



# HPE Network Automation

Software Version: 10.30  
Windows® and Linux® operating systems

## Disaster Recovery configuration guide

Document Release Date: June 2017  
Software Release Date: June 2017

  
**Hewlett Packard**  
Enterprise

## Legal Notices

### Warranty

The only warranties for Hewlett Packard Enterprise 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. HPE 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 HPE 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 Notice

© Copyright 2001-2017 Hewlett Packard Enterprise Development LP

### 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.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

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

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

Red Hat® is a registered trademark of Red Hat, Inc. in the United States and other countries.

UNIX® is a registered trademark of The Open Group.

#### **Oracle Technology — Notice of Restricted Rights**

For the full Oracle license text, see the `license-agreements` directory on the NA product DVD.

## 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

NA Disaster Recovery .....	5
Disaster Recovery Arcitecture .....	6
NA Disaster recovery initial setup .....	8
Setting up NA for Disaster Recovery .....	8
Files to Synchronize Across NA Cores .....	12
Verifying the Disaster Recovery Configuration .....	14
Switchover .....	14
Switchback .....	21
Switching Back to the Original Servers in the Primary Location .....	21
Switching Back to Different NA and Database Servers .....	26
Creating a New Disaster Recovery Location .....	34
Disaster Recovery Configuration with Horizontal Scalability over a WAN .....	36
Example Architecture .....	37
Switchover Considerations .....	37
Switchback Considerations .....	38
Send Documentation Feedback .....	40

# NA Disaster Recovery

This guide describes the recommended disaster recovery architecture for Network Automation (NA). This guide describes the procedure for configuring NA for disaster recovery. It also describes the procedures for switching over to the disaster recovery location and switching back from the disaster recovery location.

Disaster recovery planning provides for minimizing the business disruption should a significant event affect an entire data center. Possible uses for the disaster recovery configuration include the following:

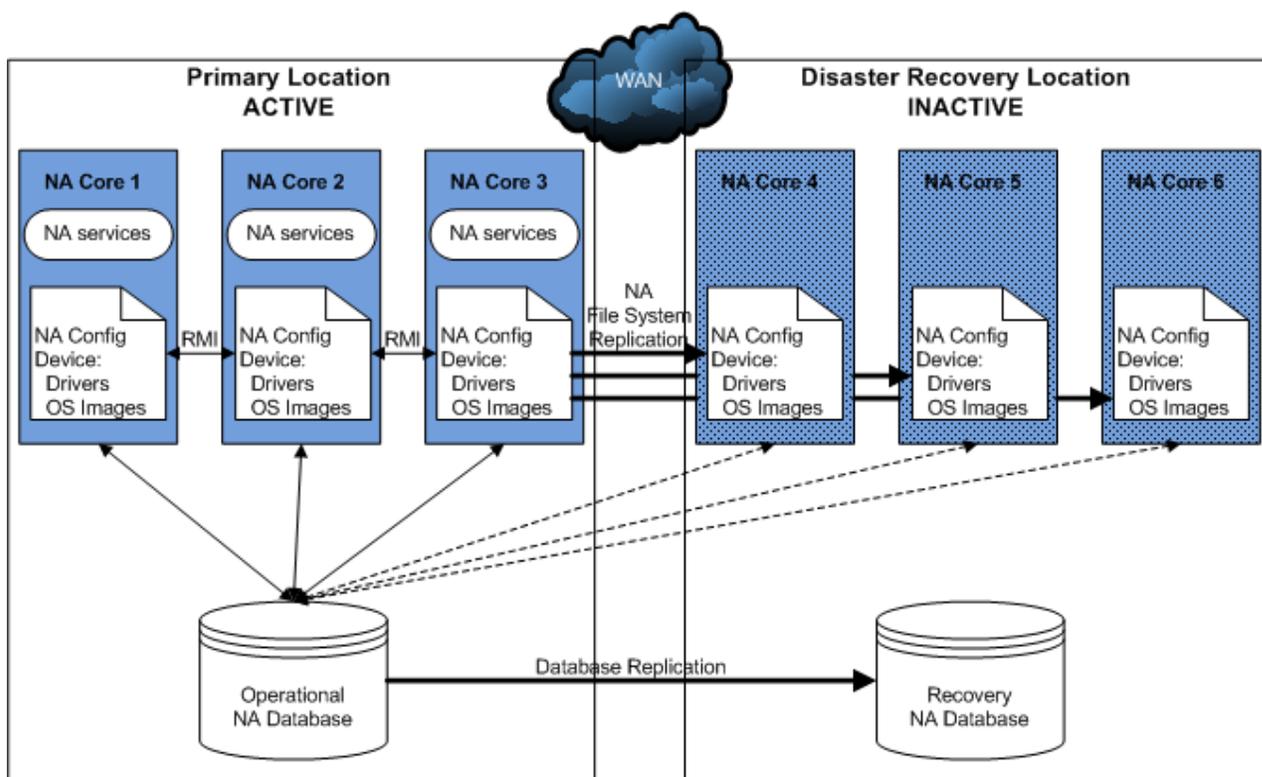
- Unexpected unavailability of a data center due to natural disaster or acts of war. In this case, any lag in data replication results in lost data at the disaster recovery location.
- Anticipated unavailability of a data center due to natural events (for example, a forecasted hurricane), facilities maintenance, or data center movement. In this case, it might be possible to avoid data loss by waiting until the NA database in the disaster recovery location is completely synchronized with the NA database in the primary location before switching over to the disaster recovery location.

Disaster recovery is different from high availability in that with disaster recovery, down time is expected. Generally, disaster recovery configuration includes both of the following processes:

1. Setting up redundant hardware and software at a disaster recovery location that is remote to the primary, operational location.
2. Providing for one-way replication of application data to the disaster recovery location.

# Disaster Recovery Arcitecture

For NA, disaster recovery configuration involves duplicating the NA environment running in a data center in the primary location to a remote data center in the disaster recovery location. The following figure shows this duplication for a three NA core Horizontal Scalability environment. (An NA core is a physical or virtual server on which the NA services and supporting configuration are installed.) Horizontal Scalability provides load sharing, high availability, and fault tolerance. The disaster recovery configuration provides continuity after a disaster.



Note the following:

- NA Horizontal Scalability (and, therefore, NA disaster recovery configuration) is supported with PostgreSQL, Oracle, or Microsoft SQL Server databases. For information about supported versions, see the *Support Matrix*.
- Each location includes one NA database. This database could be implemented as a standalone database server or as a database cluster using a technology such as Oracle Real Application Clusters (RAC).
- At any point in time, only one NA database is actively used by the NA cores. The second database

must be running to receive database updates; however, no NA cores connect to the second database.

The recommended NA disaster recovery scenario includes one-way database replication.

- One to five active NA cores (for Oracle, seven cores if you apply the patch NA\_00051) connect to the operational NA database using Horizontal Scalability. The server for each NA core has a unique IP address and hostname, so the switchover and switchback procedures include updating any configuration that connects to the NA servers. For best performance of the overall solution, it is recommended that each NA server be located in the same data center as the database server to which it is most likely to connect.

When multiple NA cores are active simultaneously, Java remote method invocation (RMI) calls synchronize the NA-specific file systems across the active NA cores. RMI calls also manage schedules for running tasks across the active NA cores.

- During disaster recovery configuration, all NA cores connect to the operational NA database. The NA cores in the disaster recovery location are then set to the inactive state. 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.
- The maximum number of active NA cores in a Horizontal Scalability environment is five (for Oracle, seven cores if you apply the patch NA\_00051).

**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.

During disaster recovery configuration, all nine NA cores might be active at one time; however, no device management occurs on the NA cores in the disaster recovery location.

- If the primary location includes one or more core gateways, the disaster recovery location must include at least one core gateway. The disaster recovery location could include up to one core gateway per NA core. The number of core gateways in the disaster recovery location need not match the number of core gateways in the primary location.
- This guide assumes that satellite gateways are located at facilities other than the primary and disaster recovery locations; therefore, it does not discuss disaster recovery configuration for satellite gateways.
- NA task management in the Horizontal Scalability environment includes the following behavior:
  - If NA tasks for a given device are bound to only one NA core (Horizontal Scalability with the core binding topology, the default behavior), when an NA administrator reassigns the sites from one NA core to another NA core, NA moves the tasks associated with that site to the receiving NA core.

- If all NA tasks for all devices are distributed across all NA cores (Horizontal Scalability with the distributed round robin topology), when an NA administrator sets an NA core to the inactive state, all tasks scheduled to run on that NA core are distributed to the remaining active NA cores.
- This disaster recovery configuration is licensed as follows:
  - One production license for the total number of managed devices for one NA core in the primary location.
  - One non-production license for the total number of managed devices for *each* additional NA core in the primary and disaster recovery locations.
  - One production license for each core gateway in the primary location.
  - One non-production license for each core gateway in the disaster recovery location.
- Some or all of the NA cores in the disaster recovery location can be actively managing the network before a disaster occurs. For more information, see ["Disaster Recovery Configuration with Horizontal Scalability over a WAN" on page 36](#).

Related topics:

## NA Disaster recovery initial setup

You can configure Network Automation (NA) in a disaster recovery setup any time after NA is configured and running satisfactorily in the primary location. The approach described in this document requires NA Horizontal Scalability functionality.

This topic assumes NA database replication in an operation–recovery configuration. The recovery database server contains a copy of the NA database. The replication technology monitors the database transactions on the operational database and periodically replicates them to the recovery database. This configuration requires that the recovery database server be powered on and running continuously. Select a database replication technology appropriate to your database type and business needs.

This guide assumes that the replication technology replicates schema updates on the operational database to the recovery database. In this way, schema updates need be applied to the operational database only.

