalability environment for one of the topology variations, follow the appropriate procedure:

- [Configure the Distributed Round Robin Variation for User Interaction](#) on page 28
- [Configure the Core Binding Variation for User Interaction](#) on page 29
- [Configure the Core Binding Variation for Local Device Tasks](#) on page 29

After completing the appropriate procedure for configuring a Horizontal Scalability topology variation, optionally add one or more NA satellites to the environment as described in [Configuring NA Satellites in a Horizontal Scalability Environment](#) on page 30.

Configure the Distributed Round Robin Variation for User Interaction

If the NA Horizontal Scalability environment uses the standard distributed round robin topology, add the variation for user interaction by following these steps:

- 1 Identify which NA cores should run device tasks. No additional configuration is needed on these NA cores.
- 2 On *each* NA core that should not run device tasks, reserve that NA core for user interaction.
 - a Log on to the NA console as an NA administrator.
 - b Open the Administrative Settings – Server page (**Admin > Administrative Settings > Server**), and then scroll to the bottom of the page.
 - c Select the **Reserve this core for user interaction** check box.



If this option is not available, see the additional information for fix QCCR1D95740 in the NA 9.11 (or later) patch `readme.txt` file.



NA does not prevent you from reserving all NA cores for user interaction. Manually confirm that at least one NA core is available to run device tasks by verifying that the **Reserve this core for user interaction** check box is *not* selected for at least one NA core.

- d Click **Save**.



When the NA Horizontal Scalability environment uses a distributed round robin topology, NA ignores the setting of the **Allow a core or cores in the mesh to run all tasks created on that core locally** check box.

Configure the Core Binding Variation for User Interaction

If the NA Horizontal Scalability environment uses the standard core binding topology, add the variation for user interaction by ensuring that no partitions are assigned to any NA core that is designated for user interaction only. Also ensure that each partition for managing devices is assigned to an NA core.



By default, the NA core failover process does not consider the presence or absence of partitions on a running NA core. Therefore, an NA core that was previously designated for user interaction only might be running device tasks after failover. For information about configuring the NA core failover sequence, see [Configuring Failover Order \(Core Binding Only\)](#) on page 42.



When the NA Horizontal Scalability environment uses a core binding topology, NA ignores the setting of the **Reserve this core for user interaction** check box.

Configure the Core Binding Variation for Local Device Tasks

If the NA Horizontal Scalability environment uses the standard core binding topology, add the variation for running local device tasks by following these steps:

- 1 On one NA core, log on to the NA console as an NA administrator.
- 2 Open the Administrative Settings – Server page (**Admin > Administrative Settings > Server**), and then scroll to the bottom of the page.
- 3 Select the **Allow a core or cores in the mesh to run all tasks created on that core locally** check box.
- 4 Click **Save**.
- 5 Restart all NA services on all NA cores.
 - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Restart**:
 - TrueControl ManagementEngine
 - TrueControl FTP Server
 - TrueControl SWIM Server
 - TrueControl Syslog Server
 - TrueControl TFTP Server
 - *UNIX*: Run the following command:

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

- 6 On each NA core that does not run pre-planned tasks but should run user-initiated device tasks, do the following:
 - a Log on to the NA console for that NA core as an NA administrator.
 - b Open the Administrative Settings – Server page (**Admin > Administrative Settings > Server**), and then scroll to the bottom of the page.
 - c Select the **Allow this core to run all tasks created on it locally** check box.
 - d Click **Save**.



To obtain the benefit of this variation, pre-planned tasks must be created on the NA cores on which the **Allow this core to run all tasks created on it locally** check box is *not* selected.

Configuring NA Satellites in a Horizontal Scalability Environment

An NA Horizontal Scalability environment can include one or more NA satellites. An NA satellite consists of a satellite gateway installed on a remote server and a satellite agent deployed to the satellite gateway. An NA satellite communicates with an NA core through a core gateway that is installed on or near the NA server, depending on operating system compatibility.

Each core gateway can be associated with only one NA core and can communicate with multiple NA satellites. Likewise, each NA satellite can communicate with multiple core gateways. Communication between an NA satellite and a core gateway occurs across a tunnel that is specific to that NA satellite/core gateway pair. This communication includes NA core-initiated tasks and NA satellite-initiated syslog message forwarding. For a given NA satellite, each tunnel to a core gateway is configured with a different cost so the syslog messages are forwarded to only one NA core, the NA core associated with the running core gateway that has the lowest cost tunnel.

The number and placement of the NA satellites depends on the network configuration. The number of core gateways used to connect NA to the NA satellites depends on the Horizontal Scalability topology.

- *Distributed Round Robin:*

In a distributed round robin topology, one core gateway must be installed for each NA core. This model ensures that each NA core can access all managed devices.

