

HP Network Automation Software

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

Software Version: 9.20

Horizontal Scalability

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



Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notices

© Copyright 2008–2012 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

AMD is a trademark of Advanced Micro Devices, Inc.

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

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

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

UNIX® is a registered trademark of The Open Group.

Oracle Technology — Notice of Restricted Rights

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

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

Acknowledgements

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

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

Documentation Updates

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

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

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

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

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

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

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

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

Support

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

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

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

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

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

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

To find more information about access levels, go to:

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

To register for an HP Passport ID, go to:

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

Contents

| | | |
|----------|-------------------------------------------------------|----|
| 1 | Introduction | 7 |
| | Terminology | 7 |
| | Supported Databases and System Requirements | 7 |
| | What is NA Horizontal Scalability? | 8 |
| | Architecture | 9 |
| | Topologies | 9 |
| | Configuration with NA Satellites | 12 |
| | Load-Balancing | 12 |
| | Known Issues and Limitations | 13 |
| 2 | Configuring Horizontal Scalability | 15 |
| | Configuring NA Horizontal Scalability on Two NA Cores | 15 |
| | Scripts | 17 |
| | Running Scripts on Oracle | 17 |
| | Running Scripts on SQL Server | 17 |
| | Adding Additional NA Cores | 18 |
| | Verifying Installation and Setup | 19 |
| | Configuring Topology After Setup | 19 |
| | Post Installation Setup | 20 |
| | Uninstall Procedures | 21 |
| | Upgrading Horizontal Scalability | 21 |
| 3 | System Administration | 23 |
| | Managing Your Horizontal Scalability NA Mesh | 23 |
| | NA-Generated Events | 23 |
| | Using the NA Distributed System Pages | 24 |
| | Distributed Monitor Results Page | 24 |
| | Site Reassignment Page | 25 |
| | List Cores Page | 25 |
| | Edit Core Page | 25 |
| | Renew Configuration Options Page | 26 |

1 Introduction

This guide describes the HP Network Automation Software (NA) Horizontal Scalability functionality. It includes information on:

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



This document is updated as new information becomes available. To check for recent updates, or to verify that you are using the most recent edition of a document, go to:

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

For more information, see [Documentation Updates](#) on page 3.

Terminology

The following terms are used throughout this guide:

- **NA Core:** A single NA Management Engine, associated services (Syslog and TFTP), and a single database. An NA Core can manage multiple Sites.
- **Partition:** A set of devices. A Partition is managed by one (and only one) NA Core. Multiple Partitions can be managed by a single NA Core. For information on segmenting devices, see the *NA User Guide*.
- **NA Mesh:** Multiple NA Cores connected via replication.

Supported Databases and System Requirements

For information about the databases on which the NA Horizontal Scalability solution is supported, see the *NA Support Matrix*.