This topic includes:

- ["Setting up NA for Disaster Recovery" below](#)
- ["Files to Synchronize Across NA Cores" on page 12](#)
- ["Verifying the Disaster Recovery Configuration" on page 14](#)

## Setting up NA for Disaster Recovery

To perform initial setup of an NA disaster recovery configuration, follow this general outline:

- ["Task 1: Prepare the Primary Location for Disaster Recovery Configuration" on the next page](#)
- ["Task 2: Configure Database Replication" on the next page](#)

- ["Task 3: Install and Configure NA in the Disaster Recovery Location" on the next page](#)
- ["Task 4: Finish the Disaster Recovery Configuration" on page 12](#)

### Task 1: Prepare the Primary Location for Disaster Recovery Configuration



1. In the primary location, start with a running NA deployment. This deployment can be a single NA core or a Horizontal Scalability environment containing up to five NA cores (for Oracle, seven if you apply the NA\_00051 patch). This deployment can also include NA satellite functionality.
2. *Optional.* Consider the risk that if the primary location is not accessible during switchover to the disaster recovery location, the NA cores in the primary location cannot be deactivated. In this case, two NA cores (one each in the primary and disaster recovery locations) might run the same task. To mitigate this risk, on each NA server in the primary location, disable automatic starting of the NA services. See [Disable All Services](#).
3. Prepare to stop NA in the primary location.
  - a. Notify users to log out.
  - b. Log on to the NA console for one of the NA cores in the primary location.
  - c. Pause tasks scheduled to start during the disaster recovery configuration process (until [Task 4, step 1](#)). Include time for the currently running tasks to complete. Also include time for database synchronization. For example:
    - i. 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**.
    - ii. On the Task Search Results page, pause each listed task.
  - d. On the Running Tasks page (**Tasks > Running Tasks**), examine the list of running tasks.

If any critical tasks are running, wait for them to complete before continuing with step 4, next.
4. Stop all NA services on all NA cores in the primary location. See [Start, Stop, or Restart All Services](#).

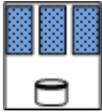
### Task 2: Configure Database Replication



1. If the NA database in the primary location was created for the postgres user (PostgreSQL), SYSTEM user (Oracle), or the SA user (SQL Server), move the NA database to a custom user with the privileges described in the *Install and Upgrade guide*. Do the following:
  - a. Create a new tablespace or database instance dedicated to NA on the database server in the primary location.
  - b. Use database tools to copy the NA schema to the new tablespace or database instance.
2. Use database tools to create a copy of the operational NA database in the disaster recovery location.

	<p>Note the following:</p> <ul style="list-style-type: none"><li>• The NA database user in the disaster recovery location must have the same name and permissions as the NA database user in the primary location.</li><li>• Copy the NA schema tables only.</li></ul> <p>For example, you might follow this process:</p> <ol style="list-style-type: none"><li>a. Install the database software.</li><li>b. Create a database user with the same name and permissions as the NA database user for the operational database in the primary location.</li><li>c. Export the NA database from the operational database in the primary location.</li><li>d. Import the NA database to the recovery database in the disaster recovery location.</li></ol> <ol style="list-style-type: none"><li>3. Configure one-way database replication from the operational database in the primary location to the recovery database in the disaster recovery location. Use a database replication technology appropriate to your database type and business needs. Follow the documentation for that technology.</li><li>4. Set up a schedule for performing regular trimming of the replication log files.</li></ol>
--	--

### Task 3: Install and Configure NA in the Disaster Recovery Location

	<ol style="list-style-type: none"><li>1. In the disaster recovery location, install the NA cores as additional NA cores connected through Horizontal Scalability to the operational database in the primary location. Make sure that the total number of active and inactive cores in the disaster recovery configuration does not exceed nine. To use more than nine cores in this environment, contact HPE Software Support. If you reached this step from a switchback procedure, complete the script edits to account for having removed one or more NA cores from the Horizontal Scalability environment. If you reached this step from a switchback procedure, complete the script edits to account for having removed one or more NA cores from the Horizontal Scalability environment. Connecting the NA cores in the disaster recovery location to the NA database in the primary location ensures that the databases remain synchronized. These NA cores will be stopped later in this procedure.</li><li>2. <i>Optional.</i> Perform any NA core-specific tuning. Restart the NA services as needed.</li><li>3. <i>Optional.</i> Configure the managed devices (or the syslog relay) to send syslog messages to one NA core in the disaster recovery location.</li><li>4. If the primary location includes NA Satellite functionality, do the following:<ol style="list-style-type: none"><li>a. In the disaster recovery location install one or more core gateways. Install at least one core gateway in the disaster recovery location to continue communication with the existing gateway mesh. Optionally install additional core gateways, up to one core gateway per NA core. During installation, configure each core gateway in the disaster recovery location as follows:<ul style="list-style-type: none"><li>◦ Use the same Gateway Crypto Data file as for the core gateways in the</li></ul></li></ol></li></ol>
---	---

	<p>primary location.</p> <ul style="list-style-type: none"><li>o Assign the same realm name, typically Default Realm, to each core gateway.</li><li>o After installation, update the <code>opswgw.EgressFilter</code> entry to match the following: <pre>opswgw.EgressFilter=tcp:*:443:127.0.0.1:*,tcp:*:22:NAS:, tcp:*:23:NAS:,tcp:*:513:NAS:,tcp:*:443:NAS:,tcp:*:80:NAS:</pre></li></ul>
	<p>4. b. For each satellite in the gateway mesh, update the satellite configuration to enable communication with a core gateway in the disaster recovery location.</p> <ul style="list-style-type: none"><li>i. Connect to the satellite gateway server as the root user.</li><li>ii. Change to the following directory: <pre>/etc/opt/opsware/opswgw-&lt;gateway_name&gt;</pre></li><li>iii. Back up the <code>opswgw.properties</code> file to a location outside the gateway installation directory.</li><li>iv. In a text editor, such as <code>vi</code>, open the <code>opswgw.properties</code> file.</li><li>v. Add an <code>opswgw.TunnelSrc</code> entry that points to a core gateway in the disaster recovery location.  Configure this secondary connection with a higher route cost so it is used only when the core gateway in the primary location is unavailable. For example: <pre>opswgw.TunnelSrc=&lt;core_gateway1_IP&gt;:2001:100:0: /var/opt/opsware/crypto/opswgw-RemoteGw/opswgw.pem opswgw.TunnelSrc=&lt;core_gateway2_IP&gt;:2001:200:0: /var/opt/opsware/crypto/opswgw-RemoteGw/opswgw.pem</pre></li><li>vi. Restart the satellite gateway. With this configuration, no additional work is needed to enable the core gateways during switchover or switchback. This guide expects that the satellites are remote to the primary and disaster recovery locations. If necessary, set up additional satellites in the gateway mesh for redundancy.</li></ul> <p>5. Deactivate the NA cores in the disaster recovery location.</p> <ul style="list-style-type: none"><li>a. Connect as an NA administrator to the NA proxy on one of the NA cores in the disaster recovery location.</li><li>b. Run the following command: <pre>list core</pre></li><li>c. From the <code>list core</code> command output, determine the core IDs of the new NA cores, and then run the following command for each of these core IDs: <pre>set core status -status inactive -coreid &lt;coreid&gt;</pre></li></ul> <p>6. Stop all NA services on all NA cores in the disaster recovery location. See <a href="#">Start, Stop, or Restart All Services</a>.</p>

## Task 4: Finish the Disaster Recovery Configuration



1. In the primary location, start all NA services on all active NA cores (if they are not already running). See [Start, Stop, or Restart All Services](#).
2. Resume the tasks that were paused in Task 1, step 3.
  - a. Log on to the NA console for one of the NA cores in the primary location.
  - b. On the Search for Task page (**Reports > Search For > Tasks**), for the Task Status field, select Paused.
  - c. On the Task Search Results page, resume each listed task.  
Later instances of periodic tasks remove the paused instances of those tasks.
3. Notify users to resume use of the NA console on the NA cores in the primary location.
4. Configure replication of the NA file system from one of the active NA cores to the inactive NA cores (or an intermediate server) as described in "[Files to Synchronize Across Cores](#)" on page 1.
5. *Optional*. If you chose to synchronize the NA files to an intermediate server, to conserve resources, power down the NA servers in the disaster recovery location.
6. *Optional*. If the database replication technology supports reverse replication from the database in the disaster recovery location to the database in the primary location, prepare for, but *do not enable*, reverse replication.

## Files to Synchronize Across NA Cores

While most NA data is stored in the NA database, some files on the NA core support the NA console and NA functions. A complete disaster recovery scenario must include replication of these files. [Table 1](#) lists the files to consider for replication among the NA cores.

Configure file replication using tools appropriate to your environment. Maintain file ownership and permissions during replication.

Set up a regularly scheduled server-level job to copy the files listed in [Table 1](#) from one of the active NA cores (in the primary location) to all of the inactive NA cores in the disaster recovery location. This copy could be initiated by an active NA core (push) or by each inactive NA core (pull).

Alternatively, the script might copy files from an active NA core to an intermediate server. This approach is useful if you want to leave the inactive NA cores powered off until they are needed. In this case, the procedure for switching over to the disaster recovery location includes copying the files from the intermediate server to each inactive NA core.

To set core-specific values of options in the synchronized files, on each core, create a file named `adjustable_local_options.rcx` that contains the core-specific settings.

### Example file replication script

For example, the following script pulls files from an active NA core. It uses the `rsync` command, which compares the versions of a file on each server and copies only those files that have changed. This script would be located on each inactive NA core. It copies files from the active NA core whose core ID is `core1`.

```
C1=core1
rsync -avz $C1:/opt/NA/jre/site_options.rcx /opt/NA/jre
rsync -avz $C1:/opt/NA/jre/logging.rcx /opt/NA/jre
rsync -avz $C1:/opt/NA/jre/adjustable_options.rcx /opt/NA/jre
rsync -avz $C1:/opt/NA/jre/distributed.rcx /opt/NA/jre
rsync -avz $C1:/opt/NA/jre/securityfilter_additional_init.rcx /opt/NA/jre
rsync -avz $C1:/opt/NA/server/lib/drivers/ /opt/NA/server/lib/drivers
rsync -avz $C1:/opt/NA/server/images/ /opt/NA/server/images
```

**Table 1: Files to Synchronize Across NA Cores**

Category	Files
<p>RCX files, which are located in the following directory:</p> <ul style="list-style-type: none"> <li>• <i>Windows</i>: &lt;NA_HOME&gt;\jre</li> <li>• <i>Linux</i>: &lt;NA_HOME&gt;/jre</li> </ul>	<p>Specifically, at least the following files:</p> <ul style="list-style-type: none"> <li>• <code>site_options.rcx</code> (NA server behavior)</li> <li>• <code>logging.rcx</code> (NA logging levels)</li> <li>• <code>adjustable_options.rcx</code> (Customer-specified configuration options)</li> <li>• <code>distributed.rcx</code> (Distribution settings)</li> <li>• <code>securityfilter_additional_init.rcx</code> (Customer-specified filters for URL strings)</li> </ul> <p>Also include other RCX files that have been customized.</p> <p>Do <i>not</i> include the <code>appserver.rcx</code> file, which contains paths to the local system. If this file has been modified, copy the changed blocks to the <code>adjustable_options.rcx</code> file for synchronization across all NA cores or to the <code>adjustable_local_options.rcx</code> file for one NA core.</p>
<p>HPE-developed device drivers (*.rdp), which are located in a directory as specified by the <code>driver/dir</code> option in the <code>site_options.rcx</code> file, typically:</p> <ul style="list-style-type: none"> <li>• <i>Windows</i>: &lt;NA_HOME&gt;\server\lib\drivers</li> <li>• <i>Linux</i>: &lt;NA_HOME&gt;/server/lib/drivers</li> </ul>	<p>Synchronize all files in the <code>drivers</code> directory.</p>
<p>Device drivers developed outside of HPE, which are located in a directory as specified by the <code>driver/extension/dir</code> option in the <code>site_options.rcx</code> file.</p>	<p>Synchronize all files in the identified directory.</p>
<p>Device operating system images, which are</p>	<p>Synchronize all files in the <code>images</code> directory.</p>

**Table 1: Files to Synchronize Across NA Cores, continued**

Category	Files
<p>located in a directory as specified by the <code>deploy/repository/root</code> option in the <code>site_options.rcx</code> file, typically:</p> <ul style="list-style-type: none"><li>• <i>Windows</i>: <code>&lt;NA_HOME&gt;\server\images</code></li><li>• <i>Linux</i>: <code>&lt;NA_HOME&gt;/server/images</code></li></ul> <p>This directory is created with the creation of the first image set in NA.</p>	

## Verifying the Disaster Recovery Configuration

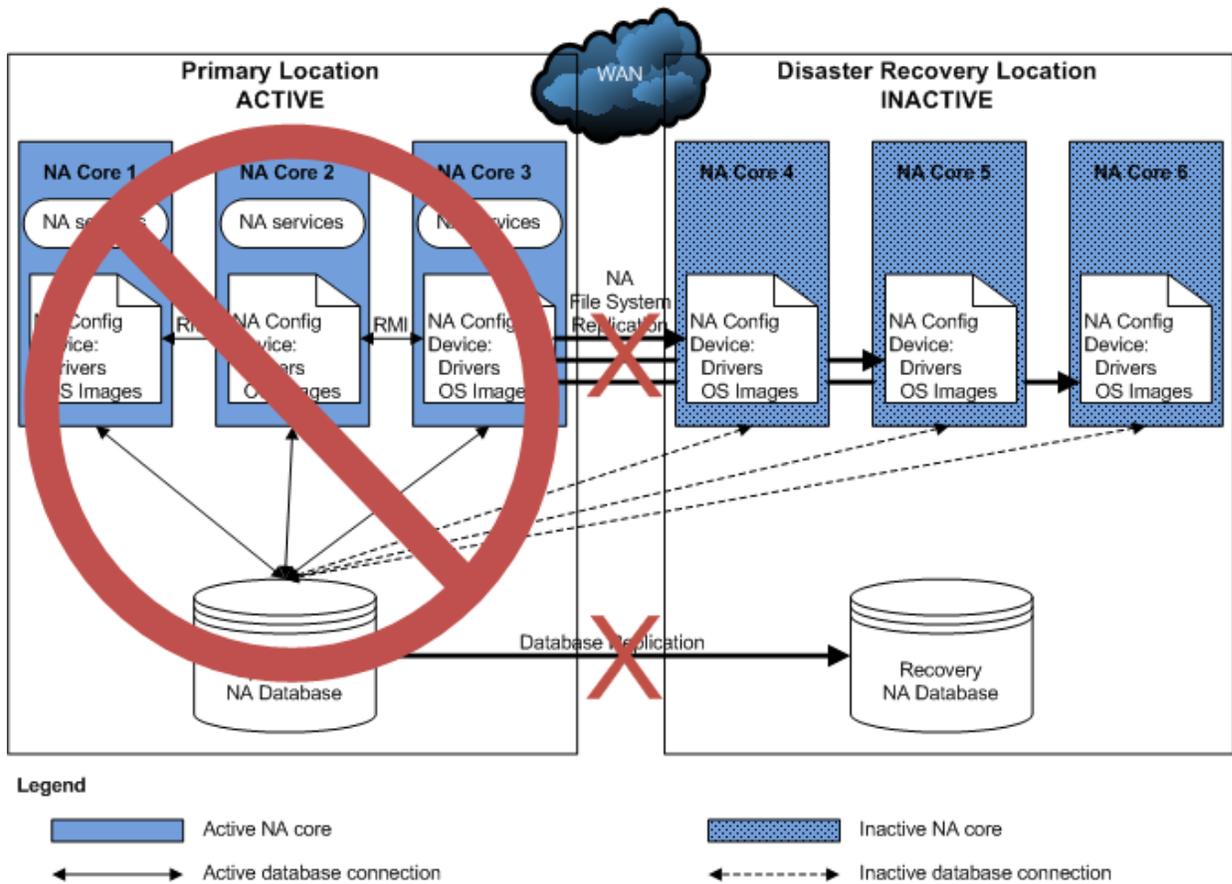
To verify the initial setup of an NA disaster recovery configuration, the database administrator (DBA) can follow this general outline:

1. Verify that database replication works correctly.  
Use database tools to confirm that the numbers of tables and records in the two NA databases are the same.
2. Check the database replication logs.
  - Are there any replication errors?
  - Is there a problem with any of the NA tables?
  - Are there errors regarding the primary key?
3. Examine the replication lag, which is the time difference between when a transaction is recorded in the primary and disaster recovery database.  
If the lag is unacceptably large, tune database replication. For information, see the documentation for your database replication technology.  
For Oracle GoldenGate, consider tuning the `TCPBUFSIZE`, `TCPFLUSHBYTES`, and `COMPRESS` arguments to the `RMTHOST` parameter.
4. Verify NA file system replication by comparing the file sizes and timestamps in the primary and disaster recovery locations.
5. Verify NA core status on the List Core page () or by using the `list core` command on the NA proxy.

## Switchover

When the primary location becomes unavailable, an administrator can follow the procedure described in this section to switch use of Network Automation (NA) over to the disaster recovery location. In the case of an unplanned disaster event, NA will be unavailable until switchover is complete and the most recent database updates might be lost. In the case of a planned unavailability, replication can be fully completed before switchover begins, and NA downtime can be very short with no data loss.

The following figure shows the state of the disaster recovery configuration immediately after an event has occurred. The primary location is unavailable, and the disaster recovery has not yet gone live.



If the primary location is *not* accessible, to switch over from the primary location to the disaster recovery location, complete the following tasks in order:

- "Task 1: Plan to Disable NA in the Primary Location" on the next page
- "Task 3: Enable Use of the Database in the Disaster Recovery Location" on page 17
- "Task 4: Enable NA in the Disaster Recovery Location" on page 18
- "Task 5: Finish Switchover" on page 20

If the primary location is accessible, to switch over from the primary location to the disaster recovery location, complete the following tasks in order:

- "Task 2: Disable NA in the Primary Location" on the next page
- "Task 3: Enable Use of the Database in the Disaster Recovery Location" on page 17
- "Task 4: Enable NA in the Disaster Recovery Location" on page 18
- "Task 5: Finish Switchover" on page 20

## Task 1: Plan to Disable NA in the Primary Location



If the primary location is not currently accessible, make plans to disable NA functionality in the primary location as soon as that location becomes accessible. These plans might include any or all of the following:

- Disabling automatic starting of the NA services. See [Disable All Services](#).
- Physical changes to the NA server (for example, disconnecting the power source or the network cable).

Continue with [Task 3](#).

## Task 2: Disable NA in the Primary Location



If the primary location is accessible, disable NA by completing the steps in this task. If you anticipate losing connectivity to the primary location, complete as many of these steps as possible while connectivity remains. If you are unable to complete this task before losing connectivity to the primary location, also consider the information in [Task 1: Plan to Disable NA in the Primary Location](#).

1. On each NA server in the primary location, disable automatic starting of the NA services. See [Disable All Services](#).
2. Prepare to stop NA in the primary location.
  - a. Notify users to log out.
  - b. Log on to the NA console for one of the NA cores in the primary location.
  - c. Pause tasks scheduled to start during the switchover process (until [Task 4, step 5](#)). Include time for the currently running tasks to complete. For example:
    - i. On the Search for Task page (**Reports > Search For > Tasks**), for the Schedule Date field, set since to **Now** and until to **2 hours later**.
    - ii. On the Task Search Results page, pause each listed task.
  - d. On the Running Tasks page (**Tasks > Running Tasks**), examine the list of running tasks.

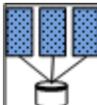
If any critical tasks are running, wait for them to complete before continuing with [step 3](#), next.
3. Stop all NA services on all NA cores in the primary location. See [Start, Stop, or Restart All Services](#).

**Caution:** Ensure that the NA services on the NA cores in the primary location remain stopped until directed otherwise in the switchback procedure.

4. Wait for NA file system replication from the primary location to complete. Verify completeness by comparing the file sizes and timestamps in the primary and

	<p>disaster recovery locations.</p> <ol style="list-style-type: none"><li>5. Wait for all database updates to replicate from the database in the primary location to the database in the disaster recovery location.</li><li>6. On the primary location database server, disable database replication to the database in the disaster recovery location.</li></ol>
--	--

## Task 3: Enable Use of the Database in the Disaster Recovery Location

	<ol style="list-style-type: none"><li>1. On the disaster recovery location database server, disable database replication from the database in the primary location.</li></ol>
	<ol style="list-style-type: none"><li>2. <i>Optional.</i> In the case of a planned switchover, if the database replication technology supports reverse replication from the database in the disaster recovery location to the database in the primary location, enable reverse replication.</li></ol> <div style="background-color: #e0e0e0; padding: 10px;"><p><b>Tip:</b> If the down time of the primary location is expected to be less than the time in which the database transaction logs fill, reverse replication can be a good way to prepare the database in the primary location for switchover. If the database transaction logs fill before the primary database becomes available, reverse replication becomes ineffective. In this case, you will need to do a complete database copy as part of switching back to the primary location.</p></div> <ol style="list-style-type: none"><li>3. If necessary, power on the NA servers in the disaster recovery location.</li><li>4. If necessary, stop all NA services on all NA cores in the disaster recovery location. See <a href="#">Start, Stop, or Restart All Services</a>.</li></ol>
	<ol style="list-style-type: none"><li>5. Connect the NA cores in the disaster recovery location to the local database (the database in the disaster recovery location). On each NA server in the disaster recovery location, in a text editor such as WordPad or vi, edit following file:<ul style="list-style-type: none"><li>• <i>Windows:</i> &lt;NA_HOME&gt;\server\ext\jboss\server\default\deploy\db_ds.xml</li><li>• <i>Linux:</i> &lt;NA_HOME&gt;/server/ext/jboss/server/default/deploy/db_ds.xml</li></ul>This file contains <i>two</i> lines defining the JdbcUrl attribute. For example:<ul style="list-style-type: none"><li>• PostgreSQL: &lt;attribute name="JdbcUrl"&gt;jdbc:postgres:jdbc:postgresql://hostname:5432/na:1521:nadb&lt;/attribute&gt;</li><li>• Oracle: &lt;attribute name="JdbcUrl"&gt;jdbc:oracle:thin:@db.example.com:1521:nadb&lt;/attribute&gt;</li></ul></li></ol>

- SQL Server: <attribute name="JdbcUrl">jdbc:sqlserver://db.example.com:1433;DatabaseName=NA;SendStringParametersAsUnicode=false</attribute>

For each JdbcUrl attribute, replace the database server name (db.example.com in the example) with the fully-qualified domain name or IP address of the database server in the disaster recovery location.

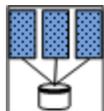
For PostgreSQL Server, also replace the database name (na in the example) with the database name for the NA database in the disaster recovery location.

For Oracle, also replace the database SID or service name (nadb in the example) with the database SID or service name for the NA database in the disaster recovery location.

For SQL Server, also replace the database name (NA in the example) with the database name for the NA database in the disaster recovery location.

This step replaces the connection between these NA cores and the database in the primary location with a connection to the database in the disaster recovery location.

## Task 4: Enable NA in the Disaster Recovery Location



To enable NA in the disaster recovery location, follow these steps:

1. If necessary, copy the NA server files from the intermediate location to the correct locations on the NA servers in the disaster recovery location. Maintain file ownership and permissions.
2. Activate the NA cores in the disaster recovery location.
  - a. Start all NA services on all NA cores in the disaster recovery location. See [Start, Stop, or Restart All Services](#).
  - b. After approximately one minute, connect using telnet or SSH as an NA administrator to the NA proxy on that NA core.
  - c. Run the following command:

```
list core
```

**Tip:** The NA cores in the primary location might be listed as fully functional. This inaccurate information results from no NA core being available to catch the status change and update the NA database.

- d. From the list core command output, determine the core IDs of the NA cores in the disaster recovery location, and then run the following command for each of these core IDs:

```
set core status -status active -coreid <coreid>
```

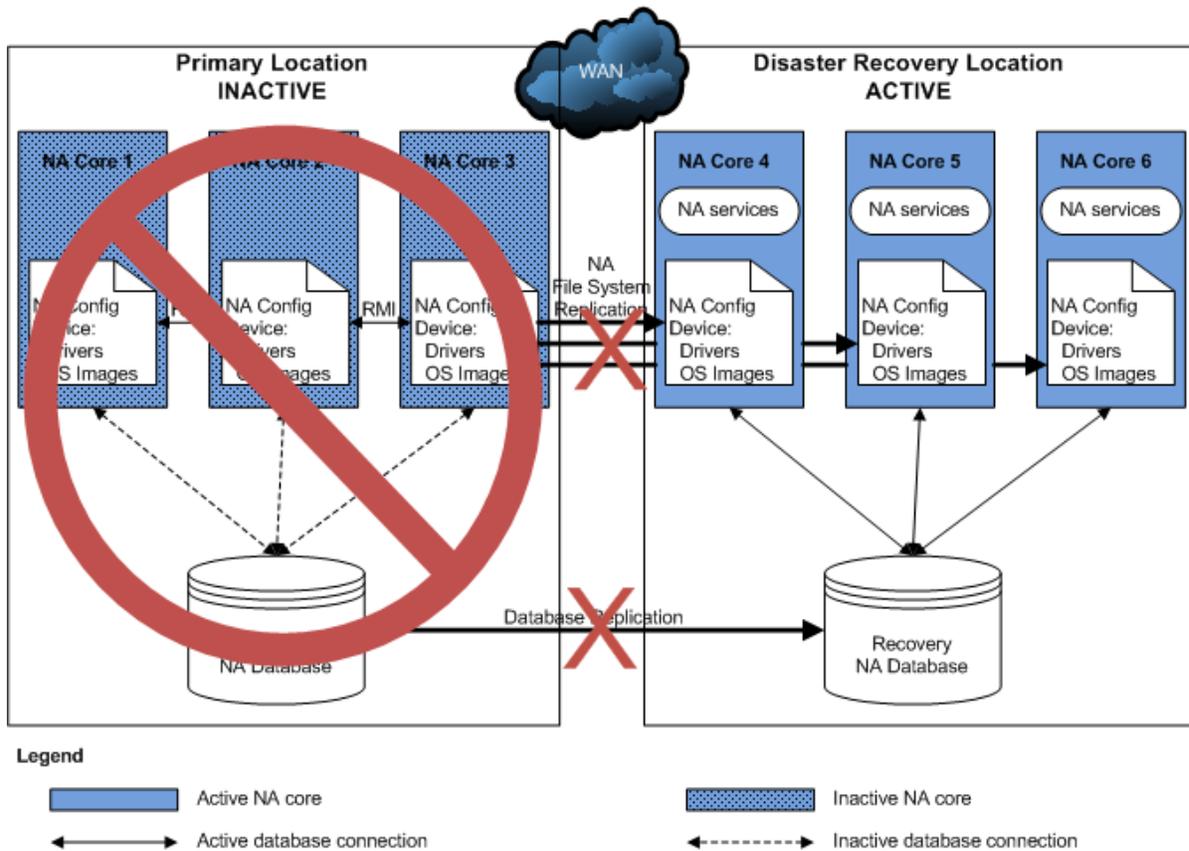
	<p><b>Tip:</b> This step changes the NA core status in the NA database in the disaster recovery location only. Unless reverse replication is running, the NA database in the primary location still shows these cores as inactive.</p> <ol style="list-style-type: none"><li>3. Update the site assignments.<ol style="list-style-type: none"><li>a. Log on to the NA console on an NA core running in the disaster recovery location.</li><li>b. In the NA console, open the <b>Site Reassignment</b> page (<b>Admin &gt; Distributed &gt; Site Reassignment</b>), and then assign all partitions to NA cores in the disaster recovery location.</li><li>c. <i>Optional.</i> On the <b>Partitions</b> page (<b>Admin &gt; Security Partitions</b>), verify the mapping of partitions to NA cores.</li><li>d. <i>Optional.</i> If you enabled reverse replication to the NA database in the primary location, verify that the RN_SITE table in the primary location database shows the same content as the RN_SITE table in the disaster recovery location.</li></ol></li></ol>
	<ol style="list-style-type: none"><li>4. Deactivate the NA cores in the primary location.<ol style="list-style-type: none"><li>a. Connect as an NA administrator to the NA proxy on an NA core running in the disaster recovery location.</li><li>b. Run the following command: <b>list core</b></li><li>c. From the <code>list core</code> command output, determine the core IDs of the NA cores in the primary location, and then run the following command for each of these core IDs: <b>set core status -status inactive -coreid &lt;coreid&gt;</b></li></ol><p><b>Note:</b> This step changes the NA core status in the NA database in the disaster recovery location only. Unless reverse replication is running, the NA database in the primary location still shows these cores as active.</p></li></ol>
	<ol style="list-style-type: none"><li>5. If necessary, resume the tasks that were paused in <a href="#">Task 2, step 2</a>.<ol style="list-style-type: none"><li>a. Log on to the NA console for one of the NA cores in the primary location.</li><li>b. On the Search for Task page (<b>Reports &gt; Search For &gt; Tasks</b>), for the Task Status field, select <b>Paused</b>.</li><li>c. On the Task Search Results page, resume each listed task.</li></ol><p><b>Tip:</b> Later instances of periodic tasks remove the paused instances of those tasks.</p></li></ol>

## Task 5: Finish Switchover



1. Verify the configuration as described in [See "Verifying the Disaster Recovery Configuration"](#).
2. Configure any applications that integrate with NA to connect to the working NA cores in the disaster recovery location.
3. Notify users to connect to the NA console on the NA cores in the disaster recovery location.
4. *Optional.* Configure the managed devices (or the syslog relay) to send syslog messages to one NA core in the disaster recovery location.  
Without syslog messages, NA will detect configuration changes at the next scheduled snapshot.
5. Verify that one NA core in the disaster recovery location is configured to perform dynamic device group membership calculations. For more information, see "Ensure Continuity of Dynamic Device Group Calculation" the *Horizontal Scalability guide*.

The following figure shows the results of switching over to the disaster recovery location.



# Switchback

Switchback involves synchronizing application data from the disaster recovery location to the primary location. Switchback can be scheduled for a time with the least impact.

This section describes the following switchback scenarios:

- ["Switching Back to the Original Servers in the Primary Location" below](#)
- ["Switching Back to Different NA and Database Servers" on page 26](#)
- ["Creating a New Disaster Recovery Location" on page 34](#)

## Switching Back to the Original Servers in the Primary Location

When the NA servers are available, switching back from the disaster recovery location to the primary location involves the following general process:

- ["Task 1: Disable NA in the Disaster Recovery Location" below](#)
- ["Task 2: Enable the Use of the Database in the Primary Location" on the next page](#)
- ["Task 3: Enable NA in the Primary Location" on page 23](#)
- ["Task 4: Finish Switchback" on page 25](#)

## Task 1: Disable NA in the Disaster Recovery Location



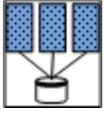
This task assumes that all NA services are stopped on all NA cores in the primary location.

To disable NA in the disaster recovery location, follow these steps:

1. Prepare to stop NA in the disaster recovery location. Do the following:
  - a. Notify users to log out.
  - b. Log on to the NA console for one of the NA cores in the disaster recovery location.
  - c. Pause tasks scheduled to start during the switchback process ([Task 3, step 5](#)). Include time for the currently running tasks to complete. Also include time for database synchronization. For example:
    - i. 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**.
    - ii. On the Task Search Results page, pause each listed task.
  - d. On the Running Tasks page (**Tasks > Running Tasks**), examine the list of

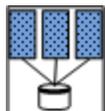
	<p>running tasks.</p> <p>If any critical tasks are running, wait for them to complete before continuing with <a href="#">step 2</a>, next.</p> <p>2. Stop all NA services on all NA cores in the disaster recovery location. See <a href="#">Start, Stop, or Restart All Services</a>.</p>
--	--

## Task 2: Enable the Use of the Database in the Primary Location

	<p>1. Synchronize the database in the primary location with the database in the disaster recovery location. Possible approaches include the following:</p> <ul style="list-style-type: none"> <li>• Use database tools to copy the NA database from the disaster recovery location to the primary location. Copy the NA schema only.</li> </ul> <p>For example, you might follow this process:</p> <ol style="list-style-type: none"> <li>i. Export the NA database from the recovery database server in the disaster recovery location.</li> <li>ii. Wipe the NA database from the database server in the primary location.</li> <li>iii. Import the NA database to the database server in the primary location.</li> </ol>
	<p>1. • If reverse replication from the database in the disaster recovery location to the database in the primary location is running, analyze the reverse replication transaction logs.</p> <p>If the transaction logs have overflowed, reverse replication becomes ineffective. In this case, do the following:</p> <ol style="list-style-type: none"> <li>i. Disable reverse replication to the database in the primary location.</li> <li>ii. Use database tools to copy the NA database from the disaster recovery location to the primary location. Copy the NA schema only.</li> </ol> <p>If the transaction logs are within bounds, do the following:</p> <ol style="list-style-type: none"> <li>i. Wait for all database updates to replicate to the database in the primary location.</li> <li>ii. After replication is complete, disable reverse replication to the database in the primary location.</li> </ol>
	<p>2. Re-enable database replication from the primary location to the disaster recovery location.</p>
	<p>3. Connect the NA cores in the disaster recovery location to the database in the primary location.</p> <p>On each NA server in the disaster recovery location, in a text editor, edit following file:</p> <ul style="list-style-type: none"> <li>• <i>Windows</i>: &lt;NA_HOME&gt;\server\ext\jboss\server\default\deploy\db_ds.xml</li> </ul>

	<ul style="list-style-type: none"><li>• <i>Linux</i>: &lt;NA_HOME&gt;/server/ext/jboss/server/default/deploy/db_ds.xml</li></ul> <p>This file contains <i>two</i> lines defining the JdbcUrl attribute. For example:</p> <ul style="list-style-type: none"><li>• PostgreSQL: &lt;attribute name="JdbcUrl"&gt;jdbc:postgres:jdbc:postgresql://hostname:5432/na:1521:nadb&lt;/attribute&gt;</li><li>• Oracle: &lt;attribute name="JdbcUrl"&gt;jdbc:oracle:thin:@db.example.com:1521:nadb&lt;/attribute&gt;</li><li>• SQL Server: &lt;attribute name="JdbcUrl"&gt;jdbc:sqlserver://db.example.com:1433;DatabaseName=NA;SendStringParametersAsUnicode=false&lt;/attribute&gt;</li></ul> <p>For each JdbcUrl attribute, replace the database server name (db.example.com in the example) with the fully-qualified domain name or IP address of the database server in the primary location.</p> <p>For PostgreSQL Server, also replace the database name (na in the example) with the database name for the NA database in the primary location.</p> <p>For Oracle, also replace the database SID or service name (nadb in the example) with the database SID or service name for the NA database in the primary location.</p> <p>For SQL Server, also replace the database name (NA in the example) with the database name for the NA database in the primary location.</p> <p>This step replaces the connection between these NA cores and the database in the disaster recovery location with a connection to the database in the primary location.</p>
	<ol style="list-style-type: none"><li>4. Ensure that the NA servers in the primary location are powered on.</li><li>5. If necessary, stop all NA services on all NA cores in the primary location. See <a href="#">Start, Stop, or Restart All Services</a>.</li><li>6. Verify that the NA cores in the primary location are connected to the local database (the database in the primary location) as described in step 3.</li></ol>

## Task 3: Enable NA in the Primary Location



To enable NA in the primary location, follow these steps:

1. If the configuration of the NA cores in the disaster recovery location has changed since switchover, update the NA core files on each NA server in the primary location. Maintain file ownership and permissions. See [Table 1](#). You will restart the NA services in [step 5](#). You do not need to do so now.
2. Activate the NA cores in the primary location.
  - a. Start all NA services on all NA cores in the primary location. See [Start, Stop, or Restart All Services](#).
  - b. After approximately one minute, connect using telnet or SSH as an NA

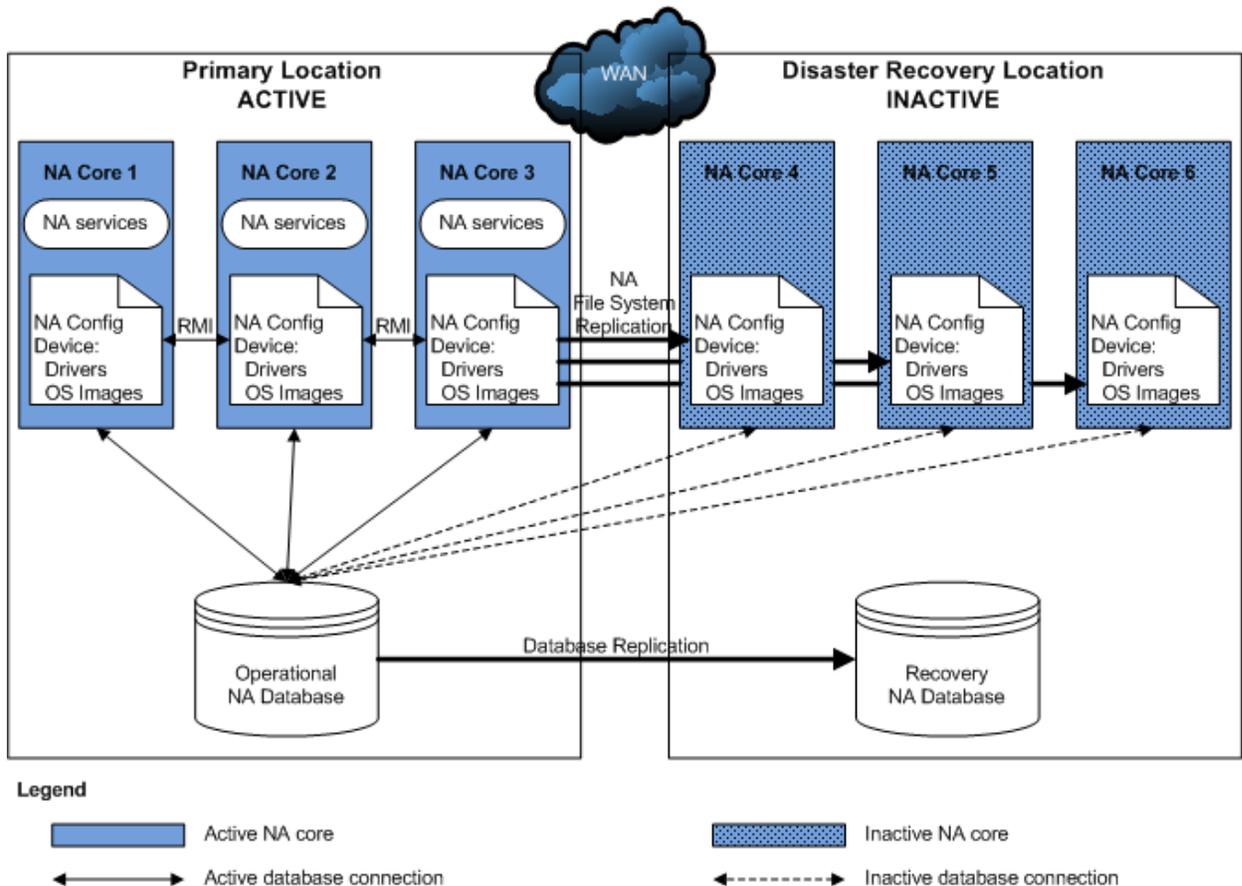
	<p>administrator to the NA proxy on that NA core.</p> <ol style="list-style-type: none"><li>c. Run the following command: <b>list core</b></li><li>d. From the <code>list core</code> command output, determine the core IDs of the NA cores in the primary location, and then run the following command for each of these core IDs: <b>set core status -status active -coreid &lt;coreid&gt;</b></li></ol> <p>This step changes the NA core status in the NA database in the primary location. Because replication is running, the NA database in the disaster recovery location also shows these cores as active.</p>
	<ol style="list-style-type: none"><li>3. Update the site assignments.<ol style="list-style-type: none"><li>a. Log on to the NA console on an NA core running in the primary location.</li><li>b. Update the site assignments. In the NA console, open the <b>Site Reassignment</b> page (<b>Admin &gt; Distributed &gt; Site Reassignment</b>), and then assign all partitions to NA cores in the primary location.</li><li>c. <i>Optional.</i> On the <b>Partitions</b> page (<b>Admin &gt; Security Partitions</b>), verify the mapping of partitions to NA cores.</li></ol></li><li>4. Deactivate the NA cores in the disaster recovery location.<ol style="list-style-type: none"><li>a. Connect as an NA administrator to the NA proxy on an NA core running in the primary location.</li><li>b. Run the following command: <b>list core</b></li><li>c. From the <code>list core</code> command output, determine the core IDs of the NA cores in the disaster recovery location, and then run the following command for each of these core IDs: <b>set core status -status inactive -coreid &lt;coreid&gt;</b></li></ol><p>This step changes the NA core status in the NA database in the primary location. Because replication is running, the NA database in the disaster recovery location also shows these cores as inactive.</p></li></ol>
	<p>If necessary resume the tasks that were paused in <a href="#">Task 1, step 1</a>.</p> <ol style="list-style-type: none"><li>a. Log on to the NA console for one of the NA cores in the primary location.</li><li>b. On the Search for Task page (<b>Reports &gt; Search For &gt; Tasks</b>), for the Task Status field, select Paused.</li></ol> <p>On the Task Search Results page, resume each listed task</p> <ol style="list-style-type: none"><li>5. Later instances of periodic tasks remove the paused instances of those tasks.</li></ol>

## Task 4: Finish Switchback



1. Configure any applications that integrate with NA to connect to the working NA cores in the primary location.
2. Notify users to connect to the NA console on the NA cores in the primary location.
3. Re-enable file system replication from the primary location to the disaster recovery location.
4. *Optional.* If you chose to synchronize the NA files to an intermediate server, to conserve resources in the disaster recovery location, power down the NA servers in the disaster recovery location.
5. *Optional.* Configure the managed devices (or the syslog relay) to send syslog messages to one NA core in the primary location and to no NA cores in the disaster recovery location.
6. Verify that one NA core in the disaster recovery location is configured to perform dynamic device group membership calculations. For more information, see “Ensure Continuity of Dynamic Device Group Calculation” the *Horizontal Scalability guide*.

The following figure shows the results of switching back to the primary location.



## Switching Back to Different NA and Database Servers

This section describes how to switch back to different primary servers than those from which the switchover to the disaster recovery location occurred. This situation applies to the following cases:

- New hardware has been provisioned in the original primary location. Any or all of the NA cores could be running on newly-provisioned servers. Additionally, the NA database might be running on a newly-provisioned server.
- The original primary location is no longer available, so a different site is now being used as the primary location. All NA servers and the database server are newly-provisioned.
- In this procedure, the following terms apply:
- The primary location is the data center that will receive the NA deployment switched back from the disaster recovery location. This location could be the original primary location data center with newly-provisioned servers, or it could be a different data center with newly-provisioned servers.
  - The existing NA cores are NA cores in the original primary location data center that are still available for switchback.
  - The new NA cores are NA cores running on newly-provisioned servers in the primary location.
  - The disaster recovery location is data center currently hosting the existing NA deployment.

Switching back from the disaster recovery location to one or more new servers in the primary location involves the following general process:

- ["Task 1: Disable NA in the Disaster Recovery Location" on page 21](#)
- ["Task 2: Enable the Use of the Database in the Primary Location" on page 22](#)
- ["Task 3: Enable NA in the Primary Location" on page 23](#)
- ["Task 4: Finish Switchback" on the previous page](#)

### Task 1: Disable NA in the Disaster Recovery Location



This task assumes that all NA services are stopped on all NA cores in the primary location.

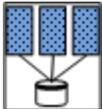
To disable NA in the disaster recovery location, follow these steps:

1. Determine the core IDs of the NA cores in the original primary location that are no longer available.
  - a. Connect as an NA administrator to the NA proxy on an NA core running in the disaster recovery location.
  - b. Run the following command:

	<p><b>list core</b></p> <ol style="list-style-type: none"><li>c. Note the core ID for each unavailable NA server.</li></ol> <ol style="list-style-type: none"><li>2. Prepare to stop NA in the disaster recovery location. Do the following:<ol style="list-style-type: none"><li>a. Notify users to log out.</li><li>b. Log on to the NA console for one of the NA cores in the disaster recovery location.</li><li>c. Pause tasks scheduled to start during the switchback process (<a href="#">Task 3, step 9</a>). Include time for the currently running tasks to complete. Also include time for database synchronization. For example:<ol style="list-style-type: none"><li>i. On the Search for Task page (<b>Reports &gt; Search For &gt; Tasks</b>), for the Schedule Date field, set since to <b>Now</b> and until to <b>4 hours later</b>.</li><li>ii. On the Task Search Results page, pause each listed task.</li></ol></li><li>d. On the Running Tasks page (<b>Tasks &gt; Running Tasks</b>), examine the list of running tasks. If any critical tasks are running, wait for them to complete before continuing with <a href="#">step 3</a>, next.</li></ol></li><li>3. Stop all NA services on all NA cores in the disaster recovery location. See <a href="#">Start, Stop, or Restart All Services</a>.</li></ol>
--	---

## Task 2: Enable the Use of the Database in the Primary Location

	<ol style="list-style-type: none"><li>1. If the database server is newly-provisioned, create the NA schema on that database server. Note the following:<ul style="list-style-type: none"><li>• The NA database user in the disaster recovery location must have the same name and permissions as the NA database user in the primary location.</li><li>• Copy the NA schema tables only.</li></ul>For example, you might follow this process:<ol style="list-style-type: none"><li>a. Install the database software.</li><li>b. Create a database user with the same name and permissions as the NA database user for the database in the disaster recovery location.</li></ol></li><li>2. Synchronize the database in the primary location with the database in the disaster recovery location. Possible approaches include the following:<ul style="list-style-type: none"><li>• Use database tools to copy the NA database from the disaster recovery location to the primary location. Copy the NA schema only. For example, you might follow this process:<ol style="list-style-type: none"><li>i. Export the NA database from the recovery database server in the disaster</li></ol></li></ul></li></ol>
---	--

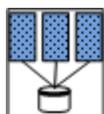
	<p>recovery location.</p> <ol style="list-style-type: none"><li>ii. Wipe the NA database from the database server in the primary location.</li><li>iii. Import the NA database to the database server in the primary location.</li></ol> <ul style="list-style-type: none"><li>• If reverse replication from the database in the disaster recovery location to the database in the new primary location is running, analyze the reverse replication transaction logs.</li></ul> <p>If the transaction logs have overflowed, reverse replication becomes ineffective. In this case, do the following:</p> <ol style="list-style-type: none"><li>i. Disable reverse replication to the database in the primary location.</li><li>ii. Use database tools to copy the NA database from the disaster recovery location to the primary location. Copy the NA schema only.</li></ol> <p>If the transaction logs are within bounds, do the following:</p> <ol style="list-style-type: none"><li>i. Wait for all database updates to replicate to the database in the primary location.</li><li>ii. After replication is complete, disable reverse replication to the database in the primary location.</li></ol>
	<p>3. Configure database replication as follows:</p> <ul style="list-style-type: none"><li>• If the NA database server is newly-provisioned, update database replication as follows:<ol style="list-style-type: none"><li>i. On the disaster recovery location database server, remove the configuration for database replication from the original primary location database server.</li><li>ii. Configure one-way database replication from the database in the primary location to the recovery database in the disaster recovery location.</li></ol></li><li>• If the NA database server is <i>not</i> newly-provisioned, re-enable database replication from the primary location to the disaster recovery location.</li></ul>
	<p>4. Connect the NA cores in the disaster recovery location to the database in the primary location.</p> <p>On each NA server in the disaster recovery location, in a text editor, edit following file:</p> <ul style="list-style-type: none"><li>• <i>Windows</i>: &lt;NA_HOME&gt;\server\ext\jboss\server\default\deploy\db_ds.xml</li><li>• <i>Linux</i>: &lt;NA_HOME&gt;/server/ext/jboss/server/default/deploy/db_ds.xml</li></ul> <p>This file contains <i>two</i> lines defining the JdbcUrl attribute. For example:</p> <ul style="list-style-type: none"><li>• PostgreSQL: &lt;attribute name="JdbcUrl"&gt;jdbc:postgres:jdbc:postgresql://hostname:5432/na:1521:nadb&lt;/attribute&gt;</li><li>• Oracle: &lt;attribute name="JdbcUrl"&gt;jdbc:oracle:thin:</li></ul>

	<pre>@db.example.com:1521:nadb&lt;/attribute&gt;</pre> <ul style="list-style-type: none"><li>• SQL Server: &lt;attribute name="JdbcUrl"&gt;jdbc:sqlserver://db.example.com:1433;DatabaseName=NA;SendStringParametersAsUnicode=false&lt;/attribute&gt;</li></ul> <p>For each JdbcUrl attribute, replace the database server name (db.example.com in the example) with the fully-qualified domain name or IP address of the database server in the primary location.</p> <p>For PostgreSQL Server, also replace the database name (na in the example) with the database name for the NA database in the primary location.</p> <p>For Oracle, also replace the database SID or service name (nadb in the example) with the database SID or service name for the NA database in the primary location.</p> <p>For SQL Server, also replace the database name (NA in the example) with the database name for the NA database in the primary location.</p> <p>This step replaces the connection between these NA cores and the database in the disaster recovery location with a connection to the database in the primary location.</p>
	<ol style="list-style-type: none"><li>5. Ensure that all existing NA servers in the primary location are powered on.</li><li>6. If necessary, stop all NA services on all NA cores in the primary location. See <a href="#">Start, Stop, or Restart All Services</a>.</li><li>7. Connect the existing NA cores in the primary location to the local database (the database in the primary location) as follows:<ul style="list-style-type: none"><li>• If the NA database server is newly-provisioned, update the db_ds.xml file as described in <a href="#">step 4</a>.</li><li>• If the NA database server is <i>not</i> newly-provisioned, verify the db_ds.xml file as described in <a href="#">step 4</a>.</li></ul></li><li>8. In the primary location database, remove the unavailable NA servers from the NA database.<ol style="list-style-type: none"><li>a. Locate the list of core IDs for the unavailable NA servers, as determined in <a href="#">Task 1, step 1</a>.</li><li>b. For each unavailable NA server in the original primary location, run the following script as appropriate for your database type.<p>You do not need to run any scripts if you use a PostgreSQL database.</p><p><b>Oracle Script</b></p><pre>UPDATE RN_SITE SET OwningCoreID = 1 WHERE OwningCoreID = &lt;coreID&gt;; UPDATE RN_SITE SET ManagingCoreID = 1 WHERE ManagingCoreID = &lt;coreID&gt;; UPDATE RN_SCHEDULE_TASK SET CoreID = 1 WHERE CoreID = &lt;coreID&gt;; DELETE FROM RN_CORE WHERE CoreID = &lt;coreID&gt;; COMMIT;</pre><p><b>SQL Server Script</b></p></li></ol></li></ol>

```
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.

## Task 3: Enable NA in the Primary Location



To enable NA in the primary location, follow these steps:

1. If the configuration of the NA cores in the disaster recovery location has changed since switchover, update the NA core files on each existing NA server in the primary location. See [Table 1](#).

**Tip:** You will restart the NA services in [step 8](#). You do not need to do so now.

2. In the primary location, install the new NA cores as additional NA cores connected (through Horizontal Scalability) to the database in the primary location. Install the same NA version and patch level as in on the existing NA cores.
3. Configure the new NA cores as follows:
  - a. Stop all NA services on the new NA cores. See [Start, Stop, or Restart All Services](#).
  - b. Put into place any modified files from those listed in [Table 1](#). Options include:
    - o Copy the files from one of the NA servers in the disaster recovery location.
    - o Retrieve the files from a backup.
  - c. Start all NA services on the new NA cores. See [Start, Stop, or Restart All Services](#).
  - d. *Optional.* Perform any NA core-specific tuning. Restart the NA services as needed.

4. If the NA environment includes NA Satellite functionality, do the following:
  - a. If necessary, in the primary location install one or more core gateways.  
Install at least one core gateway in the primary location to continue communication with the existing gateway mesh. Optionally install additional core gateways, up to one core gateway per NA core. During installation, configure each core gateway in the primary location as follows:
    - o Use the same Gateway Crypto Data file as for the core gateways in the original primary location.

- o Assign the same realm name, typically Default Realm, to each core gateway.
- o After installation, update the `opswgw.EgressFilter` entry to match the following:  

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

For information, see the *Satellite guide*.

- b. If necessary, in the primary location reconnect each core gateway installed on a system other than an NA server to an NA core.

For each core gateway that remains from the original primary location configuration and is not installed on an NA server, connect that core gateway with an NA server.

- c. For each satellite in the gateway mesh, update the satellite configuration to enable communication with a core gateway in the primary location.
  - i. Connect to the satellite gateway server as the root user.
  - ii. Change to the following directory:  

```
/etc/opt/opsware/opswgw-<gateway_name>
```
  - iii. Back up the `opswgw.properties` file to a location outside the gateway installation directory.
  - iv. In a text editor, such as `vi`, open the `opswgw.properties` file.
  - v. Modify the `opswgw.TunnelSrc` entry that points to a core gateway in the original primary location to now point to a core gateway in the new primary location.

For example, change:

```
opswgw.TunnelSrc=<core_gateway1_IP>:2001:100:0:  
/var/opt/opsware/crypto/opswgw-RemoteGw/opswgw.pem  
opswgw.TunnelSrc=<core_gateway2_IP>:2001:200:0:  
/var/opt/opsware/crypto/opswgw-RemoteGw/opswgw.pem
```

To:

```
opswgw.TunnelSrc=<core_gateway11_IP>:2001:100:0:  
/var/opt/opsware/crypto/opswgw-RemoteGw/opswgw.pem  
opswgw.TunnelSrc=<core_gateway2_IP>:2001:200:0:  
/var/opt/opsware/crypto/opswgw-RemoteGw/opswgw.pem
```

- vi. Restart the satellite gateway.
- 5. Activate the existing NA cores in the primary location.
    - a. Connect as an NA administrator to the NA proxy on an NA core running on a newly-provisioned NA server.
    - b. Run the following command:  

```
list core
```
    - c. From the `list core` command output, determine the core IDs of the original NA cores in the primary location, and then run the following command for each of these core IDs:

	<p><b>set core status -status active -coreid &lt;coreid&gt;</b></p> <p>This step changes the NA core status in the NA database in the primary location. Because replication is running, the NA database in the disaster recovery location also shows these cores as active.</p> <ol style="list-style-type: none"><li>6. Update the site assignments.<ol style="list-style-type: none"><li>a. Log on to the NA console for the NA core that was just restarted in the primary location.</li><li>b. Update the site assignments. In the NA console, open the <b>Site Reassignment</b> page (<b>Admin &gt; Distributed &gt; Site Reassignment</b>), and then assign all partitions to NA cores in the primary location.</li><li>c. <i>Optional.</i> On the <b>Partitions</b> page (<b>Admin &gt; Security Partitions</b>), verify the mapping of partitions to NA cores.</li></ol></li><li>7. Deactivate the NA cores in the disaster recovery location.<ol style="list-style-type: none"><li>a. Connect as an NA administrator to the NA proxy on the NA core that was recently restarted in the primary location.</li><li>b. Run the following command: <b>list core</b></li><li>c. From the <code>list core</code> command output, determine the core IDs of the NA cores in the disaster recovery location, and then run the following command for each of these core IDs: <b>set core status -status inactive -coreid &lt;coreid&gt;</b></li></ol><p>This step changes the NA core status in the NA database in the primary location. Because replication is running, the NA database in the disaster recovery location also shows these cores as inactive.</p></li></ol>
	<ol style="list-style-type: none"><li>8. If necessary, resume the tasks that were paused in <a href="#">Task 1, step 2</a>.<ol style="list-style-type: none"><li>a. Log on to the NA console for one of the NA cores in the primary location.</li><li>b. On the Search for Task page (<b>Reports &gt; Search For &gt; Tasks</b>), for the Task Status field, select Paused.</li><li>c. On the Task Search Results page, resume each listed task.</li></ol><p>Later instances of periodic tasks remove the paused instances of those tasks.</p></li></ol>

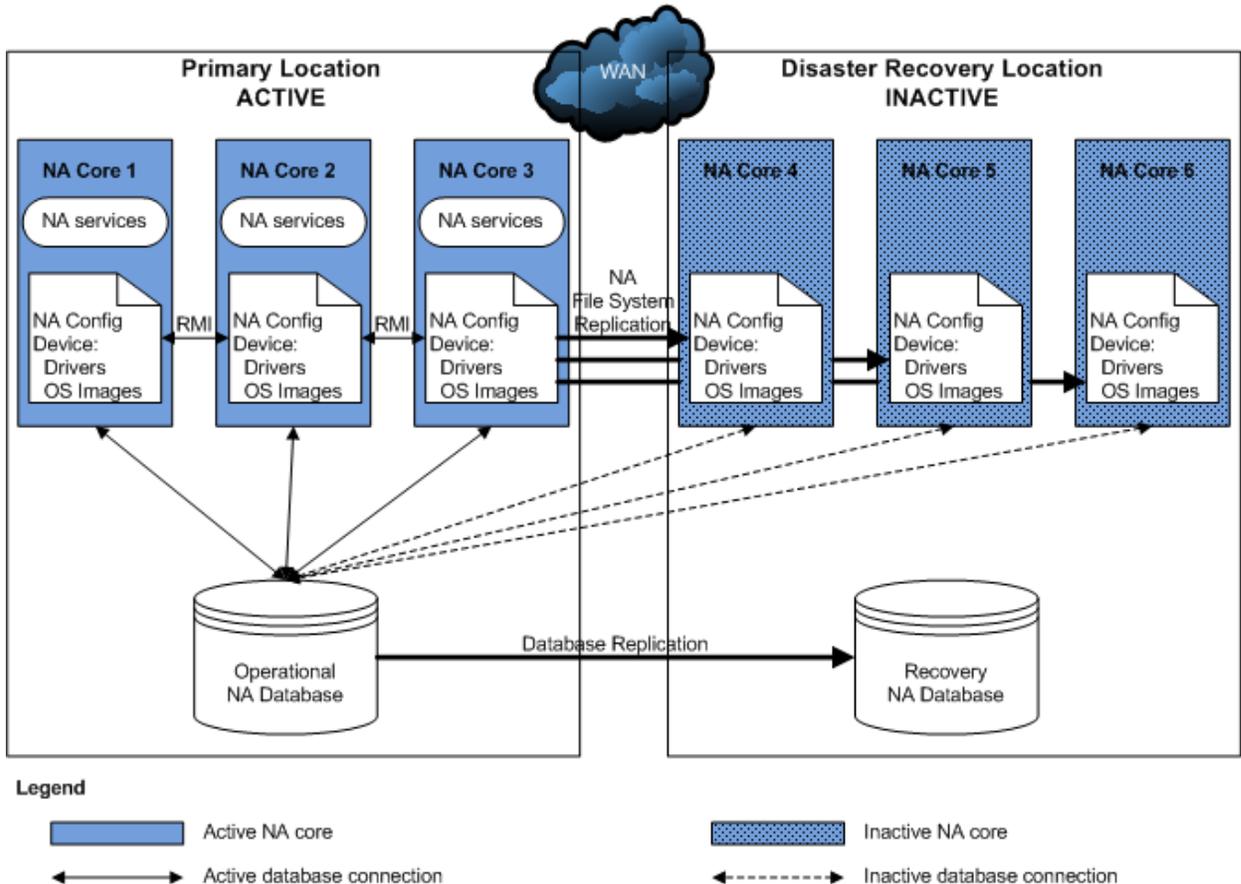
## Task 4: Finish Switchback

	<ol style="list-style-type: none"><li>1. Configure any applications that integrate with NA to connect to the working NA cores in the primary location.</li><li>2. Notify users to connect to the NA console on the NA cores in the primary location.</li><li>3. Configure replication of the NA file system from one of the active NA cores to the inactive NA cores (or an intermediate server) as described in "<a href="#">Files to Synchronize Across Cores</a>" on page 1.</li><li>4. <i>Optional.</i> If you chose to synchronize the NA files to an intermediate server, to</li></ol>
---	--

conserve resources in the disaster recovery location, power down the NA servers in the disaster recovery location.

5. *Optional.* Configure the managed devices (or the syslog relay) to send syslog messages to one NA core in the primary location and to no NA cores in the disaster recovery location.
6. Verify that one NA core in the disaster recovery location is configured to perform dynamic device group membership calculations. .
7. *Optional.* For the NA cores on the newly-provisioned NA servers, consider the risk that if the primary location is not accessible during switchover to the disaster recovery location, the NA cores in the primary location cannot be deactivated. In this case, two NA cores (one each in the primary and disaster recovery locations) might run the same task. To mitigate this risk, on each NA server in the primary location, disable automatic starting of the NA services. See [Disable All Services](#).

The following figure shows the results of switching back to the new primary location.



## Creating a New Disaster Recovery Location

This section describes how to set the current (original) disaster recovery location to be the new primary location with a new disaster recovery location. In this procedure, the following terms apply:

- The original primary location is the data center that is no longer available.
- The new primary location is the data center hosting the existing NA deployment. This location is the original disaster recovery location.
- The new disaster recovery location is the data center that will receive the new NA deployment.

This configuration involves the following process:



1. *Optional.* Consider the risk that if the new primary location is not accessible during switchover to the disaster recovery location, the NA cores in the primary location cannot be deactivated. In this case, two NA cores (one each in the primary and disaster recovery locations) might run the same task. To mitigate this risk, on each NA server in the primary location, disable automatic starting of the NA services. See [Disable All Services](#).
2. Determine the core IDs of the NA cores in the original primary location that are no longer available.
  - a. Connect as an NA administrator to the NA proxy on an NA core running in the disaster recovery location.
  - b. Run the following command:  

```
list core
```
  - c. Note the core ID for each unavailable NA server.
3. *Optional.* Perform any tuning needed to prepare the NA cores in the new primary location for daily use. Restart the NA services as needed.
4. *Optional.* Configure the managed devices (or the syslog relay) to send syslog messages to one NA core in the new primary location.
5. Prepare to stop NA in the new primary location.
  - a. Notify users to log out.
  - b. Log on to the NA console for one of the NA cores in the new primary location.
  - c. Pause tasks scheduled to start during the disaster recovery configuration process (until [Task 4, step 1](#)). Include time for the currently running tasks to complete. Also include time for database synchronization. For example:
    - i. 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**.
    - ii. On the Task Search Results page, pause each listed task.
  - d. On the Running Tasks page (**Tasks > Running Tasks**), examine the list of running tasks.

If any critical tasks are running, let them run to completion before continuing with [step 6](#).

6. In the new primary location, stop all NA services on all NA cores. See [Start, Stop, or Restart All Services](#).
7. Clean up database replication as follows:
  - a. On the new primary location database server, remove the configuration for database replication from the original primary location database server.
  - b. If reverse replication from the database in the original disaster recovery location to the database in the original primary location was configured, remove that configuration.
8. In the new primary location database, remove the unavailable NA servers from the NA database.
  - a. Locate the list of core IDs for the unavailable NA servers, as determined in [step 2](#).
  - b. For each unavailable NA server in the original primary location, run the following script as appropriate for your database type.

You do not need to run any scripts if you use a PostgreSQL database.

**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.

9. Remove any remaining configuration for replication of the NA file system from one of the NA cores in the original primary location to the original disaster recovery NA cores (or an intermediate server).
10. Beginning with "[Task 2: Configure Database Replication](#)" on [page 1](#) and through the end of "[NA Disaster recovery initial setup](#)" on [page 8](#), complete the process for setting up NA for disaster recovery.

# Disaster Recovery Configuration with Horizontal Scalability over a WAN

A disaster recovery configuration includes one or more Network Automation (NA) cores in a disaster recovery location that is geographically remote from the primary network device management location. Some or all of the NA cores in the disaster recovery location can be actively managing the network before a disaster occurs. This approach might save equipment costs. Because the active NA cores in the disaster recovery must be connected to the NA database in the disaster recovery location, this approach does not provide continuous availability of NA.

The described approach is an implementation of Horizontal Scalability over a WAN. This implementation is not recommended because it comes with costs and limitations.

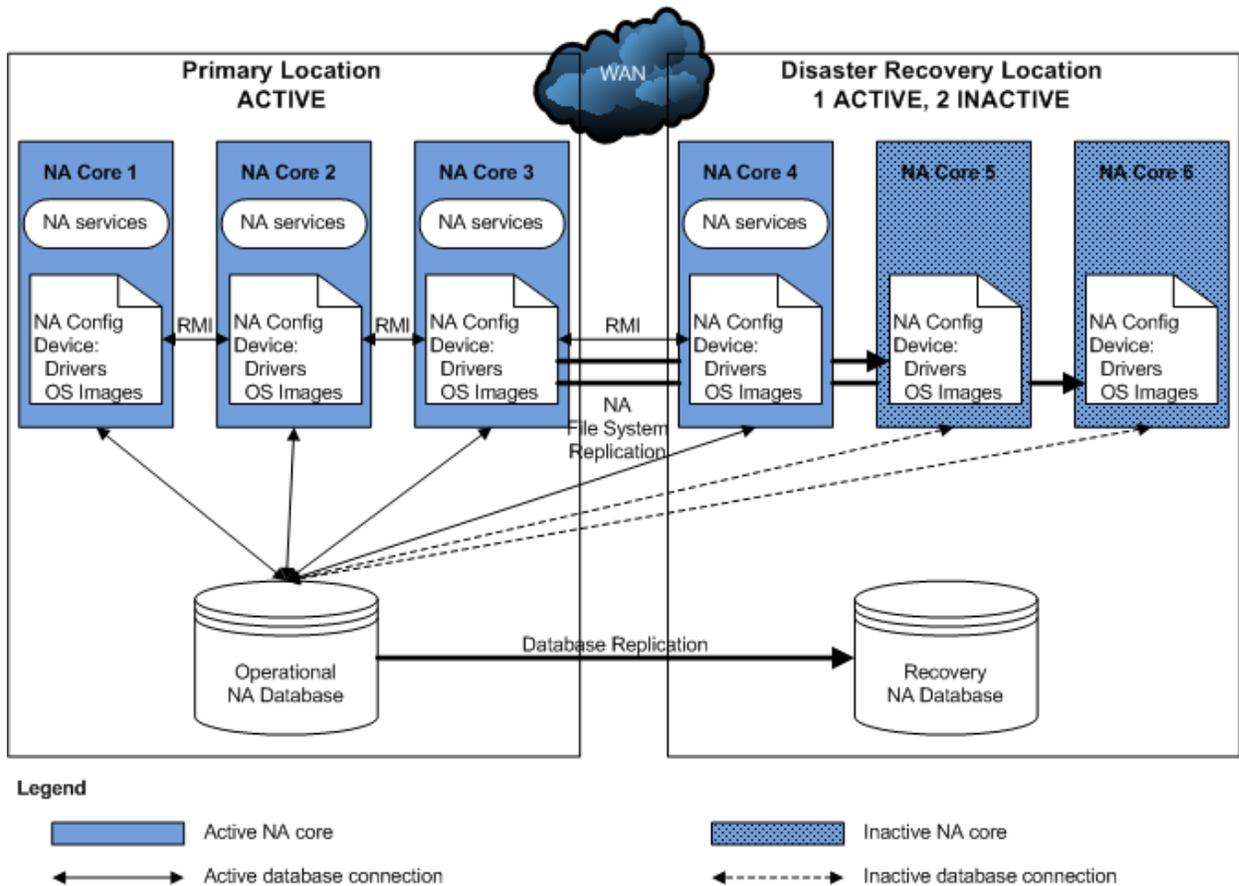
If your organization determines that the benefits of implementing Horizontal Scalability over a WAN outweigh the costs and limitations, read this appendix for supplemental information about disaster recovery configuration.

This topic includes:

- ["Example Architecture" on the next page](#)
- ["Switchover Considerations" on the next page](#)
- ["Switchback Considerations" on page 38](#)

## Example Architecture

The following figure shows an example architecture for disaster recovery with one active NA core (NA core 4) in the disaster recovery location. During general operation, NA core 4 functions as if it were in the primary location. If a disaster occurs, NA core 4 loses connection with the NA database in the primary location. Therefore, it is unable to continue network device management. The switchover step of connecting NA core 4 to the NA database in the disaster recovery location requires the temporarily shut down of NA core 4. This architecture does not provide continuous operation during the transition from the primary location to the disaster recovery location.



### Initial Setup Considerations

During initial setup of the disaster recovery configuration, treat the NA cores that will remain active in the disaster recovery location as if they were in the primary location.

### Switchover Considerations

During switchover, connect all NA cores in the disaster recovery location to the NA database in that location. Be sure that all NA cores are stopped before initiating this change. The NA cores in the

disaster recovery location that were active before switchover do not need to receive updated configuration files or be changed to active status.

Follow the steps in "[Switchover](#)" on page 14 with the following notes:

- [Task 2: Plan to Disable in the Primary Location, step 1.](#)  
Because you will be connecting the active NA cores in the disaster recovery location to the NA database in the disaster recovery location, you do not need to disable automatic starting of the NA services on these NA cores.
- [Task 2: Disable in the Primary Location, step 3.](#)  
Stop all NA services on all active NA cores. NA must be shut down during the database connection change.
- [Task 3: Enable Use of the Database in the Disaster Recovery Location, step 5.](#)  
Connect all NA cores in the disaster recovery location to the NA database in the disaster recovery location.
- [Task 4: Enable in the Disaster Recovery Location, step 1.](#)  
You do not need to copy configuration files to the NA cores in the disaster recovery location that were being used to manage the network before switchover.
- [Task 4: Enable in the Disaster Recovery Location, step 2.](#)  
You do not need to activate the NA cores in the disaster recovery location that were being used to manage the network before switchover.

## Switchback Considerations

For switchback, the supplemental information depends on the switchback scenario.

### Scenario One

For "[Switching Back to the Original Servers in the Primary Location](#)" on page 1, connect all NA cores in the disaster recovery location to the NA database in the primary location. Be sure that all NA cores are stopped before initiating this change. The NA cores in the disaster recovery location that will remain active after switchback do not need to receive updated configuration files or be changed to active status.

Follow the steps in the procedure for "[Switching Back to the Original Servers in the Primary Location](#)" on page 1 with the following notes:

- [Task 2: Enable the Use of the Database in the Primary Location, step 3.](#)  
Connect all NA cores in the disaster recovery location to the NA database in the primary location.
- [Task 3: Enable in the Primary Location](#)  
Treat the NA cores that will remain active in the disaster recovery location as if they were in the primary location. You do not need to copy configuration files to these NA cores. You do not need to activate these NA cores.

### Scenario Two

For "[Switching Back to Different and Database Servers](#)" on page 1, connect all NA cores in the disaster recovery location to the NA database in the primary location. Be sure that all NA cores are stopped before initiating this change. The NA cores in the disaster recovery location that will remain

active after switchback do not need to receive updated configuration files or be changed to active status.

Follow the steps in the procedure for "[Switching Back to Different and Database Servers](#)" on page 1 with the following notes:

- [Task 2: Enable the Use of the Database in the Primary Location, step 3.](#)  
Connect all NA cores in the disaster recovery location to the NA database in the primary location.
- [Task 3: Enable in the Primary Location](#)  
Treat the NA cores that will remain active in the disaster recovery location as if they were existing NA cores in the primary location. You do not need to copy configuration files to these NA cores. You do not need to activate these NA cores.

### **Scenario Three**

For "[Creating a New Disaster Recovery Location](#)" on page 1, treat all NA cores in the disaster recovery location as being in the new primary location. If you choose to configure NA cores that will remain active in the new disaster recovery location, do the following:

- Add all NA cores in the new disaster recovery location to the Horizontal Scalability environment as directed.
- Do not deactivate the NA cores that will remain active in the new disaster recovery location. (Treat these NA cores as if they were in the primary location.)

## Send Documentation 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 Disaster Recovery configuration guide (Network Automation 10.30)**

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@hpe.com](mailto:network-management-doc-feedback@hpe.com).

We appreciate your feedback!