Variation for user interaction: Because an NA core that is reserved for user interaction does not run device tasks, this NA core does not require an associated core gateway.



In a well-configured distributed round robin topology, each NA core can reach each NA satellite. No extra configuration work is needed to support NA core failover.

- *Core Binding:*

In a core binding topology, one core gateway must be installed for each NA core that requires access to one or more devices being managed by an NA satellite.

Variation for user interaction: Because an NA core that is designated for user interaction only does not run device tasks, this NA core does not require an associated core gateway.

Variation for local device tasks: Because an NA core that is configured to run local device tasks might access any managed device, this NA core requires an associated core gateway.



In a core binding topology, an NA core might not be connected to a core gateway. Such an NA core cannot communicate with any NA satellites. If an NA core with one or more NA satellite partitions were to fail over to an NA core with no core gateway, all tasks for the devices in the NA satellite partitions would fail.

In this case, alternatives for supporting NA core failover include the following:

- Install a core gateway on each NA core, and configure the core gateway with a tunnel to each NA satellite.
- Customize the NA core failover order to ensure that an NA core with one or more NA satellite partitions fails over only to an NA core with a core gateway that is connected to the NA satellites. For more information, see [Configuring Failover Order \(Core Binding Only\)](#) on page 42.

To configure NA satellites in a Horizontal Scalability environment, follow this general process:

- 1 Install a core gateway on or near one NA core as described in “Installing a Core Gateway” in the *NA Satellite User's Guide*.

Installing the first core gateway creates the Gateway Crypto Data file, which is needed for satellite gateway installation. Note the path and name of the Gateway Crypto Data file. Also note the password used to create the Gateway Crypto Data file.

- 2 As needed for the Horizontal Scalability topology, install a core gateway on or near each additional NA core in the Horizontal Scalability environment as described in “Installing a Core Gateway” in the *NA Satellite User's Guide*.
 - Assign the realm name of the first core gateway, typically Default Realm, to each additional core gateway.
 - After installation, update the `opswgw.EgressFilter` entry to match the following:

```
opswgw.EgressFilter=tcp:*:443:127.0.0.1:* , tcp:*:22:NAS: , tcp:*:23:NAS: , tcp:*:513:NAS: , tcp:*:443:NAS: , tcp:*:80:NAS:
```

- 3 Install a satellite gateway for each group of managed devices as described in “Installing a Satellite Gateway” in the *NA Satellite User's Guide*.
- 4 Connect each satellite gateway to each core gateway with which it should communicate.

For each satellite gateway, update the satellite configuration to enable communication with each core gateway in the Horizontal Scalability environment as follows:

- a Connect to the satellite gateway server as the root user.
- b Change to the following directory:


```
/etc/opt/opsware/opswgw-<gateway_name>/opswgw.properties
```
- c Back up the `opswgw.properties` file to a location outside the gateway installation directory.
- d In a text editor, such as `vi`, open the `opswgw.properties` file.

- e Configure an `opswgw.TunnelSrc` entry for each core gateway this satellite gateway should communicate with. Configure each tunnel with a different cost to specify order for forwarding syslog messages. For example:

```
opswgw.TunnelSrc=<core_gateway1_IP>:2001:100:0:/var/opt/opsware/
crypto/opswgw-RemoteGw/opswgw.pem
```

```
opswgw.TunnelSrc=<core_gateway2_IP>:2001:200:0:/var/opt/opsware/
crypto/opswgw-RemoteGw/opswgw.pem
```

- f Restart the satellite gateway:

```
/etc/init.d/opswgw-<gateway_name> restart
```

- 5 Connect each NA core to its core gateway as described in “Configuring NA to Communicate with the Core Gateway” in the *NA Satellite User's Guide*.
- 6 For each satellite gateway, deploy the satellite agent to the satellite gateway as described in “Adding a Remote Agent to a Satellite Gateway Host” in the *NA Satellite User's Guide*.
Deploy the satellite agent from one of the core gateways listed in the satellite gateway's `opswgw.TunnelSrc` entries.
- 7 On the Partitions page (**Admin > Security Partitions**) in the NA console, assign one or more partitions to each NA satellite.

You must associate each partition with on NA core. In a distributed round robin topology, NA ignores this association.

For more information, see “Partitions” in the NA help.

Configuring Dynamic Device Group Calculation in a Horizontal Scalability Environment

By default, all NA cores in the Horizontal Scalability environment calculate the members of each dynamic device group. If your environment includes dynamic device groups, for optimum system performance, disable dynamic device group calculation on all but one NA core.



If this one NA core fails, enable dynamic device group calculation on a different NA core for the duration of the failover (as described in [Ensure Continuity of Dynamic Device Group Calculation](#) on page 45). After failback, update the configuration to again designate only one NA core to do dynamic device group calculation.

To disable dynamic group calculation on an NA core, follow these steps:

- 1 Back up the `adjustable_options.rcx` file to a location outside the `<NA_HOME>` directory.
- 2 In the `adjustable_options.rcx` file, add the following line:

```
<option name="dynamic_group/disable">true</option>
```
- 3 Save the `adjustable_options.rcx` file.

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

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



Because the `adjustable_options.rcx` file is specific to the NA core, reloading the server options does not synchronize the new setting among the NA cores.

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:



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>;
```



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.

3 NA Core Failover in a Horizontal Scalability Environment

Failover is the process of moving certain responsibilities from a system that has failed to one that is operational. Failback is the process of reversing the moves after the failed system is operational again.

In an HP Network Automation Software (NA) Horizontal Scalability environment, NA core failover moves the responsibility for running device tasks from an NA core that is no longer running to one or more an NA cores that are running. (If the Horizontal Scalability environment uses a core binding topology, NA core failover also moves partitions.)

As of version 9.11, NA provides automated failover of a stopped NA core to the running NA cores in a Horizontal Scalability environment. While an NA core is stopped, NA does not send RMI messages to that NA core. Additionally, NA does not assign new tasks to that NA core. This automated NA core failover provides high availability of NA. Because the number of NA cores running tasks decreases, system throughput might decrease. Non-critical tasks can be paused to maintain system performance on the running NA cores.

By default, failover is enabled for all Horizontal Scalability environments.

This chapter contains the following topics:

- [Enabling Failover](#) on page 36
- [NA Core Failover Behavior](#) on page 37
- [Optional Actions after NA Core Failover Occurs](#) on page 44
- [Failing Back to the Original NA Core](#) on page 46

Enabling Failover

Installing NA 9.11 or the latest consolidated patch on all NA cores and completing the configuration described in the additional information for fix QCCR1D95740 in the patch `readme.txt` file enables NA failover. No additional configuration is needed.

NA core failover is not appropriate for all NA Horizontal Scalability environments. For example, in a core binding topology in which all NA cores cannot access all devices, you might want to disable NA core failover. For more information, see the following procedures:

- [Disabling Failover](#) on page 36
- [Re-Enabling Failover](#) on page 36

Disabling Failover

To disable NA core failover, follow these steps on only one NA core:

- 1 Back up the `distributed.rcx` file to a location outside the `<NA_HOME>` directory.
- 2 In a text editor, such as WordPad or `vi`, open the `distributed.rcx` file, and then add the following line:

```
<option name="distributed/enable_auto_failover">false</option>
```

- 3 Save the `distributed.rcx` file.
- 4 Reload the `.rcx` settings by running the `reload server options` command from the NA proxy.

NA synchronizes the change to the other NA cores in the Horizontal Scalability environment.

Re-Enabling Failover

To r-enable NA core failover, follow these steps on only one NA core:

- 1 Back up the `distributed.rcx` file to a location outside the `<NA_HOME>` directory.
- 2 In a text editor, such as WordPad or `vi`, open the `distributed.rcx` file, and then locate the following line:

```
<option name="distributed/enable_auto_failover">false</option>
```

- 3 Edit this line to set the option to true:
- 4 Save the `distributed.rcx` file.
- 5 Reload the `.rcx` settings by running the `reload server options` command from the NA proxy.

NA synchronizes the change to the other NA cores in the Horizontal Scalability environment.

NA Core Failover Behavior

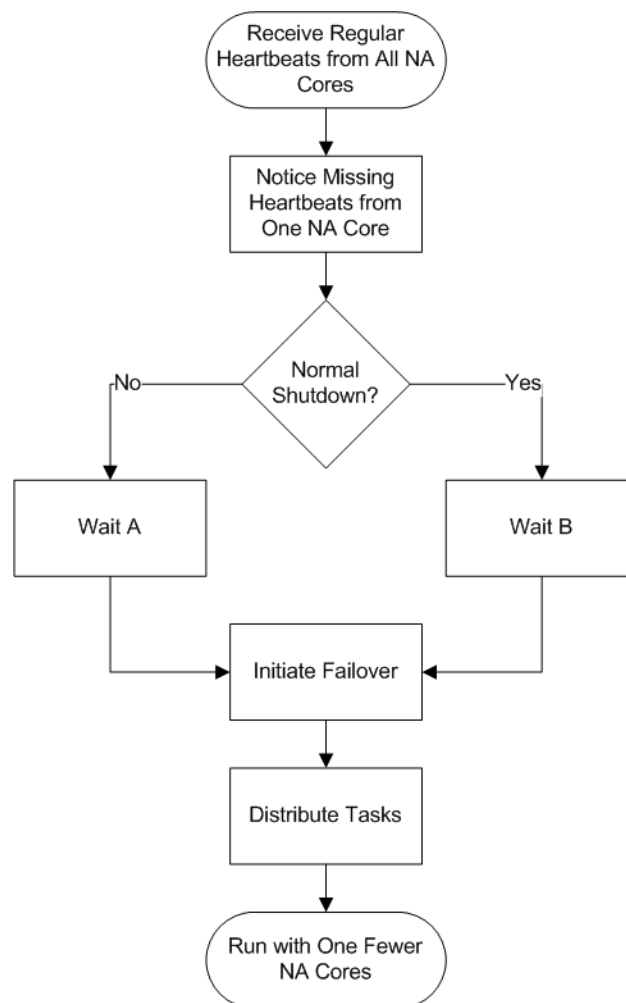
After an NA core stops, a running NA core notices the stopped NA core and waits a for communications to the stopped NA core to cease before initiating NA core failover. NA core failover moves the tasks from the stopped NA core to one or more running NA cores, depending on the Horizontal Scalability topology. The failover process sets the status of tasks that were on the stopped NA core depending on the current task status.

Figure 6 on page 37 shows the flow of the NA core failover process.

This section describes the NA core failover process in more detail. It includes the following aspects of NA core failover:

- [Delay Before Initiating Failover](#) on page 38
- [Task Status](#) on page 38
- [Movement of Tasks](#) on page 39
- [NA Failover Events](#) on page 42
- [Configuring Failover Order \(Core Binding Only\)](#) on page 42

Figure 6 NA Core Failover Conceptual Flowchart



Delay Before Initiating Failover

After an NA core stops, a running NA core notices the stopped NA core and waits for communications to the stopped NA core to cease before initiating NA core failover. The length of time between when an NA core stops and when another NA core initiates NA core failover depends on the situation:

- If the NA core shutdown is unintentional (for example, the NA server loses power), NA waits 10 minutes before initiating the NA core failover process. This time is represented by **Wait A** in [Figure 6](#) on page 37.

In this case, NA initiates the failover process and then creates the following event: Distributed System — Abnormal Shutdown of Core.

- If the NA core shutdown is intentional (for example, an NA administrator stops the NA services), NA waits 30 minutes before initiating the NA core failover process. This time is represented by **Wait B** in [Figure 6](#) on page 37. This longer wait time provides a maintenance window with the expectation that the NA services might restart soon after being stopped.

In this case, during the shutdown process, the stopped NA core creates the following event: Distributed System — Normal Shutdown of Core. Another NA core initiates the failover process and then creates the following event: Distributed System — Processed Normal Shutdown of Core.

Task Status

The failover process sets the status of tasks that were on the stopped NA core depending on the current task status.

- NA processes the tasks that were running on the stopped NA core as follows:
 - If any valid retries remain for a task, move that task to a running core (as described in [Movement of Tasks](#) on page 39), and then set the task status to PENDING.
 - If no valid retries remain for a task, keep that task associated with the stopped NA core and set the task status to FAILED.

An NA administrator can review the failed tasks to determine which of these tasks should be re-run.



Group tasks that were running on the stopped NA core remain associated with the stopped NA core. NA tracks all associated child tasks and updates the status of each group task in the NA database after the child tasks complete.

- NA moves the tasks that were pending or waiting on the stopped NA core to the running NA cores in the Horizontal Scalability environment (as described in [Movement of Tasks](#) on page 39). NA sets the task status for each of these tasks to PENDING.

Movement of Tasks

NA core failover moves the tasks from the stopped NA core to one or more running NA cores, depending on the Horizontal Scalability topology.

- In a distributed round robin topology, the NA core failover process distributes the tasks from the stopped NA core to *all* running NA cores that accept device tasks as shown in [Figure 7](#) on page 40.

An NA core that is reserved for user interaction never receives tasks from the stopped NA core.



In a distributed round robin topology, failover from a stopped NA core cannot occur in the following situations:

- No NA core is running.
- All running NA cores are reserved for user interaction.

- In a core binding topology, the NA core failover process moves all partitions (and all associated device tasks) from the stopped NA core to *one* running NA core that accepts device tasks as shown in [Figure 8](#) on page 41.

An NA core that is reserved for running local device tasks never receives tasks from the stopped NA core. An NA core that is designated for user interaction only might receive tasks from the stopped NA core. For information about which NA core receives tasks from a stopped NA core, see [Configuring Failover Order \(Core Binding Only\)](#) on page 42.









In a core binding topology, failover from a stopped NA core cannot occur in the following situations:


- No NA core is running.
- All running NA cores are reserved for running local device tasks.
- The NA cores specified in the `distributed.rcx` file as the failover targets are not running.

Figure 7 Standard Distributed Round Robin Task Failover

Legend

-  Running NA core
-  Stopped NA core
-  Running task
-  Pending task
-  Waiting task
-  Failed task

Distributed Round Robin Normal Operation



Distributed Round Robin After Failover of Core 3








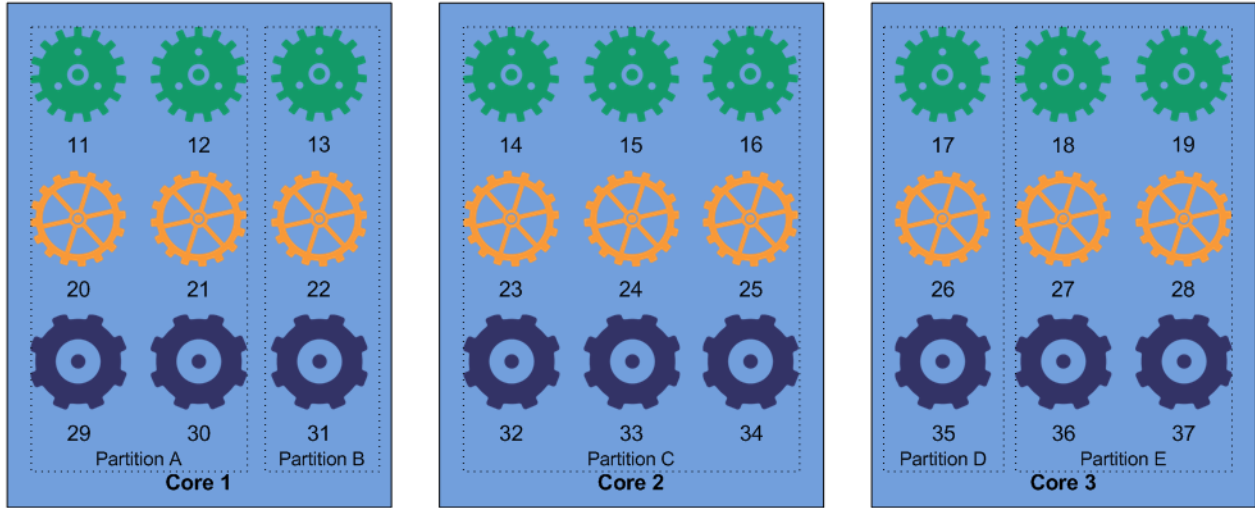


Figure 8 Standard Core Binding Task Failover

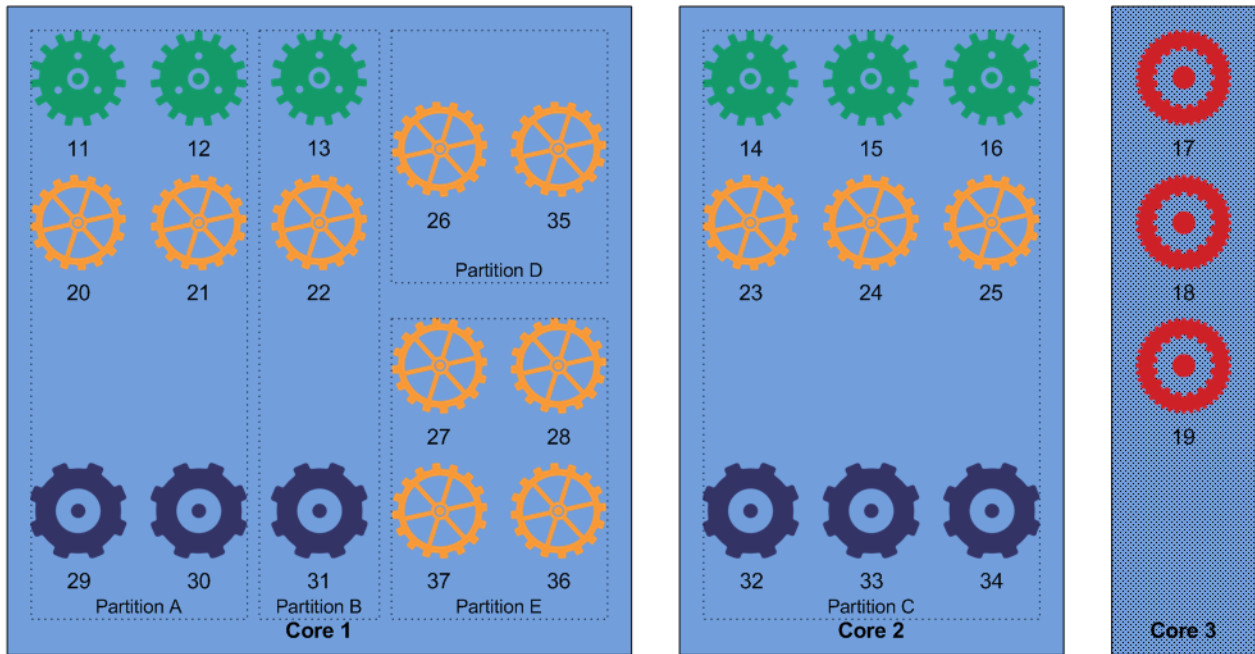
Legend

-  Running NA core
-  Stopped NA core
-  Running task
-  Pending task
-  Waiting task
-  Failed task

Core Binding Robin Normal Operation



Core Binding After Failover of Core 3



NA Failover Events

NA 9.11 adds three events related to NA core failover. These events are only available in Horizontal Scalability environments.

- Distributed System — Abnormal Shutdown of Core

A running NA core detected and responded to the unintentional shutdown of another NA core in the Horizontal Scalability environment.

- Distributed System — Normal Shutdown of Core

An NA core communicated its transition before shutting down.

- Distributed System — Processed Normal Shutdown of Core

A running NA core responded to the normal shutdown of another NA core in the Horizontal Scalability environment.

NA users can search for these events. NA users can also configure event notification rules based on these events.

There are no failback events.

Configuring Failover Order (Core Binding Only)

In a core binding topology, by default each NA core fails over to the next running NA core in the sequence. For example, in a three NA core Horizontal Scalability environment, the default failover order is as follows:

- Core 1 fails over to Core 2. If Core 2 is not running, Core 1 fails over to Core 3.
- Core 2 fails over to Core 3. If Core 3 is not running, Core 2 fails over to Core 1.
- Core 3 fails over to Core 1. If Core 1 is not running, Core 3 fails over to Core 2.

By default, the NA core failover process does not consider the presence or absence of partitions on a running NA core. Therefore, an NA core that was previously designated for user interaction only might be running device tasks after failover.

Reasons to customize the NA core failover order include:

- Some managed devices are not accessible from one or more NA cores. In this case, set each NA core to fail over to another NA core that can reach the same devices.
- No partitions are assigned to an NA core that is designated for user interaction only. To maintain this designation after NA core failover, set each NA core with partitions to fail over to another NA core with partitions.

You can change the failover order, and you can specify multiple NA cores as the failover target in case the primary target is also stopped.

To customize the default failover order, follow these steps:

- 1 On one NA core, modify the `distributed.rcx` file.
 - a Back up the `distributed.rcx` file to a location outside the `<NA_HOME>` directory.
 - b In a text editor, such as WordPad or vi, open the `distributed.rcx` file, add the following lines:

```
<option name="distributed/bind_to_core_override_1">2</option>
<option name="distributed/bind_to_core_override_2">3</option>
<option name="distributed/bind_to_core_override_3">4</option>
<option name="distributed/bind_to_core_override_4">5</option>
<option name="distributed/bind_to_core_override_5">1</option>
```

- c Delete any unneeded lines. For example, in a four NA core Horizontal Scalability environment, delete the line containing the `distributed/bind_to_core_override_5` option.
- d Customize the value of each line by replacing the default NA core number with the desired NA core number.

To specify multiple backup NA cores for one NA core, use a comma (,) to separate each additional NA core number in the order that each NA core should be tried. For example, to specify that Core 1 should fail over to Core 4 (if it is running) or to Core 2 (if Core 4 is stopped), use the following line:

```
<option name="distributed/bind_to_core_override_1">4,2</option>
```



The NA core number is the numeric portion of the NA core name on the List Cores page (**Admin > Distributed > Core List**).

- e Save the `distributed.rcx` file.
- 2 Copy the modified `distributed.rcx` file to all other NA cores in the Horizontal Scalability environment.
 - 3 Restart the NA services on all NA cores in the Horizontal Scalability environment.
 - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Restart**:
 - **TrueControl ManagementEngine**
 - **TrueControl FTP Server**
 - **TrueControl SWIM Server**
 - **TrueControl Syslog Server**
 - **TrueControl TFTP Server**
 - *UNIX*: Run the following command:


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

Optional Actions after NA Core Failover Occurs

After NA core failover occurs, consider the following optional actions:

- [Notify Users](#) on page 44
- [Configure Integrations to a Different NA Core](#) on page 44
- [Pause Non-Critical Tasks](#) on page 44
- [Locate Tasks Impacted by NA Core Failover](#) on page 44
- [Ensure Continuity of Dynamic Device Group Calculation](#) on page 45
- [Stop the Core Gateway for the Stopped NA Core](#) on page 45

Notify Users

If necessary, inform NA users that the stopped NA core is unavailable and identify a different NA core for user interaction.

Configure Integrations to a Different NA Core

If the stopped NA core is expected to be unavailable for a significant time window (as defined by the needs of your NA environment), configure any integrations that connect to that NA core to connect to a different NA core. For configuration information, see the documentation for each integration.

Pause Non-Critical Tasks

NA core failure reduces the number of NA cores running tasks. If system throughput decreases after failover, pause non-critical tasks scheduled to start during the expected duration of the NA core outage. For example:

- 1 On the Search for Task page (**Reports > Search For > Tasks**), for the Schedule Date field, set since to **Now** and until to **4 hours later**.
- 2 On the Task Search Results page, pause each listed task.

Locate Tasks Impacted by NA Core Failover

As described in [Task Status](#) on page 38, the NA core failover process sets the status of some tasks that were running on the stopped NA core to **FAILED**. You can examine the failed tasks and determine how to handle each one. To identify tasks that failed because of the NA core failover, follow these steps:

- 1 In the NA console, on the Search for Tasks page (**Reports > Search For > Tasks**), do the following:
 - Set the Start Date field to a time range near the time that the NA core stopped.
 - Set the Task Status field to **Failed**.
 - Set the Failure Type field to **Core Down**.
 - Set the Core field to the name of the NA core that stopped.

- 2 On the Task Search Results page, examine the results.



The search results might include tasks that were cancelled because the NA server rebooted.

- 3 As appropriate, re-run the identified tasks.

Ensure Continuity of Dynamic Device Group Calculation

In an NA Horizontal Scalability environment that includes dynamic groups, if only one NA core performs dynamic device group membership calculations, consider the effect of that NA core shutting down. If that NA core is expected to be unavailable for a significant time window (as defined by the needs of your NA environment), configure a different NA core to perform dynamic device group membership calculations.



After failback, update the configuration to again designate only one NA core to do dynamic device group calculation.

To re-enable dynamic device group calculation on an NA core, follow these steps:

- 1 Back up the `adjustable_options.rcx` file to a location outside the `<NA_HOME>` directory.
- 2 In the `adjustable_options.rcx` file, locate the following line:


```
<option name="dynamic_group/disable">true</option>
```
- 3 Change the located line to read:


```
<option name="dynamic_group/disable">false</option>
```
- 4 Save the `adjustable_options.rcx` file.
- 5 Restart the NA services.
 - *Windows*: Open the **Services** control panel. In the list of services, right-click **TrueControl ManagementEngine**, and then click **Restart**.
 - *UNIX*: Run the following command:

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



Because the `adjustable_options.rcx` file is specific to the NA core, reloading the server options does not synchronize the new setting among the NA cores.

Stop the Core Gateway for the Stopped NA Core

If the core gateway associated with the stopped NA core is still running, NA satellites might continue to forward syslog messages to that core gateway. Because the NA core is down, these syslog messages are lost. To prevent this loss, stop the core gateway by running the following command:

```
/etc/init.d/<coreGatewayName> stop
```

Any NA satellites that were forwarding syslog messages to the stopped core gateway now forward their syslog messages to the core gateway with the next higher cost (the current lowest cost) as configured by the `opswgw.TunnelSrc` entries in the NA satellite `opswgw.properties` file.

Failing Back to the Original NA Core

NA core failback is the process of restoring an NA core to the state it was in before it failed. The failback process differs according to Horizontal Scalability topology.

- In a distributed round-robin topology, the failback process is transparent. NA assigns tasks to the restored NA core during regular task distribution operations. Any tasks that were moved during failover remain on those NA cores.
- In a core-binding topology, the failback process includes reassigning sites to the restored NA core. This action is manual because NA does not track site moves among NA cores.

In the NA console, open the Site Reassignment page (**Admin > Distributed > Site Reassignment**), and then assign the appropriate partitions to the restored NA core.

4 System Administration

This chapter contains the following topics:

- [NA-Generated Events for Horizontal Scalability](#) on page 47
- [Using the NA Distributed System Pages](#) on page 48

NA-Generated Events for Horizontal Scalability

The following NA-generated system events relate to the health of the Horizontal Scalability functionality:

- [Distributed System – Time Synchronization Warning](#) on page 47
- [Distributed System – RMI Error](#) on page 47

Distributed System – Time Synchronization Warning

The time synchronization warning occurs when the NA server timestamps and the NA database server timestamps are not the same.

To address this situation, update the configuration of the NA servers and the NA database server to use the same time and time zone. It is recommended to synchronize the servers with an external time service.

Distributed System – RMI Error

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. This event typically occurs when there are network problems between the NA servers. To troubleshoot this problem, follow these steps:

- 1 Verify that the host that the server cannot connect to is running.
- 2 Verify that the NA instance on that host is running.
- 3 From a command line, enter `ping <host>` to ensure that network connectivity exists between the servers.
- 4 To verify that RMI connections are being accepted, at a command prompt, enter `telnet <host> to port 1099` (or whatever the RMI listen port is set to). The expected response is some data that include the text string `java.rmi.MarshalledObject`.

Failure at 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 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 server.
 - Replace `<NA.example.com>` with the fully-qualified domain name of the NA server.
 - Replace `<NA>` with the short hostname of the NA server.
- 4 Repeat [step 3](#) until the `hosts` file includes all NA servers in the distributed system environment.
- 5 To use the updated hosts information, restart the NA server.

Using the NA Distributed System Pages

Configuring NA for Horizontal Scalability enables the **Admin > Distributed** menu in the NA console. An NA administrator can use these pages to monitor and administer the NA environment. This section describes the pages that are available from the **Admin > Distributed** menu.

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 then click Monitor Results.

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

- **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. See [Edit Core Page](#) on page 50 for information.



When this condition occurs, the error logs will contain the following text:
“Fatal error - could not assign local core.”

The CoreHostname value can be either the DNS name, the `/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 value A.B.C.D should be replaced with the appropriate NAT IP address for the NA server and should match the CoreHostname value in the RN_CORE table for that NA Core.

- **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 Horizontal Scalability environment to transfer information about scheduled tasks, system settings, software images, and so on.

Distributed Error List

The Distributed Error List page does not apply to Horizontal Scalability environments.

Distributed Conflict List

The Distributed Conflict List page does not apply to Horizontal Scalability environments.

Site Reassignment Page

Use the Site Reassignment page to change the site-to-NA core mapping. This function is useful for failover of sites from one NA core to another and for restoring sites back to their original NA core.



As of NA 9.11, moving sites to an NA core also moves the pending and waiting tasks associated with that site to the receiving NA core. Running tasks run to completion.

To move sites from one NA core to another NA core:

- 1 Open the Site Reassignment page (**Admin > Distributed > Site Reassignment**).
- 2 In the Assign all sites managed by line, do the following:
 - a In the first list, select the core that currently owns the sites to be moved.
 - b In the second list, select the core to receive the sites.
 - c Click **Save**.

To reset the managing core for sites owned by one NA core to be the owning NA core:

- 1 Open the Site Reassignment page (**Admin > Distributed > Site Reassignment**).
- 2 In the Reset all sites owned by line, do the following:
 - a In the list, select the core that currently owns the sites.
 - b Click **Save**.

Distributed Core List Page

The Distributed Core List page lists all NA Cores in the NA Horizontal Scalability environment.

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

Table 3 List Cores Page Fields

Field	Description
Name	The name of the NA core, which includes the NA core number.
Core Hostname	The hostname of the NA server.
Status	The status of the NA core. Possible status value are: <ul style="list-style-type: none"> • Running: Fully functional—The NA core is running normally. This NA core runs user-initiated and pre-planned tasks. • Running: User interaction only, no tasks—The NA core is running normally. This NA core runs only user-initiated tasks; it is not available to run pre-planned tasks. (This status occurs only for the distributed round robin topology variation for user interaction.) • Stopped: Normal shutdown—The NA core shut down in an orderly fashion. • Stopped: Abnormal shutdown—The NA core shut down unintentionally.
Timezone Offset	The time zone offset of the NA server.
Realm	The default Realm for the NA core.
Actions	The available action is: <ul style="list-style-type: none"> • Edit — Open the Edit Core page. See Edit Core Page on page 50.

Edit Core Page

Use the Edit Core page to edit the NA core definition.

To edit the definition of an NA core:

- 1 Navigate to the Edit Core page.
 - a Open the List Cores page (**Admin > Distributed > Core List**).
 - b In the Actions column of the List Cores page, click **Edit**.
- 2 Edit the field values, and then click **Save Core**.

Table 4 Edit Core Page Fields

Field	Description
Name	The name of the NA core.
Database Identifier	The Oracle SID, the service name of the Oracle RAC cluster, or the SQL Server database name of the NA database.
Core Hostname	The hostname of the NA server.
RMI Port	The port on the NA server used for communication among the NA cores in the NA Horizontal Scalability environment.
Database Hostname	The hostname of the NA database server.
Database Port	The port on the NA database server used for communication with the NA cores.
Timezone Offset	The time zone offset of the NA server.
Replication Admin User	Does not apply to Horizontal Scalability environments.
Replication Password Confirm Replication Password	Does not apply to Horizontal Scalability environments.
Comments	Text comments about the NA core.
Realm Name	The default Realm for the NA core. For information about segmenting devices, see the <i>NA User's Guide</i> .

Device Password Rule Priority Reset Page

The Device Password Rule Priority Reset page does not apply to Horizontal Scalability environments.

Renew Configuration Options Page

Use the Renew Configuration Options page to reset the configuration options when the configuration options on an NA Core become unsynchronized from other NA servers in the NA Horizontal Scalability environment.

To resynchronize the configuration across all NA cores in the NA Horizontal Scalability environment:

- 1 Navigate to the Renew Configuration Options page (**Admin > Distributed > Renew Configuration Options**).
- 2 Click **Renew Config Options**.

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 9.11

Document title: *NA Horizontal Scalability User's Guide, February 2013*

Feedback: