HPE Network Automation Software

Software Version: 10.21
Windows® and Linux® operating systems

Horizontal Scalability Guide

Document Release Date: March 2018
Software Release Date: December 2016

Legal Notices

Warranty

Restricted Rights Legend

Copyright Notice

Trademark Notices

Oracle Technology – Notice of Restricted Rights

Acknowledgements

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

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

Documentation Updates

The title page of this document 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: https://softwaresupport.hpe.com/group/softwaresupport/search-result?keyword=.

This site requires an HP Passport account. If you do not have one, click the **Create an account** button on the HP Passport Sign in page.

Support

Visit the HPE Software Support web site at: **https://softwaresupport.hpe.com**

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

HPE Software Support 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 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 HPE 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 register for an HP Passport ID, go to **https://softwaresupport.hpe.com** and click **Register**.

To find more information about access levels, go to:
https://softwaresupport.hpe.com/web/softwaresupport/access-levels

Contents

Chapter 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 active NA cores (up to seven cores if you apply the `NA_00050` patch) can connect to one NA database. For scale information, see “Maximum Supported Managed Environment” in the *NA Support Matrix*.

An NA Horizontal Scalability environment can include active and inactive NA cores.

- An active NA core is a fully-functional instance of NA.
- While an NA core is inactive, the following conditions apply:
  - That NA core does not run tasks.
  - Users should *not* log on to the NA console.
  - Users might connect to the NA command-line interface through telnet or SSH *only* for the purpose of setting the NA core state during switchover or switchback for disaster recovery.

NA Horizontal Scalability is the base for the NA disaster recovery architecture. For information about implementing NA disaster recovery, see the *NA Disaster Recovery Configuration Guide*.

NA Horizontal Scalability also provides for high availability of NA with automated failover when an NA core becomes unavailable. For more information see "NA Core Failover in a Horizontal Scalability Environment" on page 50.

Note: 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:

https://softwaresupport.hpe.com/group/softwaresupport

For more information, see "Documentation Updates" on page 2.

This chapter includes the following topics:

- "Horizontal Scalability Architecture" on the next page
- "Horizontal Scalability Topologies" on the next page
- "Horizontal Scalability over a WAN" on page 18

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 some of NA configuration.

The maximum number of active NA cores in a Horizontal Scalability environment is five (seven if you apply the `NA_00050` patch).

> **Caution:** In a disaster recovery configuration, the maximum number of active and inactive NA cores is nine. If you want to support more than nine cores in a disaster recovery configuration, contact HPE Software Support.

It is recommended to implement Horizontal Scalability within a data center (on a LAN or a high-speed network).

Horizontal Scalability can be implemented with the NA cores and the database server separated by wide-area network (WAN) links. This implementation is not recommended because it comes with costs and limitations. For more information, see "Horizontal Scalability over a WAN" on page 18.

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 the next page
- "Core Binding" on page 13

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.

As of NA 10.10, distributed round robin 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 34.

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. 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.

> **Note:** 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.

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.

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.

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

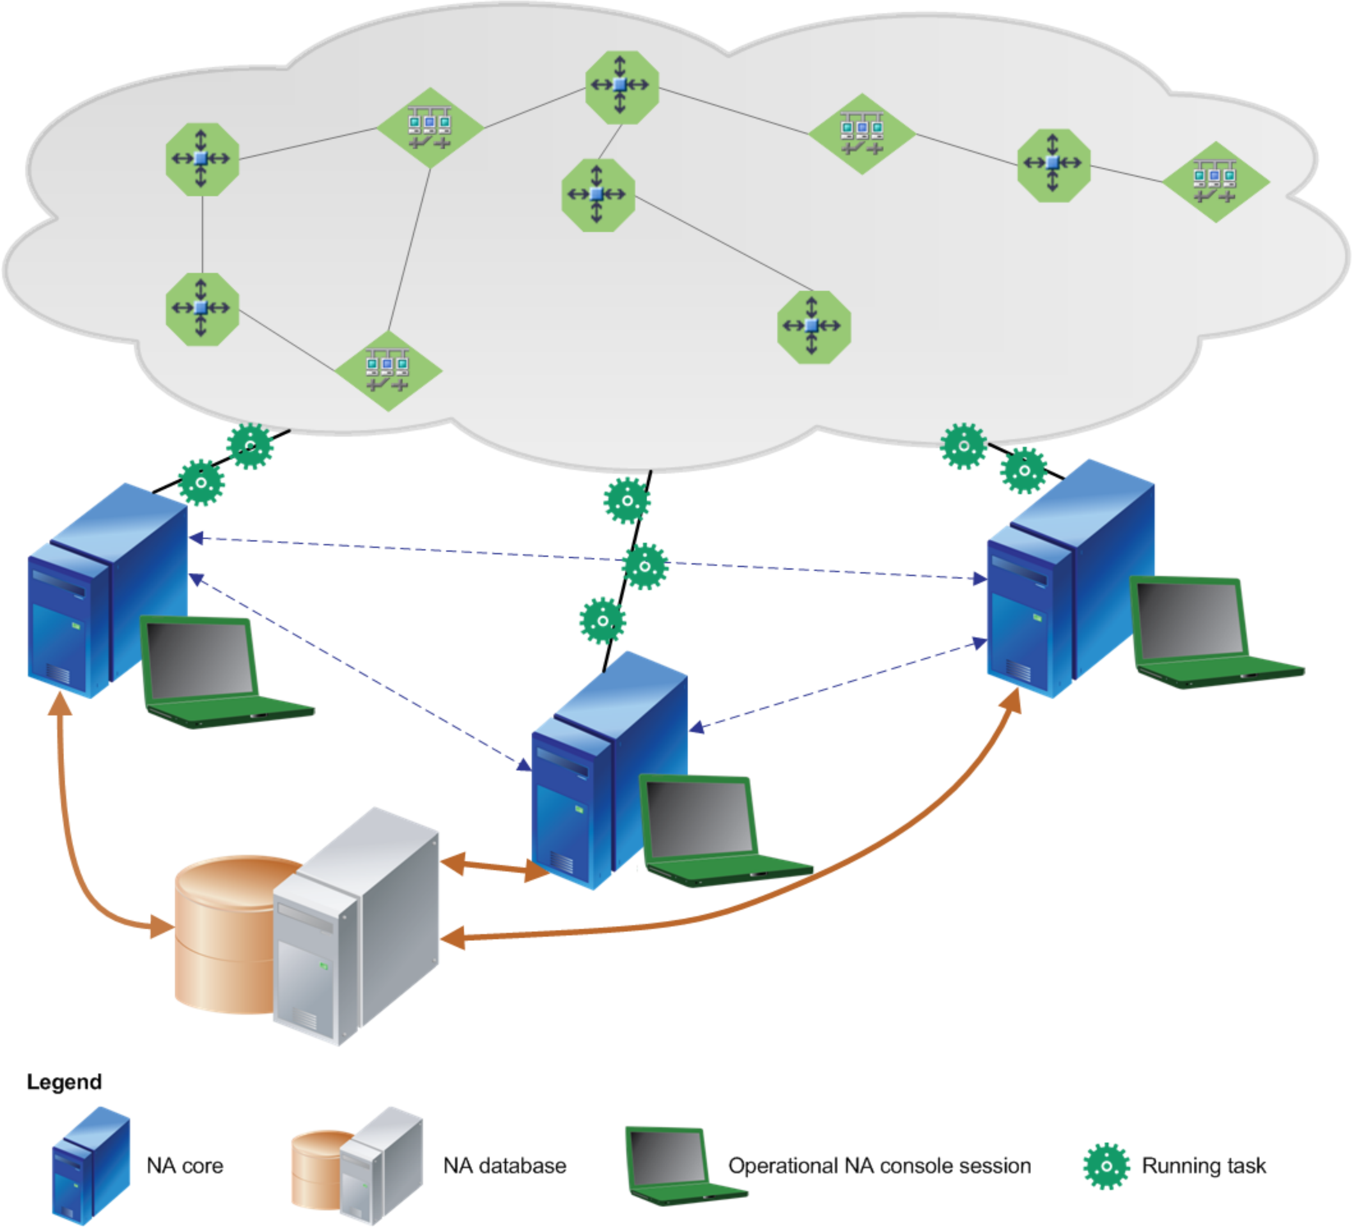

Figure 2: Example Architecture: Distributed Round Robin Variation for User Interaction

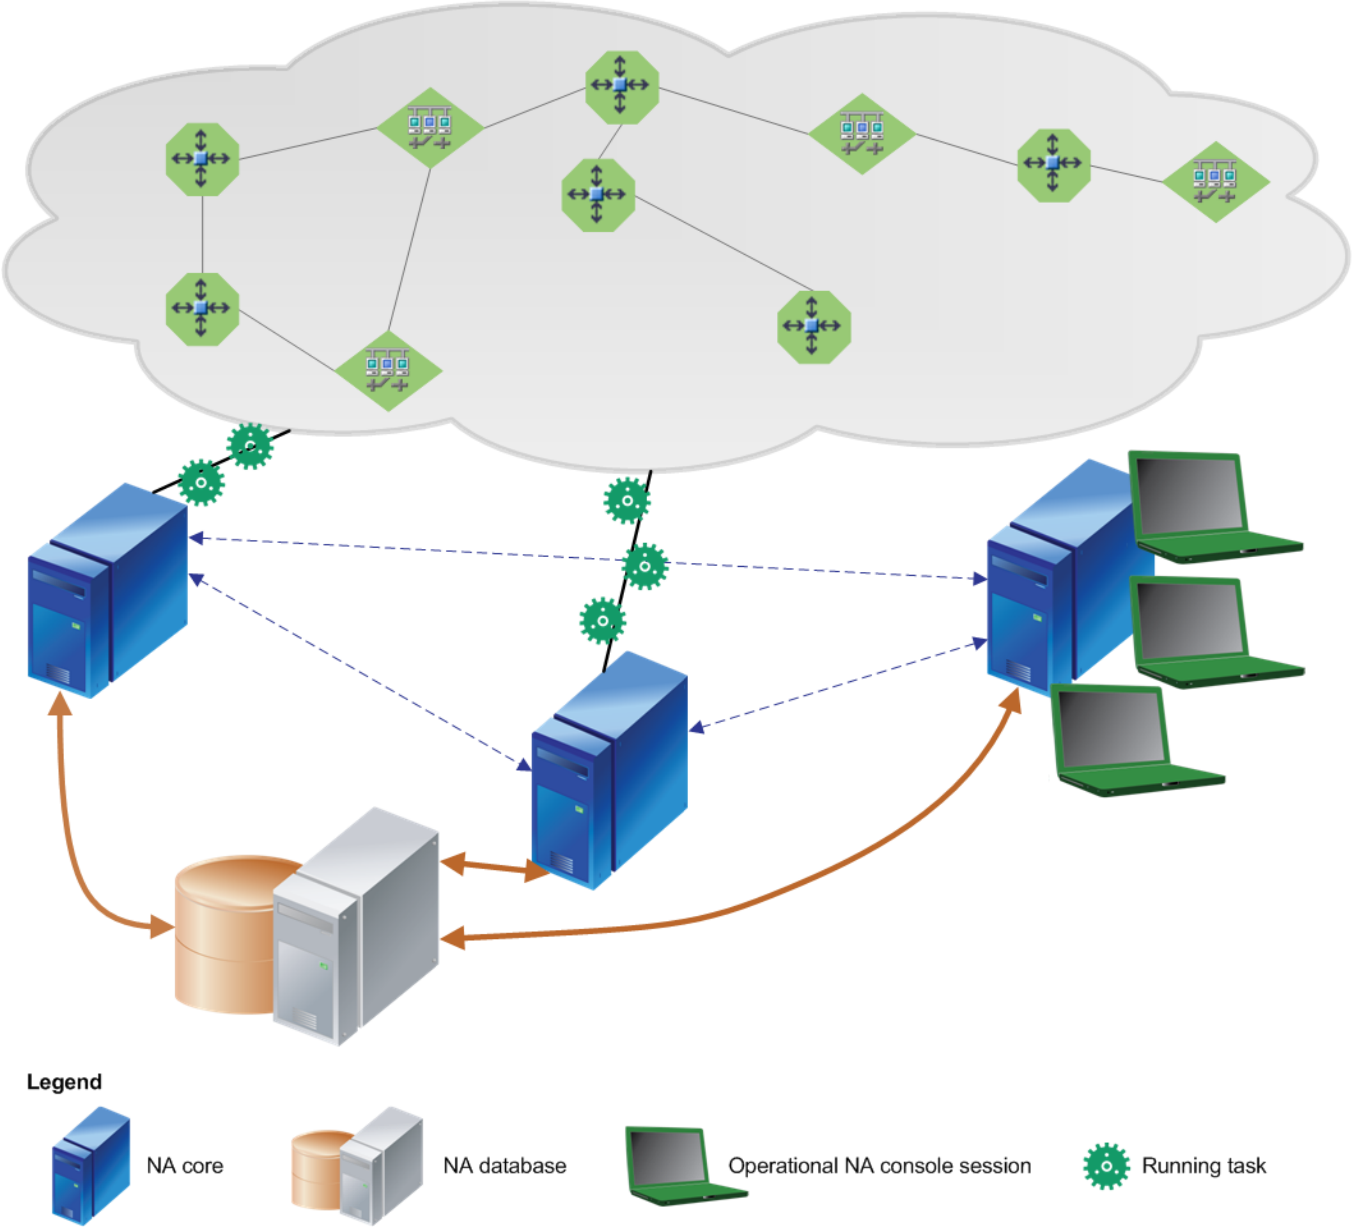

Core Binding

In the core binding topology, each NA core communicates with a fixed, non-overlapping set of devices, as shown in Figure 3. 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.

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.

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.

- 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.

  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.

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

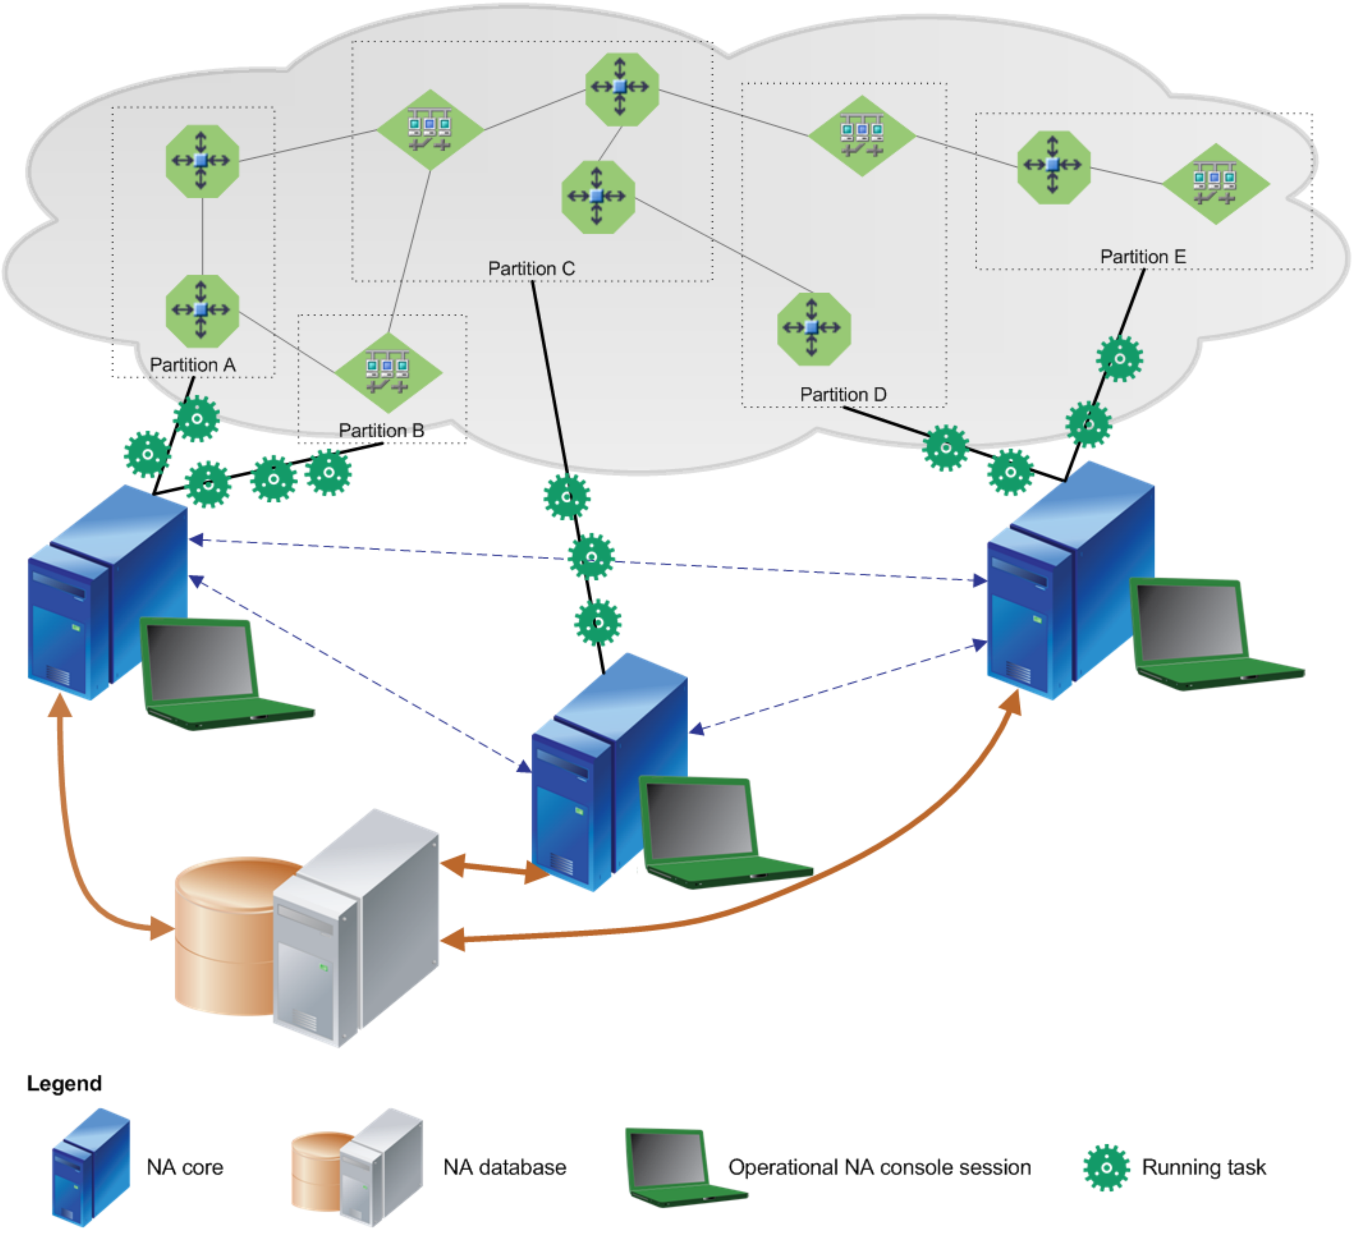

Figure 4 Example Architecture: Core Binding Variation for User Interaction

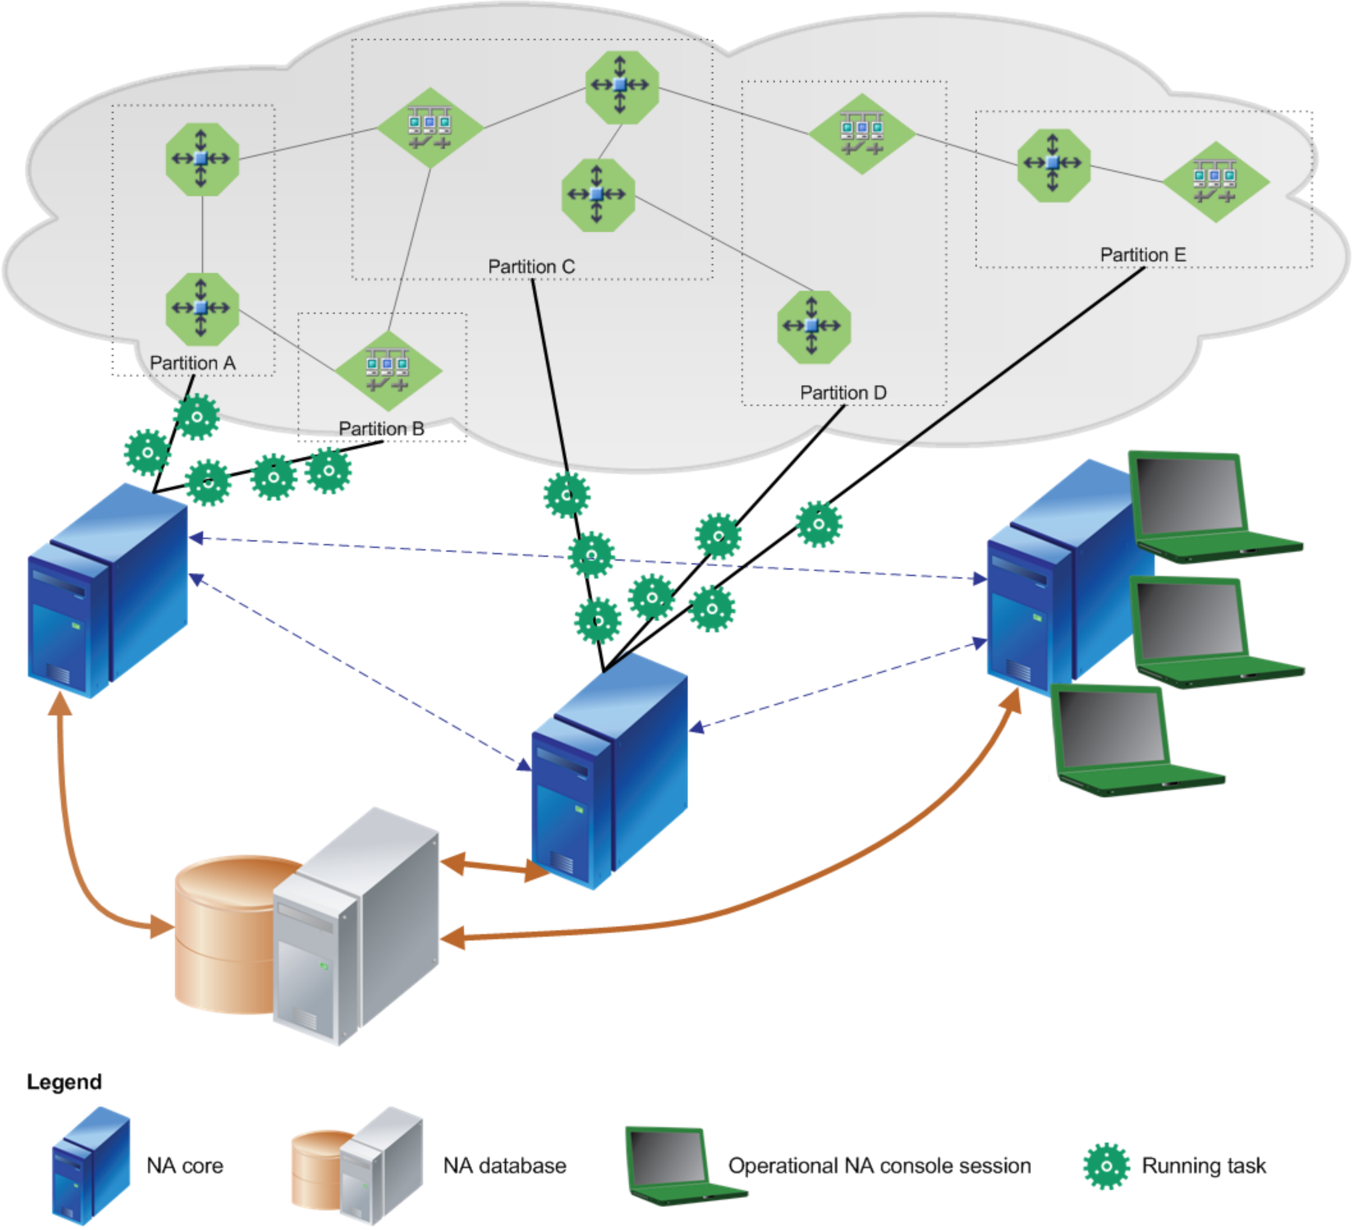

Figure 5 Example Architecture: Core Binding Variation for NA Local Device Tasks

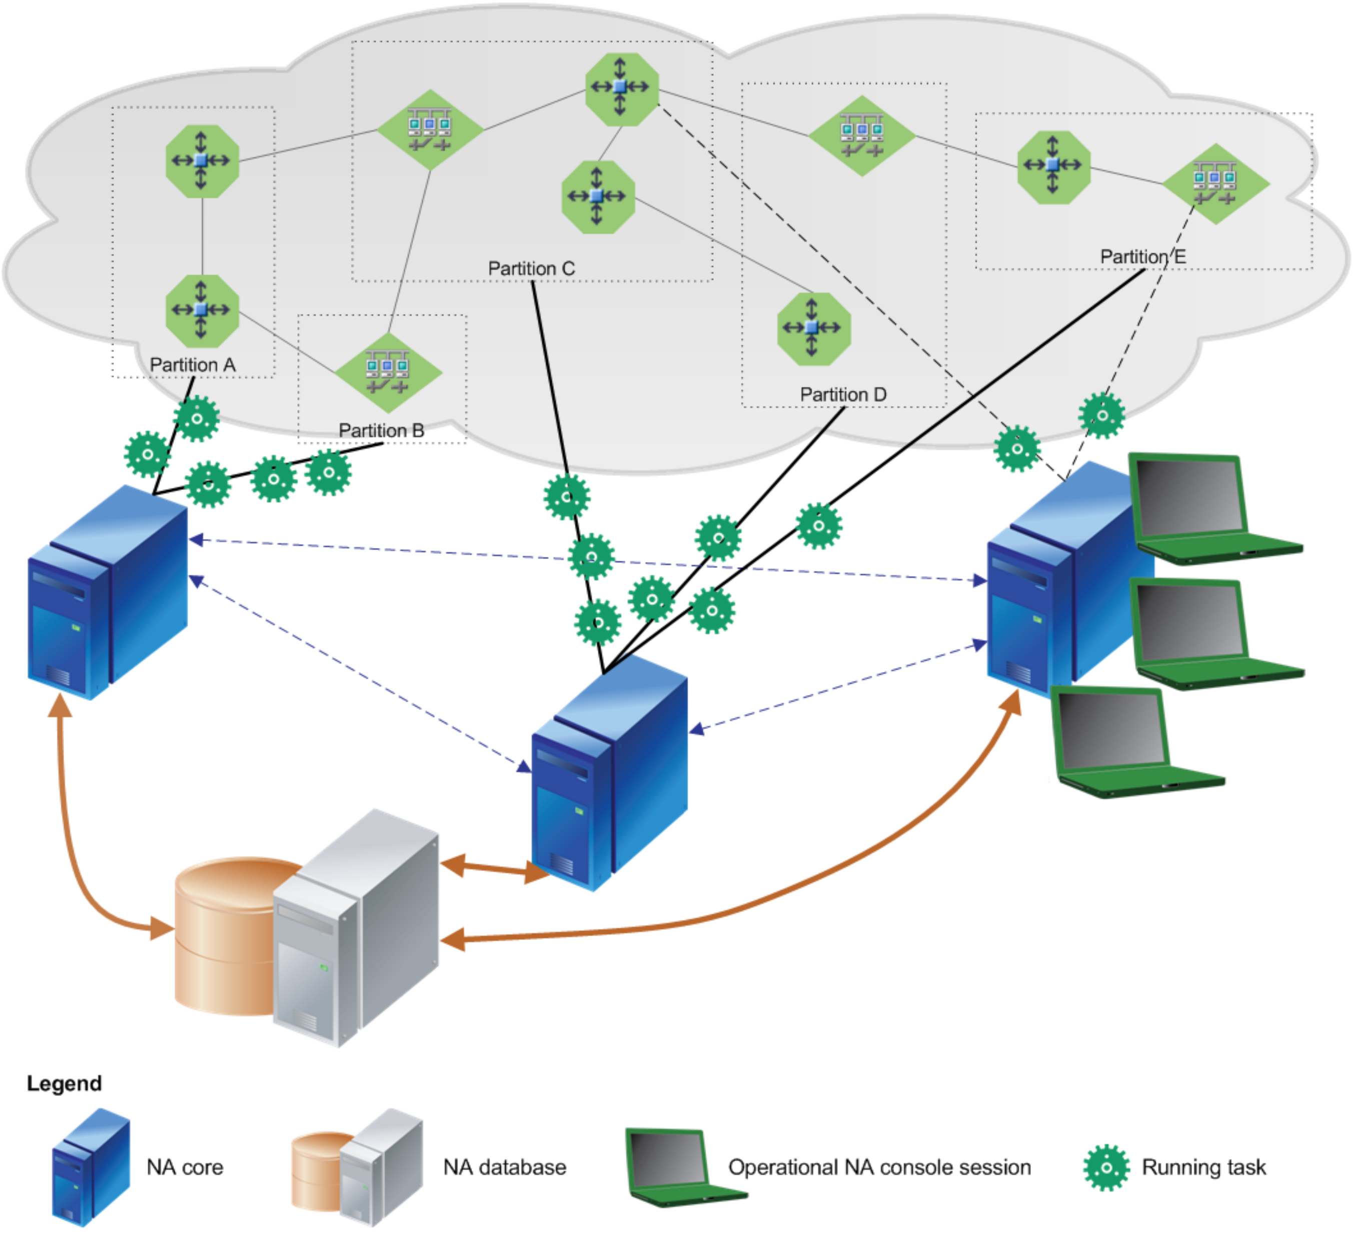

Horizontal Scalability over a WAN

NA Horizontal Scalability can be implemented over a WAN. In this implementation, one or more NA cores are connected to the NA database and the other NA cores by WAN links. See Figure 6.

> **Note:** It is recommended to implement Horizontal Scalability within a data center (on a LAN or a high-speed network). Implementing Horizontal Scalability over a WAN impacts NA performance and increases network management costs.

The advantages of Horizontal Scalability over a WAN include the following:

- Local management of some devices in remote locations.
- Local management of NA by network engineers in remote locations.
- Local access to the NA console, the NA proxy, and the NA API in remote locations.
- In disaster recovery configuration environments, less idle equipment. NA cores in the disaster recovery location can be actively managing the network before a disaster occurs.

The disadvantages of Horizontal Scalability over a WAN include the following:

- Only one topology option. The core binding topology provides for configuring which NA core runs tasks on which devices. The core binding variations for user interaction and local device tasks are available.
- Slower NA performance. Tasks run by a remote NA core could require significantly more time than for the same task run by a local NA core. Task throughput on a remote NA core is lower than that on a local NA core.
- WAN bandwidth consumption. Communications between a remote NA core and the NA database, happen over the WAN. These communications compete with all other WAN traffic.
- Configuration synchronization problems. NA does not have a mechanism for tracking configuration synchronizations that failed. If the WAN link is not highly reliable, the remote NA cores could be in a different configuration state than that of the local NA cores.

The following requirements represent the minimum network environment in which Horizontal Scalability over a WAN exhibits acceptable performance levels:

- For a WAN between a remote NA core and the NA database that also carries traffic not related to NA:
  - WAN bandwidth of 1 Gb/s or higher
  - Latency across the WAN link of 30 ms or lower
- For a dedicated WAN between a remote NA core and the NA database:
  - WAN bandwidth of 500 Mb/s or higher

When implementing Horizontal Scalability over a WAN, the following additional configuration changes are recommended:

- Use partitions to assign a smaller number of devices to a remote NA core. This configuration manages task throughput.
- Reduce the value of the maximum concurrent tasks on a remote NA core to increase system performance. Experiment to determine the best value. This configuration reduces concurrent bandwidth consumption.
- Limit the users of a remote NA console to those people local to that NA core. This choice improves NA console performance.

Figure 6 Example Deployment: Horizontal Scalability over a WAN

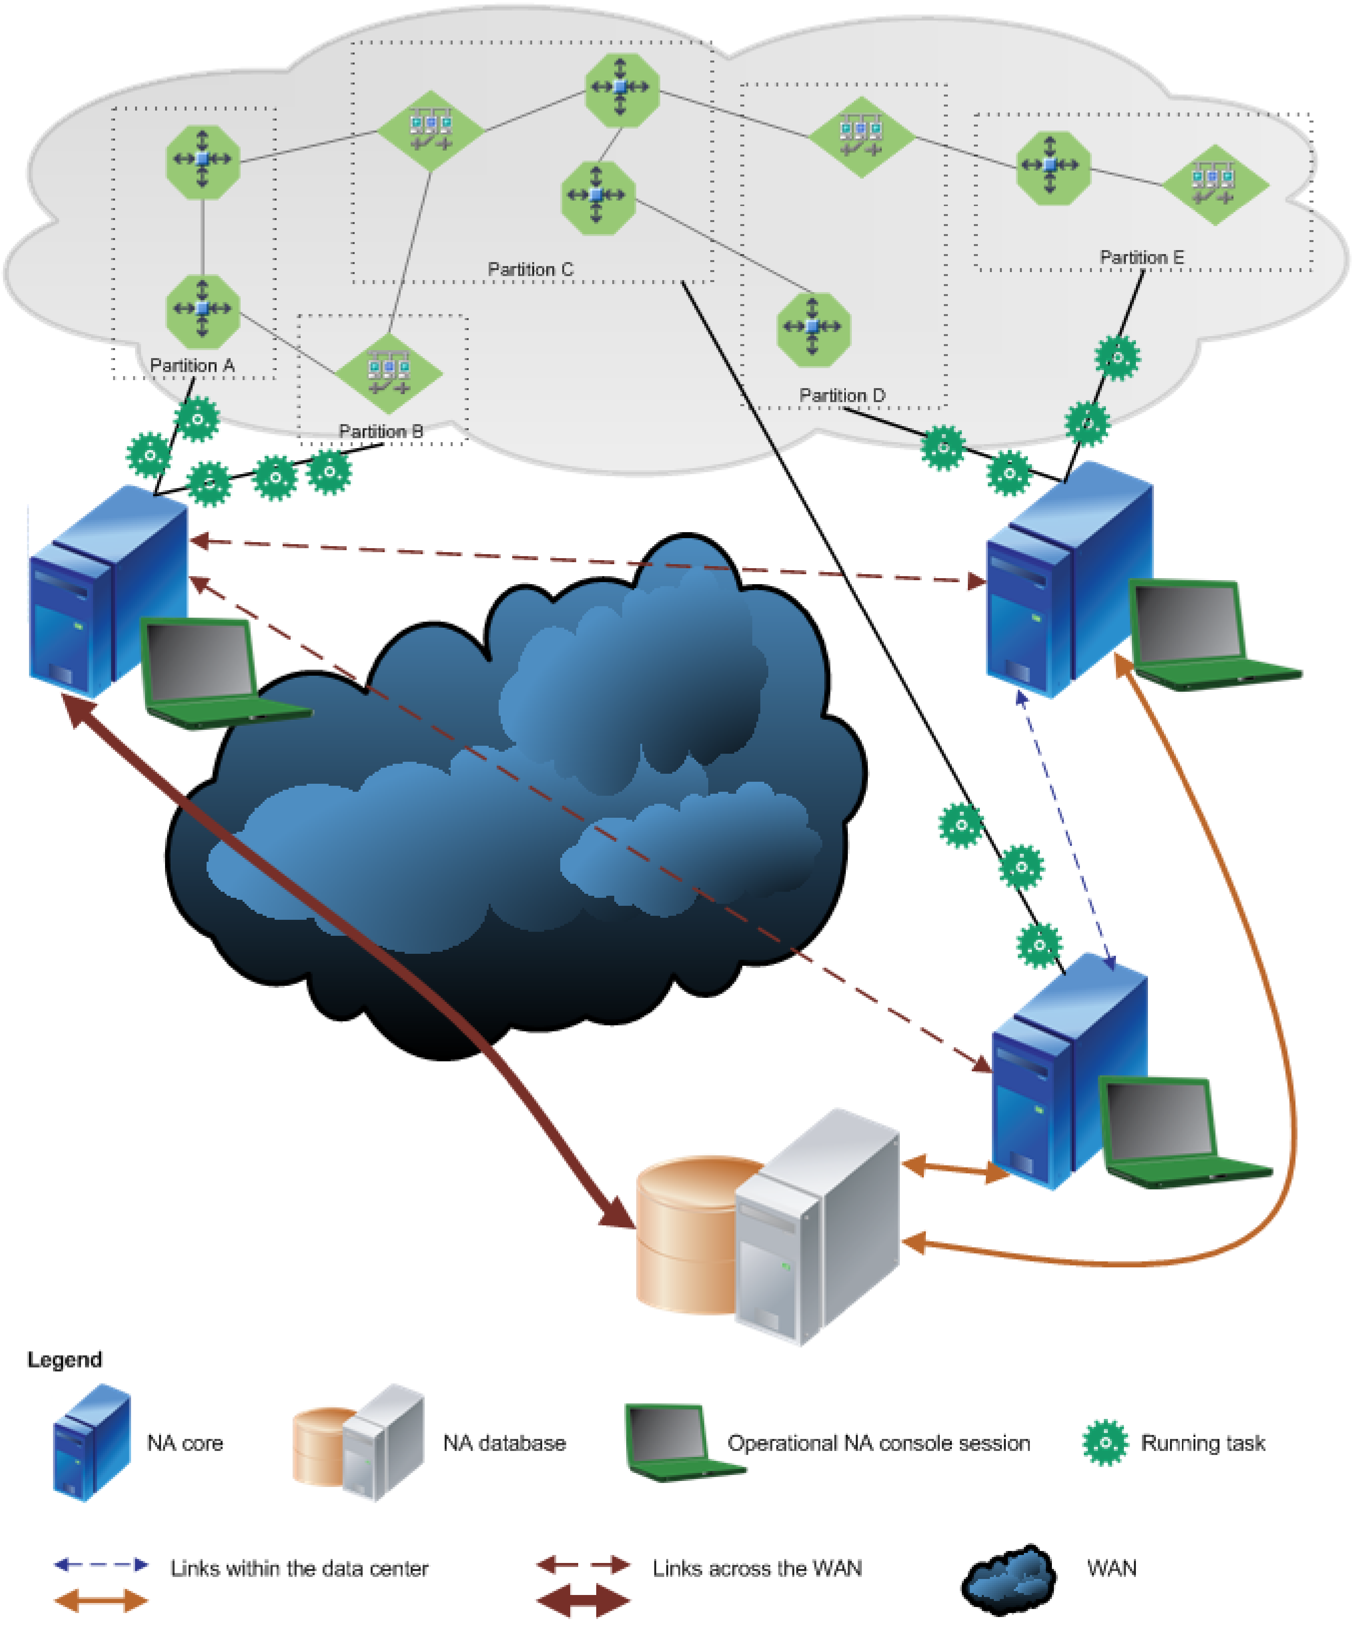

Chapter 2: Configuring Horizontal Scalability

This chapter describes how to configure an NA Horizontal Scalability environment. It includes the following topics:

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. For more information, see “Ports” in the *NA Administration Guide*.

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.

Table 2 Ports Used in an NA Horizontal Scalability Environment, continued

Protocol/Database/Port	From/To
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.
JBoss Remoting (port 4446)	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 - 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 - 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.

> **Note:** This procedure is for first-time configuration of Horizontal Scalability. If you have previously removed one or more NA cores from the Horizontal Scalability environment, follow the steps in "Adding

Additional NA Cores to the Horizontal Scalability Environment" on page 29.

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" below
- "Task 2: (New Installations Only) Install the First NA Core" on the next page
- "Task 3: Configure Horizontal Scalability for the First Two NA Cores" on the next page
- "Task 4: Configure the Standard Horizontal Scalability Topology" on page 26
- "Task 5: Optional. Configure Additional syslog Destinations for Continuity after NA Core Failover" on page 27

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 29
- "Configuring a Horizontal Scalability Topology Variation" on page 32
- "Configuring NA Satellites in a Horizontal Scalability Environment" on page 34
- "Configuring Dynamic Device Group Calculation in a Horizontal Scalability Environment" on page 43

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

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`
  - *Linux*: `/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 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 Installation and Upgrade 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.
   - *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 *`1521`* (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.

  The value of the time zone offset is derived by calculating the offset from UTC of the server. For example: a server in the Central European (CET) time zone, which is 1 hour ahead of UTC, requires the value `+1`.

5. Copy the customized initial setup script to the database server.
6. Stop all NA services on NA1. See "Start, Stop, or Restart All Services" on page 1.
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 28
   - "Running Scripts on SQL Server" on page 28
9. On the second NA server (NA2), install NA. When prompted, choose **an existing Network Automation database**, and then enter the database name from the initial setup script of step 3.
10. Stop all NA services on NA2. See "Start, Stop, or Restart All Services" on page 1.

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 32.)

1. In a text editor, such as WordPad or vi, open the distributed.rcx file in the known location of Task 3, step 1.
2. *Core Binding only*. Set the distributed/horizontalscalability option to true.

   ```
   <option name="distributed/bind_tasks_to_core">true</option>
   ```
3. 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
   - *Linux*: <NA_HOME>/jre
4. Start all NA services on both NA cores (NA1 and NA2). See "Start, Stop, or Restart All Services" on page 1.
5. *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.

> **Tip:** 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.

> **Tip:** 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.

6. Verify that the installation is working correctly as described in "Verifying Installation and Setup" on page 29.
7. 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 29.
   - 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 32.
   - 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 34.

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.

> **Note:** 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" below
- "Running Scripts on SQL Server" below

Running Scripts on Oracle

> **Note:** If you are running the customized setup script on an Oracle database installed on Linux, you must use the Oracle user account that has the **oinstall** ownership with 664 permission.

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:\`.
   - *Linux*: 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>
   ```

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

   For *<SID>*, use the Oracle SID of the NA database.
3. In the SQLPlus window, run the customized script. For example:

   ```
   run: SQL > @OracleInitialSetup.sql
   ```

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 *`1521`* (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.
5. If you have previously removed one or more NA cores from the Horizontal Scalability environment, do the following:
   a. To determine the NA core IDs that are currently in use, connect as an NA administrator to the NA proxy on an NA core, and then run the following command:

      ```
      list core
      ```

   b. To determine all NA core IDs that have ever been used, list the NA core IDs in the `RN_KEY_INCREMENTOR` table. For example:

      ```
      SELECT DISTINCT CoreID FROM RN_KEY_INCREMENTOR;
      ```

   c. To identify the NA core IDs that are available to be reused, compare the `list core` output to the results of the `RN_KEY_INCREMENTOR` table query. An NA core ID that appears in the `RN_KEY_INCREMENTOR` table but not in the `list core` output is available to be reused.
   d. In the variable replacement section of the `OracleAddServer.sql` file or the `SQLServerAddServer.sql` file, set the NA core ID to a specific value that is available to be reused.
      - Oracle: Use the following syntax:

        ```
        core_number := <value>;
        ```

For example: `core_number := 3;`

- SQL Server: Use the following syntax:

  `SET @core_number = <value>;`

  For example: `SET @core_number = 3;`

e. Also in the `OracleAddServer.sql` file or `SQLServerAddServer.sql` file, do the following:

- Insert two hyphens (--) at the beginning of the line that contains `MAX(CoreID) +1` to comment it out.
- Insert two hyphens (--) at the beginning of the `INSERT INTO RN_KEY_INCREMENTOR` line to comment it out.

f. Save the `OracleAddServer.sql` file or the `SQLServerAddServer.sql` file.

6. Copy the customized add server script to the database server.
7. Stop all NA services on all NA cores. See "Start, Stop, or Restart All Services" on page 1.
8. If NA contains production data, back up the NA file systems and the NA database.
9. 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 28
   - "Running Scripts on SQL Server" on page 28
10. On the new NA server (`NA3`), install NA. When prompted, choose **an existing Network Automation database**, and then enter the database name from the add server script of step 3.
11. Stop all NA services on `NA3`. See "Start, Stop, or Restart All Services" on page 1.
12. Repeat step 2 through step 11 to add each additional NA core to the Horizontal Scalability environment. You can add up to five NA cores (seven NA cores with the `NA_00050` patch).
13. 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`
    - *Linux*: `<NA_HOME>/jre`
14. Start all NA services on all NA cores. See "Start, Stop, or Restart All Services" on page 1.
15. *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.
16. Verify that the installation is working correctly as described in "Verifying Installation and Setup" on page 29.
17. Choose your next step:

- Optionally configure the environment with a topology variation as described in "Configuring a Horizontal Scalability Topology Variation" below.
- 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 34.

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" below
- "Configure the Core Binding Variation for User Interaction" on the next page
- "Configure the Core Binding Variation for Local Device Tasks" on the next page

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 34.

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.
   d. Click **Save**.

Tip: 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.

Note: 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 57.

Tip: 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. See "Start, Stop, or Restart All Services" on page 1.
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**.

> **Note:** 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.

Here is an example of a satellite gateway on a horizontal scalability set up:

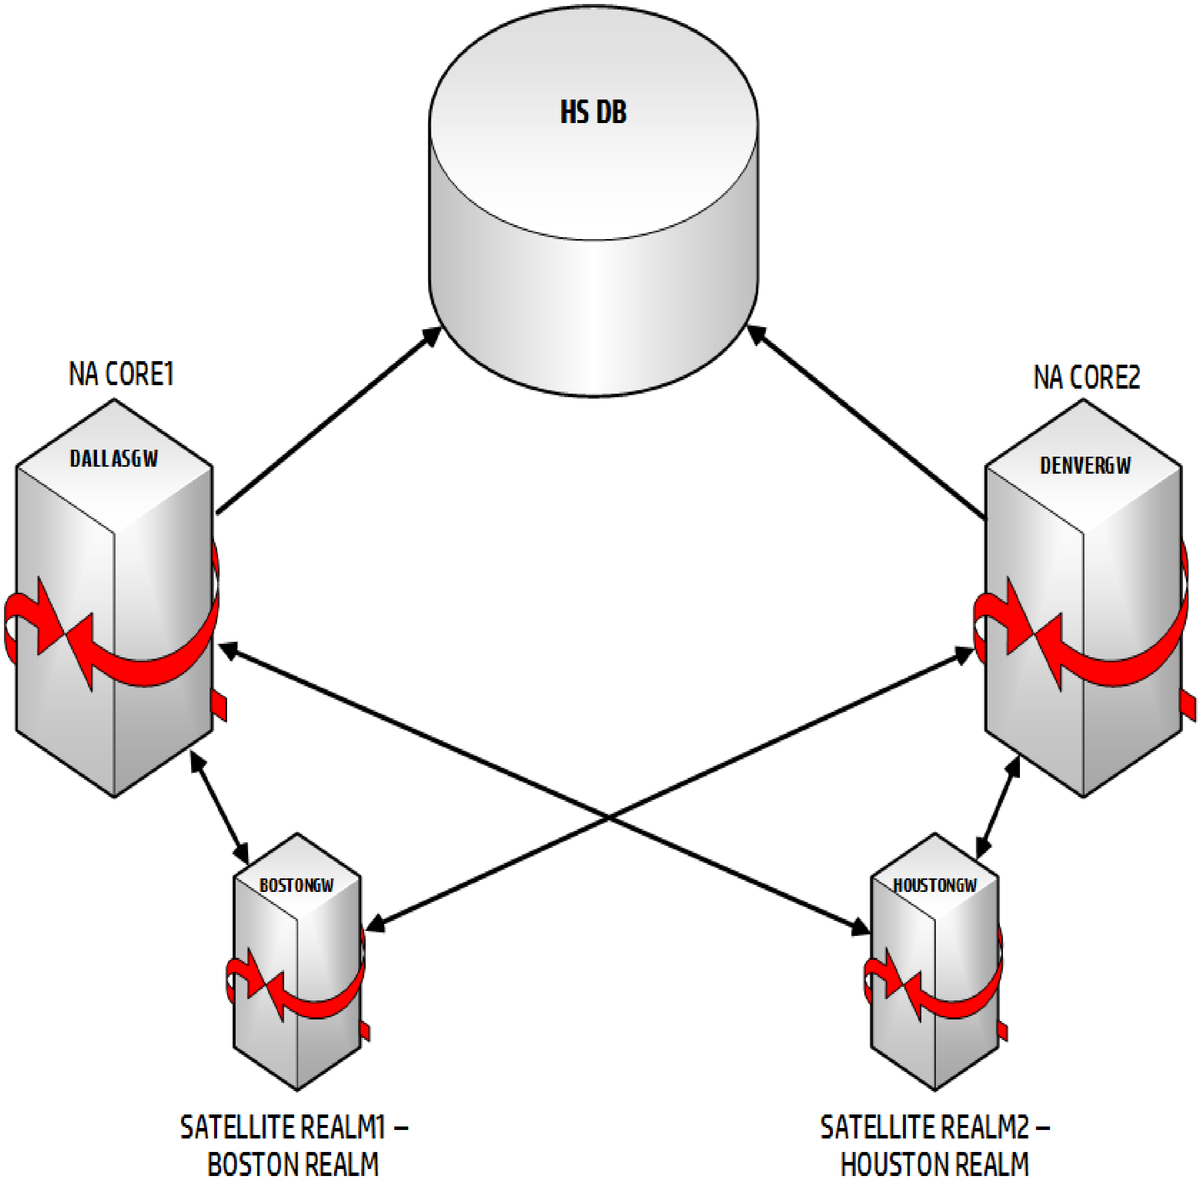
HS DB
NA CORE1
DALLASGW
NA CORE2
DENVERGW
BOSTONGW
HOUSTONGW
SATELLITE REALM1 – BOSTON REALM
SATELLITE REALM2 – HOUSTON REALM

The given example shows the setup of two Satellite or remote gateways on a two-HS core environment. The four different realms communicating with each other are as follows:

- Dallas Realm - NA Core1 with the DallasGW
- Denver Realm - NA Core2 with the DenverGW
- Boston Realm - Satellite Realm1 with the BostonGW (remote)
- Houston Realm - Satellite Realm2 with the HoustonGW (remote)

As the NA Cores are in a horizontal scalability setup, they can communicate with the devices managed by the Boston and Houston realms.

Prerequisites

- Horizontal Scalability setup on two NA Cores in the Distributed Round Robin or Core Binding topology - The number and placement of the NA satellites depend 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.

    > **Note:** 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.

    > **Note:** 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 57.

- Red Hat Enterprise Linux or SUSE Linux platform as described in the *NA Support Matrix*
- Two Red Hat Enterprise Linux or SUSE servers depending upon what is used in the Horizontal Scalability environment for the Remote Gateway setup

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

1. Install the first core gateway.
2. Install the second core gateway.
3. Install the first remote gateway.
4. Install the second remote gateway.
5. Set up each gateway.
6. Configure NA to communicate with the core gateways.
7. Deploy the remote agent.
8. View Satellite Monitor.

Installing the First Core Gateway

You must install a core gateway on or near one NA core.

To install the first core gateway, follow these steps:

1. Log on to the gateway server as the `root` user.
2. Change to the directory containing the gateway installer bundle (`nas_gw-*.zip`).
3. Unzip the gateway installer bundle.
4. Prepare the gateway server. For information about preparing the gateway server, see the *Prepare the Gateway Server for the NA Gateway Installer* topic in the *NA Satellite Guide*.
5. Run the gateway installer:

   ```
   perl install.pl
   ```
6. Supply the required information.

   Under Common Options, enter the number for the option to `Configure a new core mesh`.

   > **Note:** 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. You must use the same password in all the Core gateway and the Satellite gateway installation.

7. Supply the details based on the NA Core Application server installation:
   - If core gateway is installed on the same server as the NA core server, select `Y`. This copies the `opswgw-mngt-server.pkcs8` file to the `/opt/NA` directory.
   - If core gateway is not installed on the same server as the NA server, do the following:
     i. Store a copy of the `saOPSWgw*/certificates/opswgw-mngt-server.pkcs8` file.
     ii. After the installation of the core gateway is complete, place the file in the `/opt/NA` directory of the NA core server.

8. Supply the following details:
   - Name of the gateway (for example, DallasGW)
   - Name of the realm (for example, Dallas Realm)
9. Select 'Y' when prompted to proceed with the installation.

After the core gateway is installed, copy the `opswgw-crypto.tgz.e` from the Gateway Crypto Data file and place it in the `/tmp` directory of the same server as well as the `/tmp` directory of all the cores and gateway servers.

Installing the Second Core Gateway

To install the second core gateway, follow these steps:

1. Repeat steps 1 to 5 of "Installing the First Core Gateway" on the previous page.
2. Under Common Options, enter the number for the option to `Add a new core gateway to an existing mesh` and provide the following details:
   - The location (`/tmp` directory of the first core gateway server) of the `opswgw-crypto.tgz.e` file
   - The IP address or hostname of the first core gateway
   - Name of the core gateway (for example, DenverGW)
   - Name of the realm (for example, Denver Realm)
3. Select 'Y' when prompted to proceed with the installation.

Installing the First Remote Gateway

To install the first remote gateway, follow these steps:

1. Log on to the gateway server as the root user.
2. Change to the directory containing the gateway installer bundle (`nas_gw-*.zip`).
3. Unzip the gateway installer bundle.
4. Prepare the gateway server. For information about preparing the gateway server, see the *Prepare the Gateway Server for the NA Gateway Installer* topic in the *NA Satellite Guide*.
5. Run the gateway installer:

   ```
   perl install.pl
   ```
6. Under Common Options, enter the number for the option to `Add a new gateway to an existing mesh`. (Ensure that the option does not include the word "`core`.")
7. Provide the following information:

- The location (`/tmp` directory of the first core gateway server) of the `opswgw-crypto.tgz.e` file (for more information, see "Installing the First Core Gateway" on page 37)
- Name of the first remote gateway (for example, BostonGW)
- Name of the realm (for example, Boston Realm)

For more information about installing a remote gateway, see the *Install Remote Gateways* section of the *NA Satellite Guide*.

Installing the Second Remote Gateway

To install the second remote gateway, follow these steps:

1. Repeat steps 1 to 6 of "Installing the First Remote Gateway" on the previous page.
2. Provide the following information:
   - The location (`/tmp` directory of the second core gateway server) of the `opswgw-crypto.tgz.e` file (for more information, see "Installing the Second Core Gateway" on the previous page)
   - Name of the second Remote gateway (for example, HoustonGW)
   - Name of the realm (for example, Houston Realm)

For more information about installing a remote gateway, see the *Install Remote Gateways* section of the *NA Satellite Guide*.

Setting Up Each Gateway

This section describes the tasks to be performed on each gateway before configuring NA to communicate with the core gateways.

On the first core gateway (for example, DallasGW), do the following:

1. Under `/etc/opt/opsware/opswgw-DallasGW/opswgw.properties`, verify the following entries:
   - `opswgw.TunnelDst=2001:/var/opt/opsware/crypto/opswgw-<`*Name of the first gateway*`>/opswgw.pem`
   - `opswgw.IngressMap=127.0.0.1@NAS`
2. Use the field separator as follows:
   - For the IPv4-only format, use the field separator as `:` (for example, `opswgw.EgressFilter=tcp:*:443:127.0.0.1:*,tcp:*:22:NAS:,tcp:*:23:NAS:,tcp:*:513:NAS:`)

- For the dual-stack installation, use the field separator as **@** (for example, `opswgw.EgressFilter=tcp@*@443@::ffff:127:0:0:1@*,tcp@*@22@NAS@,tcp@*@23@NAS@,tcp@*@513@NAS@`).

  For the given example of DallasGW, the EgressFilter entry is as follows:

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

3. In case of any mismatch, add the settings accordingly and restart the gateway with the following command:

   ```
   /etc/init.d/opswgw-DallasGW restart
   ```

On the second core gateway (for example, DenverGW), do the following:

1. Under `/etc/opt/opsware/opswgw-DenverGW/opswgw.properties`, verify the following entries:

   -
     ```
     opswgw.TunnelSrc=<IP address or hostname of First Core
     Gateway>:2001:100:0:/var/opt/opsware/crypto/opswgw-DenverGW/opswgw.pem
     opswgw.TunnelDst=2001:/var/opt/opsware/crypto/opswgw-DenverGW/opswgw.pem
     ```
   - `opswgw.IngressMap=127.0.0.1@NAS` (If 127.0.0.1 is provided during installation, then change this to 127.0.0.1@NAS; else, provide the IP address of the NA server)

   -
     ```
     opswgw.EgressFilter=tcp:*:443:127.0.0.1:*,tcp:*:22:NAS:,tcp:*:23:NAS:,tcp:*:513:NA
     S:,tcp:*:443:NAS:,tcp:*:80:NAS:
     ```
2. In case of any mismatch, add the settings accordingly and restart the gateway with the following command:

   ```
   /etc/init.d/opswgw-DenverGW restart
   ```

On the first remote gateway (for example, BostonGW), do the following:

1. Under `/etc/opt/opsware/opswgw-BostonGW/ opswgw.properties`, do the following:
   a. Add an entry for the second core gateway.
   b. Enter the cost values as 100 and 200, as shown here:

   ```
   opswgw.TunnelSrc=<first Core Gateway IP address or
   hostname>:2001:100:0:/var/opt/opsware/crypto/opswgw-SatelliteGW/opswgw.pem
    opswgw.TunnelSrc=<second Core Gateway IP address or
   hostname>:2001:200:0:/var/opt/opsware/crypto/opswgw-SatelliteGW/opswgw.pem
   ```
2. Restart the remote gateway with the following command:

   ```
   /etc/init.d/opswgw-BostonGW restart
   ```

On the second remote gateway (for example, HoustonGW), do the following:

1. Under `/etc/opt/opsware/opswgw-HoustonGW/ opswgw.properties`, do the following:
   a. Enter the same TunnelSrc and cost values as that of the first remote gateway:

   ```
   opswgw.TunnelSrc=<first Core Gateway IP address or
   hostname>:2001:100:0:/var/opt/opsware/crypto/opswgw-HoustonGW/opswgw.pem
   ```

```
opswgw.TunnelSrc=<second Core Gateway IP address or
hostname>:2001:200:0:/var/opt/opsware/crypto/opswgw-HoustonGW/opswgw.pem
```

b. Restart the remote gateway with the following command:

```
/etc/init.d/opswgw-HoustonGW restart
```

Configuring NA to Communicate with the Core Gateway

To configure each NA core to communicate with the associated core gateway, follow these steps:

1. *Skip this step if the core gateway is installed on the NA core server.* Copy the `opswgw-mngt-server.pkcs8` file from the core gateway to the `<NA_HOME>` directory—typically `C:\NA` or `/opt/NA`.
2. Log on to the NA console as an administrator.
3. On the Administrative Settings - Device Access page (Admin > Administrative Settings > Device Access), under Gateway Mesh, identify the core gateway.

   For the Local Gateway Host field, enter the DNS hostname or IP address of the core gateway associated with this NA core. If the core gateway is installed on the NA core server, specify the localhost.
4. Click **Save**.

To verify whether NA is able to communicate with the core gateway, on the main menu bar of the NA console, click Admin > Gateways. The Gateway List page opens. On this page, you can view the details of the Core Gateways and the Remote Gateways that you configured.

Deploying the Remote Agent

To install the NA remote agent on a remote gateway server, follow these steps:

1. Log on to the NA console as an NA administrator.
2. On the Deploy Remote Agent page (Tasks > New Task > Deploy Remote Agent), configure the deploy agent task.

   Under Task Options, note the following:

   - For **Action**, select **Install** (or **Reinstall**).
   - For **Deploy Agent to Gateway**, select the target remote gateway from the list.
   - For **Login**, do one of the following:
     - Select **As Root**, and then enter the password for the root user on the remote gateway server.
     - Select As Non-root, select the method for gaining root access to the remote gateway server, and then enter the relevant password.

- For **Managing Core**, enter the IP address of the NA core that should receive communication from the remote gateway. This value was entered as the IP address of the core application server during installation of the core gateway. For dual-stack satellites, use only the IPv4 portion of the IP address.

- For **In Realm**, select the realm name of the core gateway associated with the managing core. In most cases, this value is `Default Realm`.

For more information, see "Deploy Remote Agent Page Fields" in the NA help.

3. Click **Save**.
4. *Dual stack satellites only*. After deploying the remote agent, set the TFTP server on the satellite to run on the IPv6 address of the remote gateway server. To achieve this, follow these steps:
   a. Connect to the remote gateway server as the root user.
   b. Change to the following directory:
   ```
   /opt/opsware/nassat/jre
   ```
   c. Back up the `nassat.rcx` file to a location outside the gateway installation directory.
   d. In a text editor, open the `nassat.rcx` file.
   e. Add the following line:
   ```
   <option name="TFTP/Server/IPv6">$ipV6address$</option>
   ```
   f. In the new line, replace `$ipV6address$` with the IP address of the remote gateway server.
   g. Save the `nassat.rcx` file.
   h. Restart the gateway by running the following command:
   ```
   /etc/init.d/nassat restart
   ```

Monitoring Satellite

You can ensure that all the satellites and the remote gateways are operational and running the same version by using the Monitor Satellite link on the Gateway List page of the NA console. To check if the configured remote gateways are operational, follow these steps:

1. Log on to NA as an administrator.
2. Select **Admin** > **Gateways**.
3. On the Gateway List page, check if the name of the core is displayed for each core gateway. Else, on the Edit Core page (**Admin** > **Distributed** > **Core List**), enter the Realm name of the core gateway, and then click **Save**. The Gateway List page displays the name of the core for the core gateway.
4. Connect to the remote gateway server as the root user, and then do the following:
   a. Change to the following directory:
   ```
   /opt/opsware/nassat/jre
   ```
   b. Back up the `nassat.rcx` file to a location outside the gateway installation directory.

c. In a text editor, open the `nassat.rcx` file.

d. Add the following:

```
<option name="syslog/usealternate/">true</option>
<array name="gateway/monitor/core_ips">
<value>core IP</value>
</array>
```

e. In the new line, replace `core IP` with the IP address of the NA core. Make sure that you add each core of the horizontal scalability environment, one after the other. For example, for a two-core horizontal scalability environment, add as follows:

```
<value>core1 IP</value>
<value>core2 IP</value>
```

You can replace the core1 IP and core2 IP with 127.0.0.1 or with a core IP depending on how you set up the satellite mesh. If you are using redundant satellites, then you must configure the NA core IPs instead of 127.0.0.1.

f. Save the `nassat.rcx` file.

g. Restart the gateway by running the following command:

```
/etc/init.d/nassat restart
```

5. On the NA console, on the Gateway List page, click the Monitor Satellite link. The Monitor Details page opens for you to view the monitor status.

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, do the following:

- Enable full cycle update on one core
- Enable event driven update on the second core
- Disable dynamic group calculation on the remaining cores

To enable the full cycle update on one core, follow these steps:

1. Identify the core on which the full cycle update is to be enabled.
2. On the core, add the following parameters in the `adjustable_options.rcx` file:

```
<option name="dynamic_group/disable">false</option>
<option name="dynamic_group/disable_event_listener">true</option>
```

```
<option name="dynamic_group/disable_queued">true</option>
<option name="monitor/DynamicDeviceGroupMonitor/enabled">true</option>
<option name="dynamic_group/queued_update_interval">30</option>
<option name="performance/device_group_commit_interval">10</option>
```

To enable the event driven update on the second core, follow these steps:

1. Identify the core on which the event driven update is to be enabled.
2. On the core, add the following parameters in the `adjustable_options.rcx` file:

   ```
   <option name="dynamic_group/disable">false</option>
   <option name="dynamic_group/disable_event_listener">true</option>
   <option name="dynamic_group/disable_queued">false</option>
   <option name="monitor/DynamicDeviceGroupMonitor/enabled">true</option>
   <option name="dynamic_group/queued_update_interval">1</option>
   <option name="performance/device_group_commit_interval">10</option>
   ```

To disable the dynamic group calculation, on the rest of the cores, add the following parameters in the `adjustable_options.rcx` file:

```
<option name="dynamic_group/disable">true</option>
<option name="dynamic_group/disable_event_listener">true</option>
<option name="dynamic_group/disable_queued">true</option>
<option name="monitor/DynamicDeviceGroupMonitor/enabled">false</option>
<option name="dynamic_group/queued_update_interval">30</option>
<option name="performance/device_group_commit_interval">10</option>
```

On all the cores (irrespective of whether the dynamic group calculation is enabled or disabled), add the following parameters under `<array name="distributed/core-specific-options">` in the `adjustable_options.rcx` file:

```
<value>dynamic_group/disable</value>
<value>dynamic_group/disable_queued</value>
<value>monitor/DynamicDeviceGroupMonitor/enabled</value>
<value>dynamic_group/queued_update_interval</value>
<value>performance/device_group_commit_interval</value>
<value>dynamic_group/update_interval</value>
<value>dynamic_group/event_driven_recalc</value>
```

Best Practices to Set Core-Specific Properties

Every upgrade or patch install overwrites the changes that you make to the `appserver.rcx` file. Therefore, it is recommended that you copy the `distributed/core-specific-options` from the `appserver.rcx` file to the `adjustable_options.rcx` file, and then enter the latest values to the `adjustable_options.rcx` file.

To set properties specific to a core, you must use the `<array name="distributed/core-specific-options">` in the `adjustable_options.rcx` file, and set the property name between the`<value></value>` pair (`<value>`*property*`</value>`). For example, `<value>dynamic_group/disable</value>`.

As the `<array name="distributed/core-specific-options">` is core-specific, you must add or update a property in the `adjustable_options.rcx` file on each core. You can view the changes made to the `.rcx` file after you do a reload of the file on each core. To reload the `.rcx` settings, do one of the following:

- In the NA web user interface, on the Admin > Administrative Settings > User Interface page, click **Save**.
- Run the `reload server options` command from the NA proxy.
- Restart the NA services.

If a property is in the `adjustable_options.rcx` file on core1, and in the `appserver.rcx` file on core2, the reload synchronizes the property only in the respective cores—it does not move the property from the `adjustable_options.rcx` file of core1 to the `adjustable_options.rcx` file of core2, or the the `appserver.rcx` file of core2 to the `appserver.rcx` file of core1.

Note: Synchronization occurs only on the active cores of the environment. If an NA core is not operational during the change replication, at a later time, use the Admin > Distributed > Renew Configuration Options page to push changes to that core.

Note: The `securityfilter_additional_init.rcx` file is also core-specific and does not get replicated with a reload of the `.rcx` settings.

When a driver is added or updated on a core, clicking the Reload Drivers button on the NA web user interface (Admin > Start/Stop Services) replicates the changes across all the active cores in the environment.

Configuring Each NA Core with Different NAT TFTP IP Addresses

Using NAT TFTP IP addresses in a Horizontal Scalability environment can require that different NAT TFTP IP addresses be used on each NA core.

To configure separate NAT TFTP IP addresses for each NA core

1. On the Edit Device page for each device whose NAT TFTP depends on which NA core is used, set the value of the **TFTP Server IP Address** field to 0.0.0.0.
2. Enable a custom data field for device groups with the API Name `nat_tftp`. The value of the Display Name field for the custom data field does not matter.
3. Put each device in a device group with the `nat_tftp` custom field set in the IP addresses of the NA

cores in the following syntax (space separated): <core_number>=<IP_address>

For example, `1=10.1.2.3 2=10.4.5.6 3=10.7.8.9` represents three NA cores with the following IP addresses:

- On NA core 1, use the NAT TFTP address 10.1.2.3.
- On NA core 2, use the NAT TFTP address 10.4.5.6.
- On NA core 3, use the NAT TFTP address 10.7.8.9.

While accessing a device with the NAT TFTP address set to 0.0.0.0, NAsearches for all device groups containing that device.

If multiple device groups include the `nat_tftp` custom field, NA uses the device group lowest in the device parent hierarchy. If the groups are not in a hierarchy, NA reports an error and cannot access the device.

If no device groups include the `nat_tftp` custom field, NA reports an error and cannot access the device.

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.

Upgrading Horizontal Scalability

> **Note:** If the NA environment includes disaster recovery configuration, verify that all inactive NA cores are connected to the operational NA database. Then, upgrade all active and all inactive NA cores as described here.

You can upgrade NA in a Horizontal Scalability environment from version 10.0x or 10.1x to version 10.10 either on the same systems on which the NA cores are installed or from different systems.

You must upgrade each NA core server as described in the *NA Installation and Upgrade Guide10.21*. (For information about upgrading on the same system, see *Upgrading to NA10.10 on the Same System*, and for information about upgrading from a different system, see *Upgrading to NA10.10 on a Different System* in the NA Installation and Upgrade Guide.)

While upgrading the NA cores, do not forget the following:

- Run the upgrade procedures in parallel. That is, complete step 1 on each NA server before initiating step 2 on any NA server, and so forth as shown in the following screenshots:

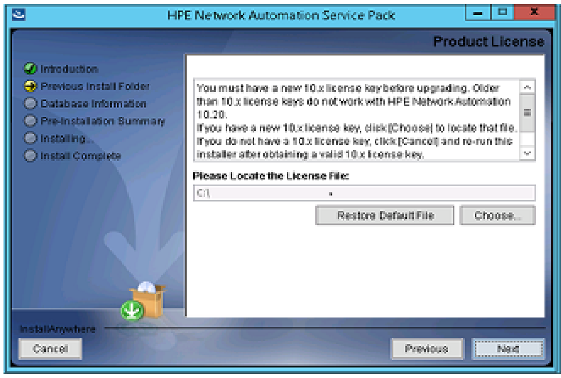

You must have a new 10.x license key before upgrading. Older than 10.x license keys do not work with HPE Network Automation 10.20.
If you have a new 10.x license key, click [Choose] to locate that file.
If you do not have a 10.x license key, click [Cancel] and re-run this installer after obtaining a valid 10.x license key.

Please Locate the License File:
C:\

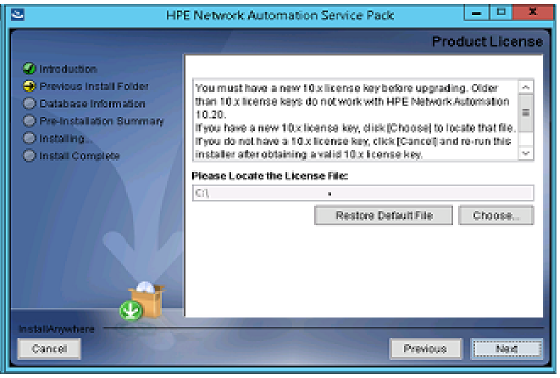

You must have a new 10.x license key before upgrading. Older than 10.x license keys do not work with HPE Network Automation 10.20.
If you have a new 10.x license key, click [Choose] to locate that file.
If you do not have a 10.x license key, click [Cancel] and re-run this installer after obtaining a valid 10.x license key.

Please Locate the License File:
C:\

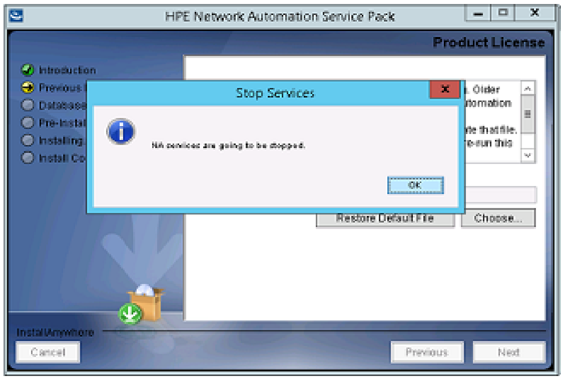

NA services are going to be stopped.

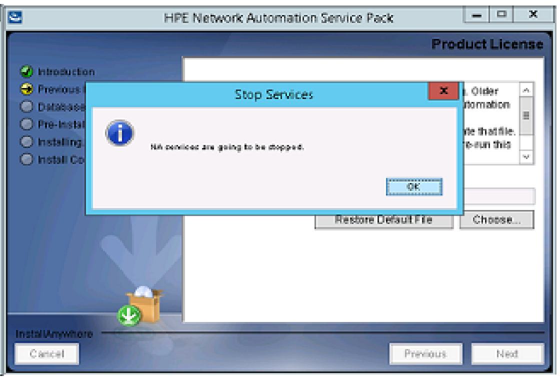

NA services are going to be stopped.

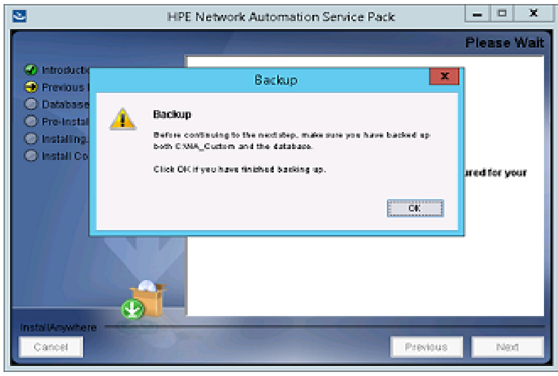

Backup

Backup

Before continuing to the next step, make sure you have backed up both C:\NA_Custom and the database.

Click OK if you have finished backing up.

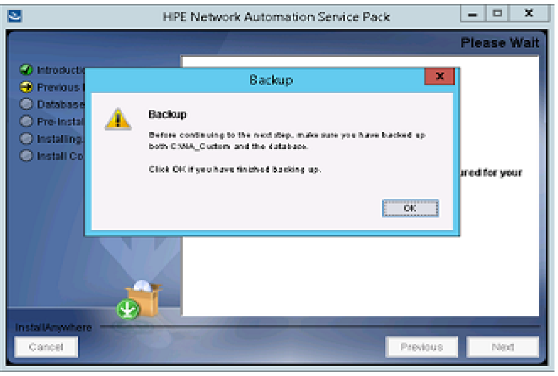

Backup

Backup

Before continuing to the next step, make sure you have backed up both C:\NA_Custom and the database.

Click OK if you have finished backing up.

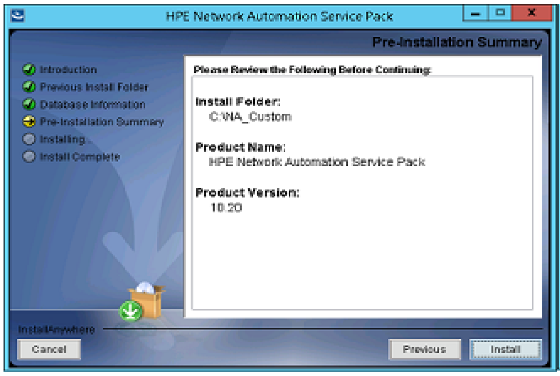

Please Review the Following Before Continuing:

Install Folder:
C:\NA_Custom

Product Name:
HPE Network Automation Service Pack

Product Version:
10.20

Please Review the Following Before Continuing:

Install Folder:
C:\NA_Custom

Product Name:
HPE Network Automation Service Pack

Product Version:
10.20

- Follow the referenced procedure through the step to run the NA 10.10 Service Pack Installer.
- To prevent database access, on each NA server, after the NA Service Pack Installer runs, stop all NA services. See "Start, Stop, or Restart All Services" on page 1.

> **Note:** If you need to upgrade the database product on the NA database server to a newer version, do so only one time, not for each NA server.

Upgrading to NA 10.20 on the Same Systems

To upgrade NA cores, from version 10.0x or 10.1x, to version 10.10 on the same systems on which they are currently installed, follow the procedure given in "Upgrading Horizontal Scalability" on the previous page, and then follow these steps:

1. On each NA core server, edit the `distributed.rcx` file as follows:
   a. If the following line does not exist, add it:
      ```
      <option name="distributed/horizontalscalability">true</option>
      ```
   b. If you want to enable distributed round robin, add or edit the following line:
      ```
      <option name="distributed/bind_tasks_to_core">false</option>
      ```
2. Start all the NA services on the master core server and make sure that you can login to NA on the master core. See "Start, Stop, or Restart All Services" on page 1.
3. Start the NA services on all the other new cores.

4. Return to the upgrade procedure in the *NA Installation and Upgrade Guide*. Start at the step after running the NA Service Pack Installer, and work through to the end of the procedure.

Upgrading to NA 10.20 from Different Systems

To upgrade NA cores, from version 10.0x or 10.1x, from different systems, follow the procedure given in "Upgrading Horizontal Scalability" on page 47, and then follow these steps:

1. Add the NA cores to the 10.21 Horizontal Scalability environment. During installation, when prompted for the database credentials, choose: use existing database.
2. Start all NA services on all the cores. See "Start, Stop, or Restart All Services" on page 1.
3. Stop the NA services on all the cores.
4. Update the database to set the `COREHOSTNAME` column in the `RN_CORE` table to the new servers (such as server A and server B).

   The example here shows how to execute this step in a standard HS environment with two cores.

   On the respective database server, run the following script:

   Oracle Script

   ```
   UPDATE RN_CORE SET COREHOSTNAME = '<Server A FQDN>' WHERE COREID = 1;
   UPDATE RN_CORE SET COREHOSTNAME = '<Server B FQDN>' WHERE COREID = 2;
   COMMIT;
   ```

   SQL Server Script

   ```
   UPDATE RN_CORE SET COREHOSTNAME = '<Server A FQDN>' WHERE COREID = 1;
   UPDATE RN_CORE SET COREHOSTNAME = '<Server B FQDN>' WHERE COREID = 2
   ```

5. Start the NA services on the master core server and make sure that you can login to NA on the master core.
6. Start the NA services on all the other new cores.

Chapter 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.)

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.

> **Note:** In a disaster recovery configuration, only active NA cores participate in NA core failover.

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

This chapter contains the following topics:

- "Enabling Failover" below
- "NA Core Failover Behavior" on the next page
- "Optional Actions after NA Core Failover Occurs" on page 58
- "Failing Back to the Original NA Core" on page 60

Enabling Failover

Installing NA 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 the next page
- "Re-Enabling Failover" on the next page

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 active NA cores in the Horizontal Scalability environment.

Re-Enabling Failover

To re-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`:

   ```
   <option name="distributed/enable_auto_failover">true</option>
   ```
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 active 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 configurable delay 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 7 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 the next page
- "Task Status" on page 53
- "Movement of Tasks" on page 53

- "NA Failover Events" on page 57
- "Configuring Failover Order (Core Binding Only)" on page 57

Figure 7 NA Core Failover Conceptual Flowchart

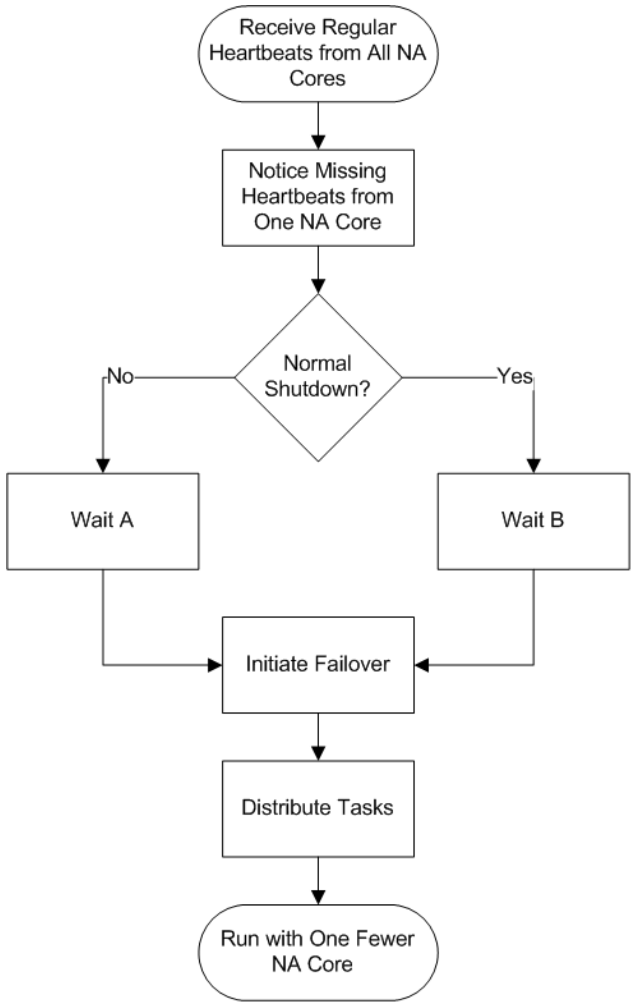
Receive Regular Heartbeats from All NA Cores
Notice Missing Heartbeats from One NA Core
Normal Shutdown?
No
Yes
Wait A
Wait B
Initiate Failover
Distribute Tasks
Run with One Fewer NA Core

Delay Before Initiating Failover

After an NA core stops, a running NA core notices the stopped NA core and waits a configurable delay 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 (by default) before initiating the NA core failover process. This time is represented by **Wait A** in Figure 7.

  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 (by default) before initiating the NA core failover process. This time is represented by **Wait B** in Figure 7. 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" below), 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.

> **Note:** 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" below). 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 8.

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

> **Note:** 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 9.
- 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 57.

Note: 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 8 Standard Distributed Round Robin Task Failover

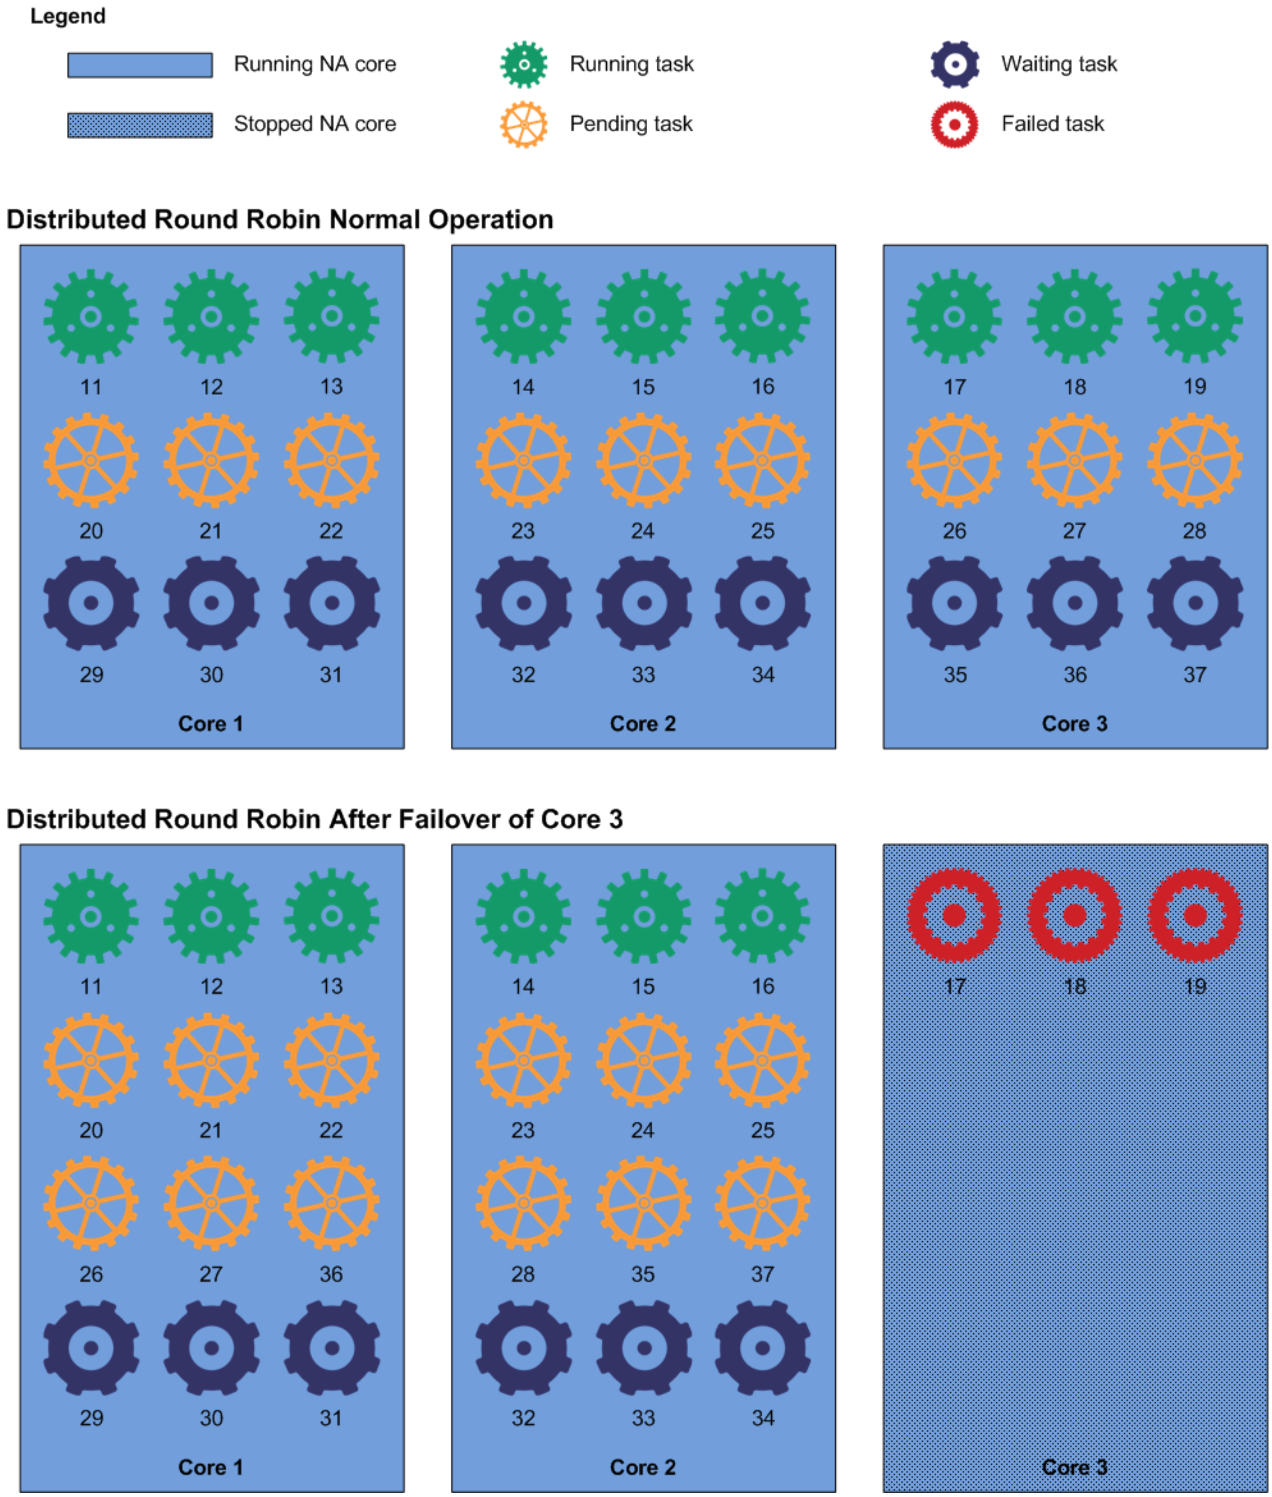

Figure 9 Standard Core Binding Task Failover

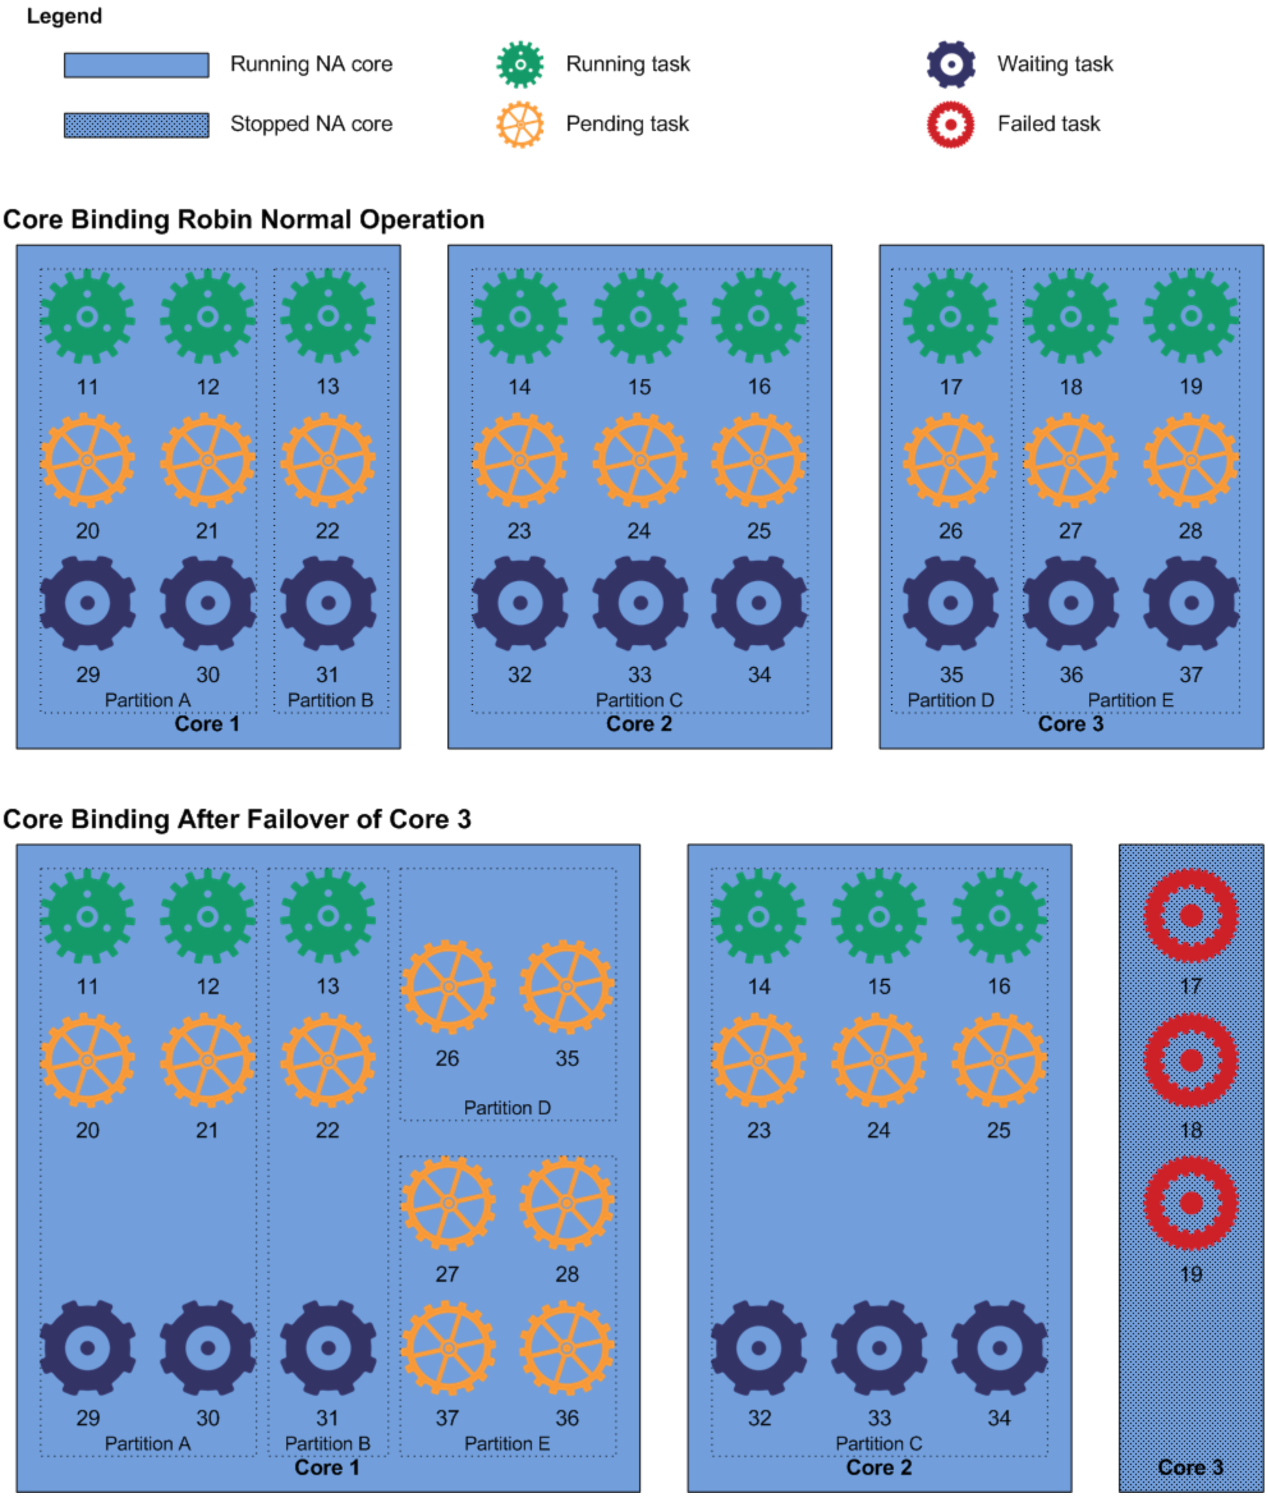

NA Failover Events

NA includes three events related to NA core failover:

- 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>
      ```

      > **Tip:** 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. See "Start, Stop, or Restart All Services" on page 1.

Optional Actions after NA Core Failover Occurs

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

- "Notify Users" on the next page
- "Configure Integrations to a Different NA Core" on the next page
- "Pause Non-Critical Tasks" on the next page
- "Locate Tasks Impacted by NA Core Failover" on the next page
- "Stop the Core Gateway for the Stopped NA Core" on page 60

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 53, 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.

   > **Tip:** The search results might include tasks that were cancelled because the NA server rebooted.

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

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.

Chapter 4: System Administration

This chapter contains the following topics:

- "NA-Generated Events for Horizontal Scalability" below
- "Using the NA Distributed System Pages" on the next page

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" below
- "Distributed System – RMI Error" below

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 core server as follows:

> **Note:** 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`
   - *Linux*: `/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 66 for information.

  > **Note:** 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.

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.

> **Tip:** 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. In a disaster recovery configuration, this page lists all active and inactive NA cores.

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: • Running: Fully functional—The NA core is running normally and active. This NA core runs user-initiated and pre-planned tasks. • Running: User interaction only, no tasks—The NA core is running normally and active. 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 and is in the active state. • Stopped: Abnormal shutdown—The NA core shut down unintentionally and is in the active state. • Running: Inactive—The NA core is running but inactive; it is not available to run any tasks. • Shutdown: Inactive—The NA core is shut down and in the inactive state.
Timezone Offset	The time zone offset of the NA server. The value of the time zone offset is derived by calculating the offset from UTC of the server. For example: a server in the Central European (CET) time zone, which is 1 hour ahead of UTC, requires the value +1.
Realm	The default Realm for the NA core.
Actions	The available action is: • Edit — Open the Edit Core page. See "Edit Core Page" on the next page.

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. The value of the time zone offset is derived by calculating the offset from UTC of the server. For example: a server in the Central European (CET) time zone, which is 1 hour ahead of UTC, requires the value +1.
Comments	Text comments about the NA core.
Realm Name	The default Realm for the NA core. For information about segmenting devices, see the *NA User Guide*. !!Note to self: Update this reference for Amazon!!

Device Password Rule Priority Reset Page

The Device Password Rule Priority Reset page is for Multimaster Distributed System environments. It 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**.

Appendix A: Common Procedures

This section describes procedures that are common to many HP Network Automation Software (NA) configuration and maintenance tasks. It includes the following topics:

- "Start, Stop, or Restart All Services" on page 1
- "Disable All Services" on page 1
- "Working with .rcx Files" on page 1

We appreciate your feedback!

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Horizontal Scalability Guide, March 2018 (Network Automation Software 10.21)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to network-management-doc-feedback@hp.com.