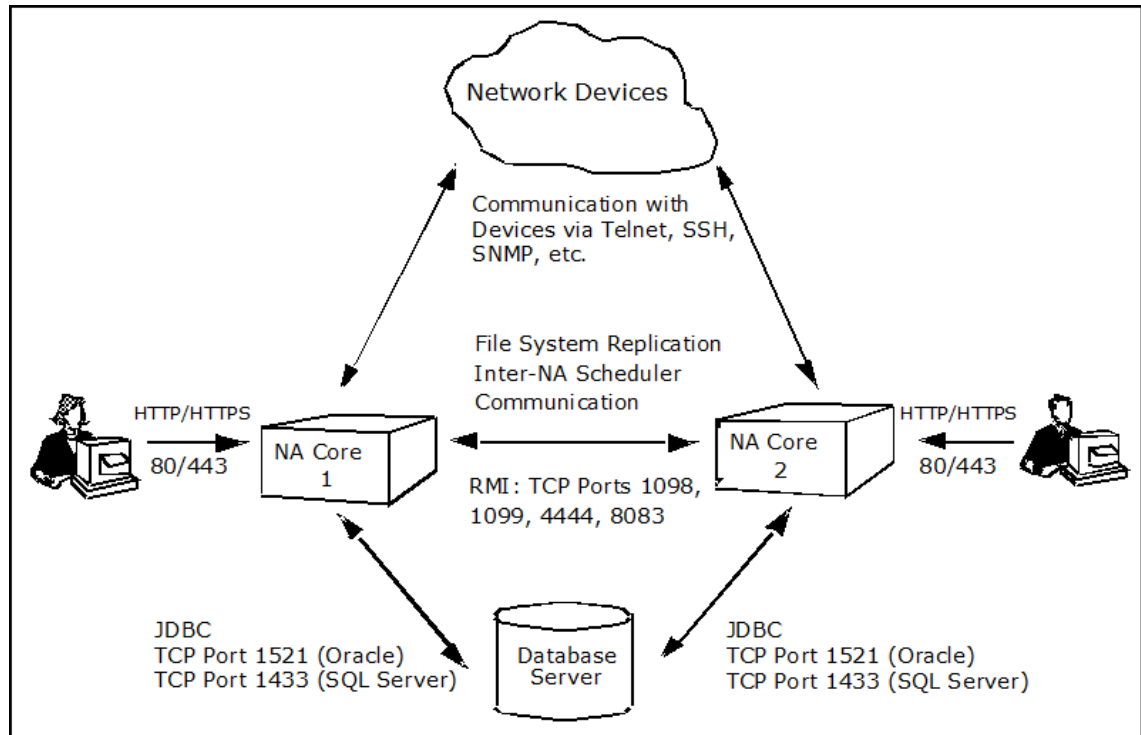
The maximum number of NA Cores supported is five.

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

What is NA Horizontal Scalability?

NA Horizontal Scalability is the ability to combine multiple NA Cores with a single database so that they work as a single logical unit to improve the performance of the overall system. NA Horizontal Scalability can be configured in different ways, depending on how you use your system and where bottlenecks are apt to occur.

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



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

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

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

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

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

These monitors generate events during error conditions. For more information, see [NA-Generated Events](#) on page 23. The events can be emailed to the NA administrator using an event rule in NA. For information on configuring email notification, see the *NA User Guide*.

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

- The NA Core's memory or CPU usage is high due to running many concurrent tasks
- The NA Core's memory or CPU usage is high due to many concurrent users on the system
- NA is not able to run as many concurrent tasks as you would like due to memory restrictions



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

Architecture

The NA Horizontal Scalability architecture has multiple components:

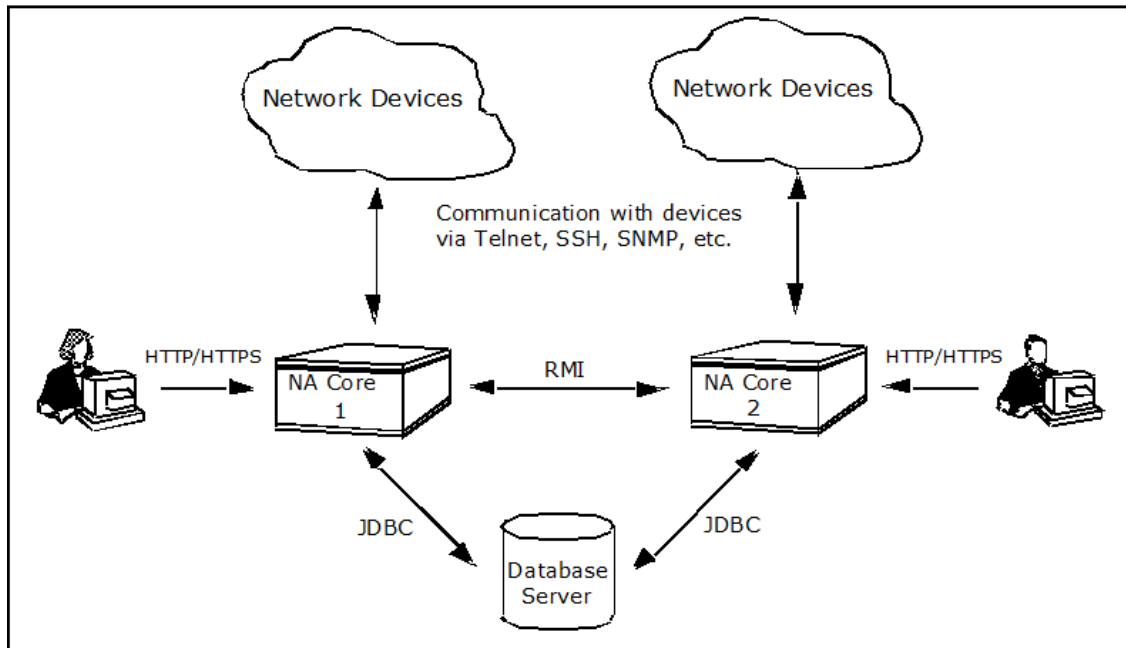
- An NA Mesh scheduler that is optimized for group task throughput. The scheduler enables group tasks to be load-balanced across multiple cores. The scheduler is not currently optimized for task failover, however.
- A file system synchronizer. Certain file system items are kept synchronized across multiple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

Topologies

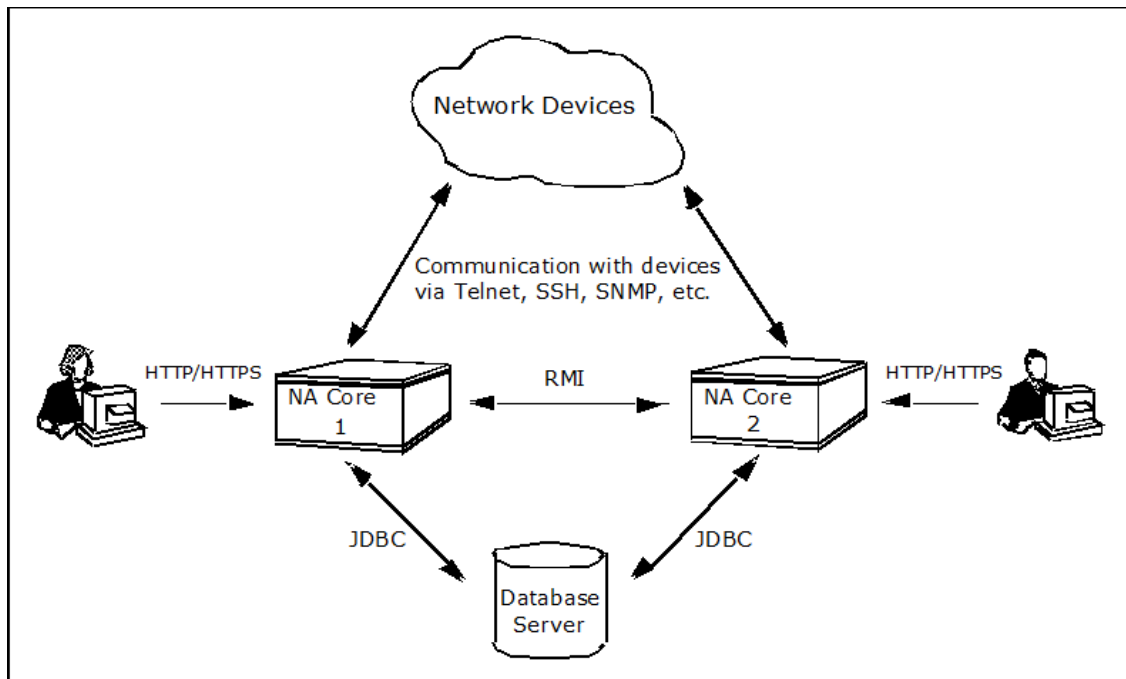
The NA Mesh can be configured in multiple topologies, depending on your requirements. For information about configuring the NA Mesh with these topologies, see [Configuring Topology After Setup](#) on page 19.

The following general topologies are supported:

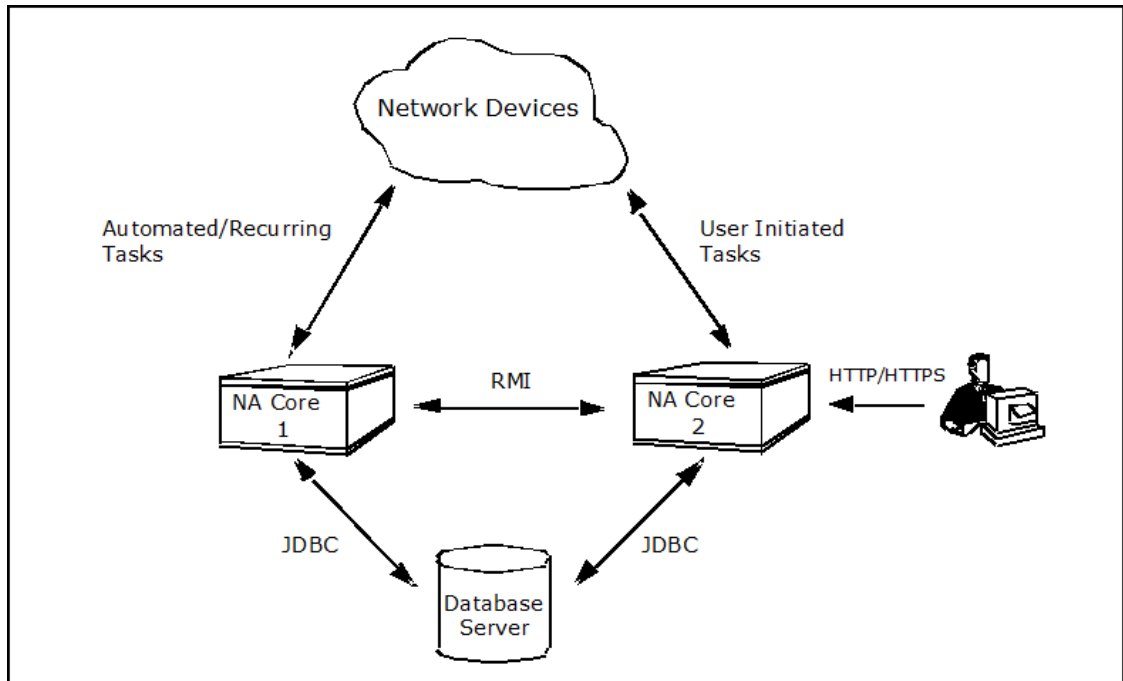
- Topology 1:** Dedicating each server to a fixed set of devices. You should use this topology if you have business units that need their own servers for performance reasons. This topology allows for global reporting on all devices.




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



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



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



- ▶ Horizontal Scalability in conjunction with Multimaster Distributed Systems is not a supported configuration due to the replication load that Horizontal Scalability generates. For more information, see the *NA Support Matrix*.

Configuration with NA Satellites

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

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

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

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

Load-Balancing

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

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

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

Known Issues and Limitations

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

If there is a firewall in the network configuration, make sure all the appropriate TCP ports are opened for traffic to tunnel through the firewalls. For example, ports 1099 must be open between the NA Cores. For detailed information on port usage, see “Protocols, Databases, and Ports” in the *NA Installation and Upgrade Guide*.

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

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

2 Configuring Horizontal Scalability

Configuring NA Horizontal Scalability on Two NA Cores

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



SQL scripts are provided in the appropriate bundle described below. The SQL scripts should be run on your database using the appropriate tool, for example SQLPlus for Oracle or SQLServer Management Studio for SQL Server. [Running Scripts on Oracle](#) on page 17 includes information for running scripts on Oracle. [Running Scripts on SQL Server](#) on page 17 includes information on running scripts on SQL Server.

- 1 Verify that the following prerequisites have been met:
 - The two application servers that will run NA have working hostnames. Note the following:
 - Each application server should have a high-speed connection to the database server.
 - For application server hardware and operating system requirements, see the *NA Support Matrix*.
 - It is recommended that the host names of the database server and both application servers are in the `hosts` file on each application server. This file is located as follows:
 - *Windows*: `<Drive>:\Windows\System32\drivers\etc\hosts`
 - *UNIX*[®]: `/etc/hosts`
 - The database server and all application servers are set to use the same time and timezone. It is recommended to synchronize the servers with an external time service.
- 2 *Optional*. If you do not have an existing SQL Server or Oracle NA database server installed, perform a clean NA install (NA1) for either the SQL Server or Oracle database server. To review the NA installation steps, see the *NA Installation and Upgrade Guide* in the `docs` folder of the installation DVD.
- 3 Record the IP addresses and the database instance name and credentials for the NA Cores and Database server.
- 4 Locate the `OracleHorizontalScalabilityBundle.zip` file or the `SQLServerHorizontalScalabilityBundle.zip` file in the `oracle_horizontal_scalability` folder or the `sql_server_horizontal_scalability` folder on the NA Multimaster and Horizontal Scalability DVD. Be sure to use the scripts for the appropriate database type.
- 5 Stop the NA Cores/daemons (NA1).

- 6 Extract either the `OracleInitialSetup.sql` file or `SQLServerInitialSetup.sql` file from the appropriate replication script bundle. You will need to edit the file to replace certain variables. The variables to be replaced are located in the script between comment lines `BEGIN VARIABLE REPLACEMENT SECTION` and `END VARIABLE REPLACEMENT SECTION`. Be sure to completely replace the variables, including replacing the angle brackets (`<>`), as follows:

`<REPLACEME_DATABASE_NAME>` Use the SQL Server database name or Oracle SID.

`<REPLACEME_DATABASE_SERVER_NAME_OR_IP>` Use the DNS hostname or static IP address of your database server.

`<REPLACEME_CORE_SERVER_NAME_OR_IP_1>` Use the DNS hostname or static IP address of the first NA Core.

`<REPLACEME_CORE_SERVER_NAME_OR_IP_2>` Use the DNS hostname or static IP address of the second NA Core.



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



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

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

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

- 7 Run the appropriate setup script, for example `SQLServerInitialSetup.sql` or `OracleInitialSetup.sql`, that you edited in [step 6](#). (See [Running Scripts on Oracle](#) on page 17 or [Running Scripts on SQL Server](#) on page 17 for information.)
- 8 On the second NA Core, install NA (NA2). When prompted, choose **Use Existing Database**. The database name/SID is the same one you used in [step 6](#).
- 9 Stop the second NA Core Management Engine (NA2). For information on how to stop the NA Management Engine, see the *NA User Guide*.
- 10 On all of the NA Cores, edit and copy the `distributed.rcx` file from the appropriate `HorizontalScalabilityBundle.zip` to `C:\NA\jre\` (Windows) or equivalent Unix directory, for example `/opt/NAfiles/jre/`.

Make sure the following options are set accordingly for Oracle.

`<options>`

...

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

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

`<option name="distributed/type">Oracle</option>`

`<option name="distributed/cache_delay">30</option> <!-- in seconds -->`

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

...

`</options>`

The following options are for SQL Server:

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



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

- 11 Restart all the NA Cores/daemons.
- 12 Verify your installation is working correctly. See [Verifying Installation and Setup](#) on page 19).

Scripts

Running Scripts on Oracle

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

- 1 Execute the following command from the appropriate shell prompt:


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



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

Running Scripts on SQL Server

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

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

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

Where the variables should be substituted as follows:

`<SERVER_NAME>` The name of your database server.

`<USER>` The user that the NA Core uses to access the database.

`<PASSWORD>` The user's password.

`<DATABASE_NAME>` The database name NA uses.

<SCRIPT_NAME> The name of the script you want to run.

Adding Additional NA Cores

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

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

<REPLACEME_DATABASE_NAME> Use the SQL Server database name or Oracle SID.

<REPLACEME_DATABASE_SERVER_NAME_OR_IP> Use the DNS hostname or static IP address of your database server.

<REPLACEME_ADDED_CORE_SERVER_NAME_OR_IP> Use the DNS hostname or static IP address of the NA Core you are adding.



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

- 6 Run the appropriate add server script, for example `OracleAddServer.sql` or `SQLServerAddServer.sql` that you edited in [step 5](#). (See [Running Scripts on Oracle](#) on page 17 or [Running Scripts on SQL Server](#) on page 17 for information.)
- 7 On the new NA Core, install NA (NA_n). When prompted, choose **Use Existing Database**. The database name/SID is the same one you used in [step 5](#) (i.e. `dbname1`).
- 8 Stop the new NA Core Management Engine (NA_n).
- 9 On the new NA Core, edit and copy the `distributed.rcx` file from the appropriate `HorizontalScalabilityBundle.zip` to `C:\NA\jre\` (Windows) or equivalent Unix directory, for example `/opt/NAfiles/jre/`.

Make sure the following options are set accordingly for Oracle.

<options>

...

<option name="distributed/enabled">true</option>

<option name="distributed/replication_enabled">>false</option>

<option name="distributed/type">Oracle</option>

```

    <option name="distributed/cache_delay">30</option> <!-- in seconds -->
    <option name="distributed/bind_tasks_to_core">true</option>
...
</options>

```

The following options are for SQL Server:

```

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

```



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

- 10 Restart all the NA Cores/daemons.

Verifying Installation and Setup

To verify installation and setup, do the following:

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


```
SELECT * FROM RN_CORE;
```

You can also view the Edit Core page in NA to review each NA Core's settings. See [Edit Core Page](#) on page 25 for information.
- 2 Edit an object on one NA server (for example, a Comments field for a device).
- 3 Verify that the updated comment exists on the second NA server.
- 4 Check the status of the Distributed Monitor in the NA UI to ensure that no problems are being reported. It could take up to five minutes for this monitor to initially run.

Configuring Topology After Setup

The following general topologies are supported:

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

See [Table 1](#).

Table 1 Topology Configurations

| Option | Topology 1 Value | Topology 2 Value | Topology 3 Value | Topology 4 Value |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|
| "distributed/bind_tasks_to_core" (in distributed.rcx) | true | false | true | true |
| Allow local task running: Admin Settings > Server > Performance Tuning Note: In a Distributed System, this option notifies the entire NA Mesh that there can be one or more NA Cores in the NA Mesh that locally run all tasks scheduled on those NA Cores. This is an NA Mesh-wide global option. If you need to change this option, you must restart all of the NA Management Engines after making the change. | Unchecked (false) | Unchecked (false) | Checked (true) | Checked (true) |
| Run tasks locally: Admin Settings > Server > Performance Tuning Note: In a Distributed System, this option notifies the NA Core to locally run all tasks created on it. This option will only take effect if the NA Mesh-wide option is set and is NA Core-specific. | Unchecked | Unchecked | <ul style="list-style-type: none"> • Checked (on the server where tasks should run locally) • Unchecked (on the other servers) | Unchecked |

Post Installation Setup

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

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

Uninstall Procedures

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



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

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

Oracle Script

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

SQL Server Script

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



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

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

Upgrading Horizontal Scalability

To upgrade NA in a Horizontal Scalability environment from version 7.5x, 7.6x, 9.00, or 9.10 to version 9.20, follow these steps:

- 1 Upgrade each NA application server as described in the “Upgrading to NA 9.20 from a Different System” chapter or the “Upgrading to NA 9.20 on the Same System” chapter of the *NA Installation and Upgrade Guide*, as appropriate.

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

Note the following:

- To prevent database access, on each NA application server, after the NA 9.20 Service Pack Installer runs, stop all NA services.
 - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Stop**:

TrueControl ManagementEngine

TrueControl FTP Server**TrueControl SWIM Server****TrueControl Syslog Server****TrueControl TFTP Server**

- *UNIX*: Run the following command:

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

- 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 application server.
- If you are moving NA to different application servers, copy the NA-specific files (the NA directory and, on UNIX, the startup files) from one existing NA application server to only one new NA application server.

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

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

Retain the permissions on the files.

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

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

Retain the permissions on the files.

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

- 2 On each NA application server, restart all NA services.
 - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Restart**:
 - **TrueControl ManagementEngine**
 - **TrueControl FTP Server**
 - **TrueControl SWIM Server**
 - **TrueControl Syslog Server**
 - **TrueControl TFTP Server**
 - *UNIX*: Run the following command:

```
/etc/init.d/truecontrol restart
```
- 3 Complete the upgrade procedure from the step to run the Data Pruning task.

3 System Administration

Managing Your Horizontal Scalability NA Mesh

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

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

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

NA-Generated Events

By default, NA generates system events. Event rules can alert you to certain error conditions requiring attention. You should examine the default “Distributed System” event rule to ensure all of the events are included in the event rule and that the event rule is configured to send the email notification to the appropriate administrator.

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

RMI Error

Event format:

Local Core: <hostname>

Remote Core: <hostname>

Error: <Exception text>

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

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

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

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



This solution is for static IP environments only.

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


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


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

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

Using the NA Distributed System Pages

When you install the Distributed System software, the NA user interface includes specific Distributed System pages to help you monitor and administer the system. The following options are included under the **Admin > Distributed** option:

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

Distributed Monitor Results Page

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

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

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

- RMI Connections - RMI (Remote Method Invocation) is Java's remote procedure call protocol. The distributed system makes RMI calls between NA servers in the NA Mesh to transfer information about scheduled tasks, system settings, software images, and so on.

- Local NA Core Definition - The local NA Core must be able to determine which entry in the RN_CORE table it is. If the “The local core for this system is undefined” error message is displayed, the CoreHostname property needs to be updated for the NA Core. This can be done using the Edit Core page. See [Edit Core Page](#) on page 25 for information.



This condition could be displayed in the error logs with the following text:
“Fatal error - could not assign local core.”

The CoreHostname value can be either the DNS, etc/hosts value, or an IP address. If you are using 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 RN_CORE table's CoreHostname value for that NA Core.

For detailed information on the Distributed Monitor Results page, see the *NA Multimaster Distributed System on Oracle Guide* or the *NA Multimaster Distributed System on SQL Server Guide*.

Site Reassignment Page

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

To open the Site Reassignment page, on the menu bar under Admin select Distributed and click Site Reassignment. The Site Reassignment opens. You can select NA Cores from the drop-down menu. For detailed information on the Site Reassignment page, see the *NA Multimaster Distributed System on Oracle Guide* or the *NA Multimaster Distributed System on SQL Server Guide*.

List Cores Page

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

To open the List Cores page, on the menu bar under Admin select Distributed and click Core List. The List Cores page opens. For detailed information on the List Cores page, see the *NA Multimaster Distributed System on Oracle Guide* or the *NA Multimaster Distributed System on SQL Server Guide*.

Edit Core Page

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

To open the Edit Core page:

- 1 On the menu bar under Admin select Distributed and click Core List. The List Cores page opens.
- 2 In the Actions column, click the Edit option. The Edit Core page opens.
- 3 Complete the following fields and click the Save Core button.

For detailed information on the Edit Core page, see the *NA Multimaster Distributed System on Oracle Guide* or the *NA Multimaster Distributed System on SQL Server Guide*.

Renew Configuration Options Page

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

To open the Renew Configuration Options page, on the menu bar under Admin select Distributed and click Renew Configuration Options. The Renew Configurations Options page opens. Click the Renew Config Options button to ensure that all options on the NA Cores are in sync with the rest of the NA Mesh.

For detailed information on the Renew Configuration Options page, see the *NA Multimaster Distributed System on Oracle Guide* or the *NA Multimaster Distributed System on SQL Server Guide*.

We appreciate your feedback!

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

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

Product name and version: NA 9.20

Document title: *NA Horizontal Scalability Guide, May 2012*

Feedback: