Disaster Recovery Configuration Guide for CiscoWorks Network Compliance Manager 1.8

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-27505-01

<THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Book Title

Contents

1 NCM Disaster Recovery Concepts

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

What is Disaster Recovery?

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 NCM database in the disaster recovery location is completely synchronized with the NCM 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 Architecture

For NCM, disaster recovery configuration involves duplicating the NCM environment running in a data center in the primary location to a remote data center in the disaster recovery location. Figure 1 shows this duplication for a three NCM core Horizontal Scalability environment. (An NCMcore is a physical or virtual server on which the NCM 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.

Figure 1 Example NCM Disaster Recovery Architecture

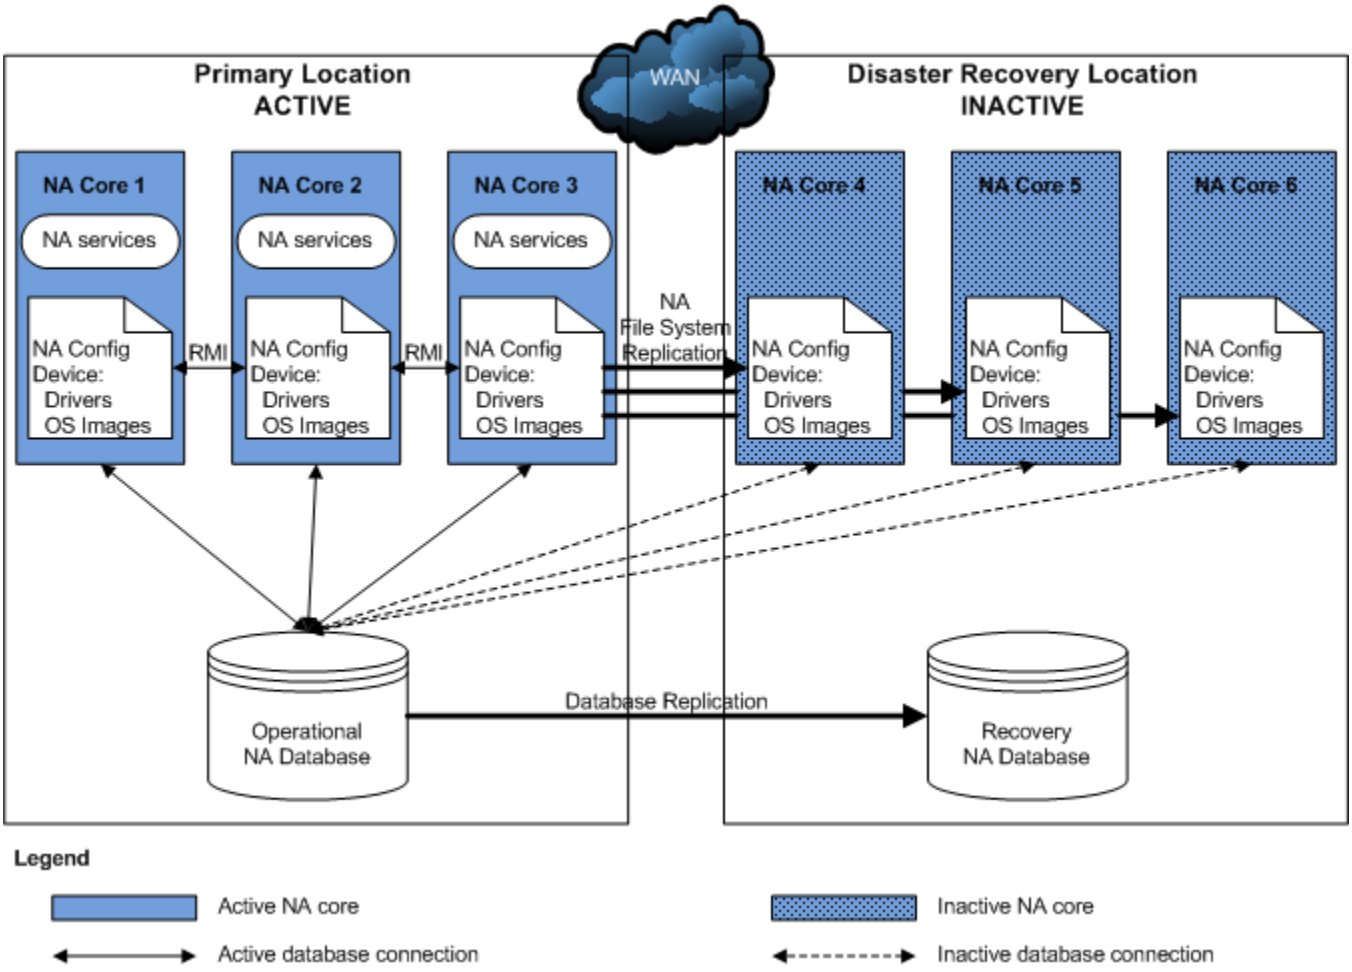

Note the following:

- Each location includes one NCM 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 NCM database is actively used by the NCM cores. The second database must be running to receive database updates; however, no NCM cores connect to the second database.

  The recommended NCM disaster recovery scenario includes one-way database replication.
- One to five active NCM cores connect to the operational NCM database using Horizontal Scalability. The server for each NCM core has a unique IP address and hostname, so the switchover and switchback procedures include updating any configuration that connects to the NCM servers. For best performance of the overall solution, it is recommended that each NCM server be located in the same data center as the database server to which it is most likely to connect.

  When multiple NCM cores are active simultaneously, Java remote method invocation (RMI) calls synchronize the NCM-specific file systems across the active NCM cores. RMI calls also manage schedules for running tasks across the active NCM cores.
- During disaster recovery configuration, all NCM cores connect to the operational NCM database. The NCM cores in the disaster recovery location are then set to the inactive state. While a NCM core is inactive, the following conditions apply:

  - That NCM core does not run tasks.
  - Users should *not* log on to the NCM Console.
  - Users might connect to the NCM command-line interface through telnet or SSH *only* for the purpose of setting the NCM core state during switchover or switchback.

- The maximum number of active NCM cores in a Horizontal Scalability environment is five. The maximum number of active and inactive NCM cores in a Horizontal Scalability environment is nine, which means that the maximum number of NCM cores in a disaster recovery scenario is five active NCM cores at the primary location and four inactive NCM cores at the disaster recovery location.

  During disaster recovery configuration, all nine NCM cores might be active at one time; however, no device management occurs on the NCM 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 NCM 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.
- As of NCM 1.8, NCM task management in the Horizontal Scalability environment includes the following behavior:
  - If NCM tasks for a given device are bound to only one NCM core (the default behavior; Horizontal Scalability topologies 1, 3, or 4), when a NCM administrator reassigns the sites from one NCM core to another NCM core, NCM moves the tasks associated with that site to the receiving NCM core.
  - If all NCM tasks for all devices are distributed in round-robin fashion across all NCM cores (Horizontal Scalability topology 2), when a NCM administrator sets a NCM core to the inactive state, all tasks scheduled to run on that NCM core are distributed to the remaining active NCM cores.
- This disaster recovery configuration is licensed as follows:
  - One production license for the total number of managed devices for one NCM core in the primary location.
  - One non-production license for the total number of managed devices for *each* additional NCM 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.

2 NCM Disaster Recovery Initial Setup

You can configure CiscoWorks Network Compliance Manager (NCM) in a disaster recovery any time after NCM is configured and running satisfactorily in the primary location. The approach described in this document requires NCM Horizontal Scalability functionality. For information about the supported database versions for NCM Horizontal Scalability, see “Databases for Horizontal Scalability” in the *NCM Support Matrix*.

This guide assumes NCM database replication in an operation–recovery configuration. The recovery database server contains a copy of the NCM 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. For information about the tested database replication technologies, see “Disaster Recovery” in the *NCM Support Matrix*.

Setting up NCM for Disaster Recovery

To perform initial setup of a NCM disaster recovery configuration, follow this general outline:

Task 1: Prepare the Primary Location for Disaster Recovery Configuration

1 In the primary location, start with a running NCM deployment. This deployment can be a single NCM core or a Horizontal Scalability environment containing up to five NCM cores. This deployment can also include NCM satellite functionality.

   NCM must be version 1.8 or later.

   For information about installing a single NCM core, see the *NCM Installation and Upgrade Guide*.

   For information about adding NCM cores to create a Horizontal Scalability environment, see the *NCM Horizontal Scalability Guide*.

2 *Optional*. Consider the risk that if the primary location is not accessible during switchover to the disaster recovery location, the NCM cores in the primary location cannot be deactivated. In this case, two NCM cores (one each in the primary and disaster recovery locations) might run the same task. To mitigate this risk, on each NCM server in the primary location, disable automatic starting of the NCM services.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, click **Properties**, and then set Startup Type to **Disabled**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *Linux*:

  ```
  mv /etc/rc.d/rc5.d/S99truecontrol /etc/S99truecontrol
  ```
- *Solaris*:

  ```
  mv /etc/rc2.d/S99truecontrol /etc/S99truecontrol
  ```

3 Prepare to stop NCM in the primary location.

a Notify users to log out.

b Log on to the NCM Console for one of the NCM cores in the primary location.

c Pause tasks scheduled to start during the disaster recovery configuration process (until Task 4, step 1 on page 15). Include time for the currently running tasks to complete. Also include time for database synchronization. For example:

- On the Search for Task page (**Reports > Search For > Tasks**), for the Schedule Date field, set since to `Until and anytime to 4 hours later`.
- 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 NCM services on all NCM cores in the primary location.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Stop**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *UNIX*: Run the following command:

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

Task 2: Configure Database Replication

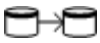

1. If the NCM database in the primary location, was created for the SYSTEM user (Oracle) or the SA user (SQL Server), move the NCM database to a custom user with the privileges described in the *NCM Installation and Upgrade Guide*. Do the following:
   a. Create a new tablespace or database instance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

   Note the following:

   - The NCM database user in the disaster recovery location must have the same name and permissions as the NCM database user in the primary location.
   - Copy the NCM schema tables only.
   - For Oracle or Oracle RAC, the SID or service name should be different between the two database servers.

     One SID *cannot* be a subset of the other SID, for example NCMRp and NCMRpBU. Instead, use SIDs that stand alone, for example NCMRp1 and NCMRp2.
   - For Microsoft SQL Server, the database name should be different between the two database servers.

   For example, you might follow this process:

   a. Install the database software.
   b. Create a database user with the same name and permissions as the NCM database user for the operational database in the primary location.
   c. Export the NCM database from the operational database in the primary location.
   d. Import the NCM database to the recovery database in the disaster recovery location.
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.

Task 3: Install and Configure NCM in the Disaster Recovery Location

1. In the disaster recovery location, install the NCM cores as additional NCM cores connected (through Horizontal Scalability) to the operational database in the primary location.

   For information, see “Adding Additional NCM Cores” in the *NCM Horizontal Scalability Guide* for NCM 1.8 or later.

If you reached this step from a switchback procedure, complete the script edits to account for having removed one or more NCM cores from the Horizontal Scalability environment.

Connecting the NCM cores in the disaster recovery location to the NCM database in the primary location ensures that the databases remain synchronized. These NCM cores will be stopped later in this procedure.

This use of Horizontal Scalability over the WAN is for disaster recovery configuration only. Daily use of Horizontal Scalability over the WAN is not supported.

2. *Optional*. Perform any NCM core-specific tuning. Restart the NCM services as needed.

3 *Optional*. Configure the managed devices to send syslog messages to one NCM core in the disaster recovery location.

4 If the primary location includes NCM Satellite functionality, do the following:

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 NCM core. During installation, configure each core gateway in the disaster recovery location as follows:

— Use the same Gateway Crypto Data file as for the core gateways in the primary location.

— Assign the same realm name, typically Default Realm, to each core gateway.

For information, see the *NCM Satellite Guide*.

b For each satellite in the gateway mesh, update the satellite configuration to enable communication with a core gateway in the disaster recovery location.

Edit the remote gateway configuration file:

`<gateway_install_dir>/opswgw-<gateway_name>/opswgw.properties`
(The default value of `<gateway_install_dir>` is `/etc/opt/opsware`.)

In the remote gateway configuration file, do the following:

— Add an `opswgw.TunnelSrc` 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:

```
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
```

— 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:
```

c Restart each remote core 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.

5 Deactivate the NCM cores in the disaster recovery location.

a Connect as a NCM administrator to the NCM proxy on one of the NCM cores in the disaster recovery location.

b Run the following command:

```
list core
```

c From the `list core` command output, determine the core IDs of the new NCM cores, and then run the following command for each core ID:

```
core status -status standby -coreid <coreid>
```

6 Stop all NCM services on all NCM cores in the disaster recovery location.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Stop**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *UNIX*: Run the following command:

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

Task 4: Finish the Disaster Recovery Configuration

1 In the primary location, start all NCM services on all active NCM cores.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Start**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *UNIX*: Run the following command:

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

2 Resume the tasks that were paused in Task 1, step 3 on page 12.

a Log on to the NCM Console for one of the NCM 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.

3 Notify users to resume use of the NCM Console on the NCM cores in the primary location.

4 Configure replication of the NCM file system from one of the active NCM cores to the inactive NCM cores (or an intermediate server) as described in Files to Synchronize Across NCM Cores on page 16.

5 *Optional*. If you chose to synchronize the NCM files to an intermediate server, to conserve resources, power down the NCM 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 NCM Cores

While most NCM data is stored in the NCM database, some files on the NCM core support the NCM Console and NCM functions. A complete disaster recovery scenario must include replication of these files. Table 1 lists the files to consider for replication among the NCM 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 NCM cores (in the primary location) to all of the inactive NCM cores in the disaster recovery location. This copy could be initiated by an active NCM core (push) or by each inactive NCM core (pull).

Alternatively, the script might copy files from an active NCM core to an intermediate server. This approach is useful if you want to leave the inactive NCM 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 NCM core.

Example file replication script

For example, the following script pulls files from an active NCM 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 NCM core. It copies files from the active NCM core whose core ID is `core1`.

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

Table 1 Files to Synchronize Across NCM Cores

Category	Files
RCX files, which are located in the following directory: • *Windows*: `%NCM_HOME%/jre` • *UNIX*: `$NCM_HOME/jre`	Specifically, at least the following files: • `site_options.rcx` (NCM server behavior) • `logging.rcx` (NCM logging levels) • `adjustable_options.rcx` (Customer-specified configuration options) • `distribution.rcx` (Distribution settings) • `securityfilter_additional_init.rcx` (Customer-specified filters for URL strings) Also include other RCX files that have been customized. **NOTE**: Do *not* include the `appserver.rcx` file, which contains paths to the local system. If this file has been modified, copy the changed blocks to the `adjustable_options.rcx` file for synchronization across all NCM cores.
Cisco-developed device drivers (`*.rdp`), which are located in a directory as specified by the `driver/dir` option in the `site_options.rcx` file, typically: • *Windows*: `%NCM_HOME%\server\lib\drivers` • *UNIX*: `$NCM_HOME/server/lib/drivers`	Synchronize all files in the `drivers` directory.
Device drivers developed outside of Cisco, which are located in a directory as specified by the `driver/extension/dir` option in the `site_options.rcx` file.	Synchronize all files in the identified directory.
Device operating system images, which are located in a directory as specified by the `deploy/repository/root` option in the `site_options.rcx` file, typically: • *Windows*: `%NCM_HOME%\server\images` • *UNIX*: `$NCM_HOME/server/images`	Synchronize all files in the `images` directory.

Verifying the Disaster Recovery Configuration

To verify the initial setup of a NCM 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 NCM databases are the same.

2. Check the database replication logs.

   - Are there any replication errors?
   - Is there a problem with any of the NCM 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. Set up a schedule for performing regular trimming of the replication log files.
5. Verify NCM file system replication by comparing the file sizes and timestamps in the primary and disaster recovery locations.

3 Switchover

When the primary location becomes unavailable, an administrator can follow the procedure described in this chapter to switch use of CiscoWorks Network Compliance Manager (NCM) over to the disaster recovery location. In the case of an unplanned disaster event, NCM 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 NCM downtime can be very short with no data loss.

Figure 2 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.

Figure 2 After a Disaster Event, Before Switchover

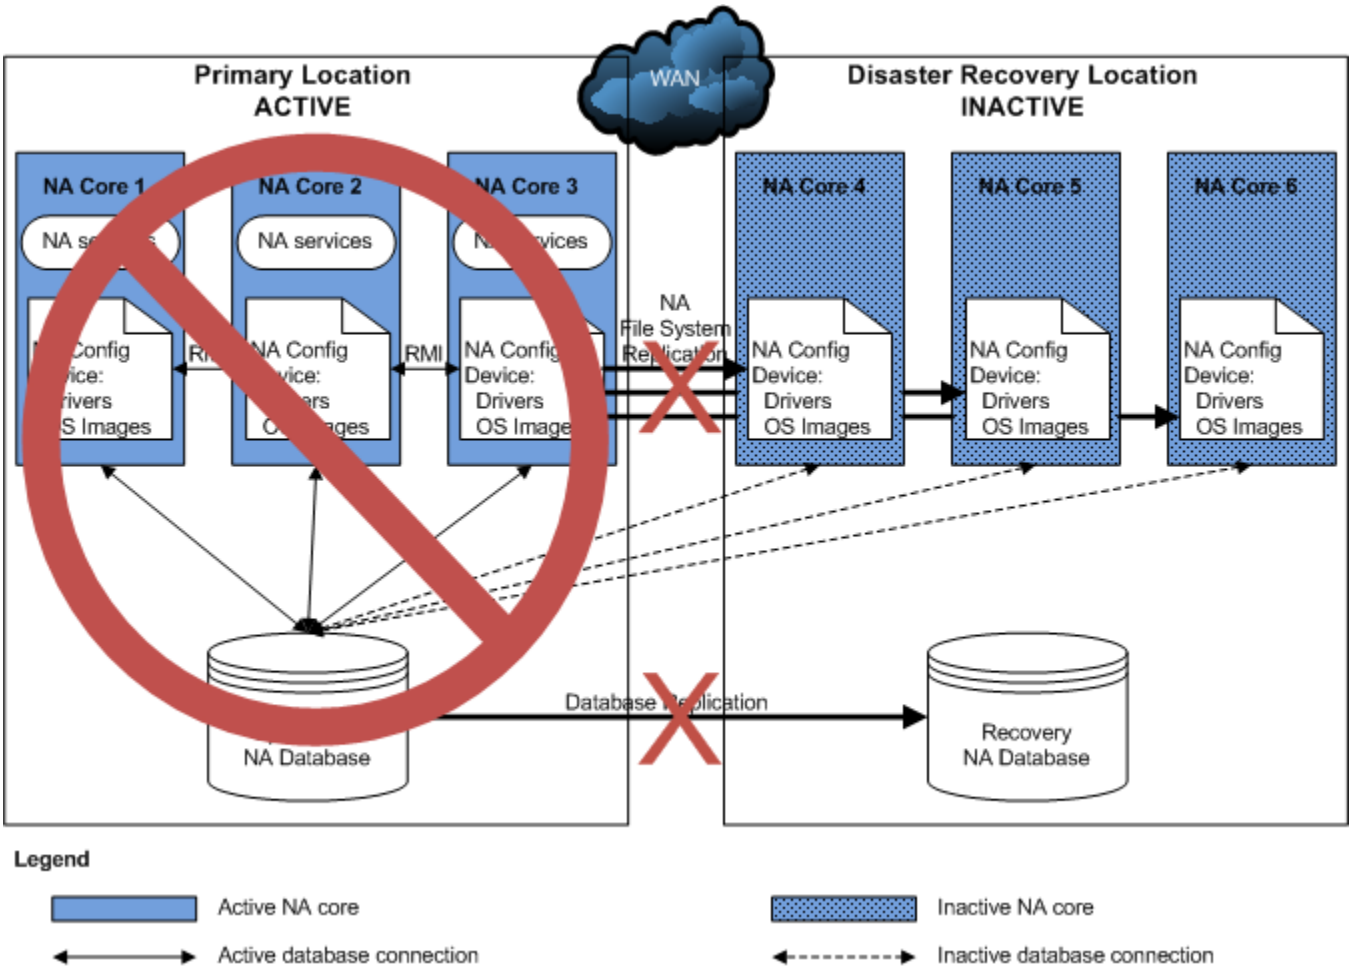

If the primary location is *not* accessible, to switch over from the primary location to the disaster recovery location, complete the following tasks in order:

- Task 1: Plan to Disable NCM in the Primary Location on page 20
- Task 3: Enable Use of the Database in the Disaster Recovery Location on page 22
- Task 4: Enable NCM in the Disaster Recovery Location on page 23
- Task 5: Finish Switchover on page 25

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 NCM in the Primary Location on page 20
- Task 3: Enable Use of the Database in the Disaster Recovery Location on page 22
- Task 4: Enable NCM in the Disaster Recovery Location on page 23
- Task 5: Finish Switchover on page 25

Task 1: Plan to Disable NCM in the Primary Location

If the primary location is not currently accessible, make plans to disable NCM 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 NCM services (as described in step 3 on page 12)
- Physical changes to the NCM server (for example, disconnecting the power source or the network cable)

Continue with Task 3 on page 22.

Task 2: Disable NCM in the Primary Location

If the primary location is accessible, disable NCM 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 NCM in the Primary Location.

1 On each NCM server in the primary location, disable automatic starting of the NCM services.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, click **Properties**, and then set Startup Type to **Disabled**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *Linux*:

  ```
  mv /etc/rc.d/rc5.d/S99truecontrol /etc/S99truecontrol
  ```
- *Solaris*:

  ```
  mv /etc/rc2.d/S99truecontrol /etc/S99truecontrol
  ```

2 Prepare to stop NCM in the primary location.

a Notify users to log out.

b Log on to the NCM Console for one of the NCM cores in the primary location.

c Pause tasks scheduled to start during the switchover process (until Task 4, step 5 on page 24). Include time for the currently running tasks to complete. For example:

— On the Search for Task page (**Reports > Search For > Tasks**), for the Schedule Date field, set since to **`Until and anytime to 2 hours later`**.

— 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 NCM services on all NCM cores in the primary location.

- *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
```

Ensure that the NCM services on the NCM cores in the primary location remain stopped until directed otherwise in the switchback procedure.

4 Wait for NCM file system replication from the primary location to complete.

Verify completeness by comparing the file sizes and timestamps in the primary and disaster recovery locations.

5 Wait for all database updates to replicate from the database in the primary location to the database in the disaster recovery location.

6 On the primary location database server, disable database replication to the database in the disaster recovery location.

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

1. On the disaster recovery location database server, disable database replication from the database in the primary location.

2. *Optional*. 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.

If the down time of the primary location is expected to be less that 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 switchback. 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.

3. If necessary, power on the NCM servers in the disaster recovery location.
4. If automatic starting of the NCM services is enabled, stop all NCM services on all NCM cores in the disaster recovery location.
   - *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
     ```

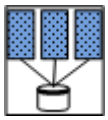

5. Connect the NCM cores in the disaster recovery location to the local database (the database in the disaster recovery location).

   On each NCM server in the disaster recovery location, in a text editor such as WordPad or vi, edit following file:

   - *Windows*:

     ```
     <NCM_HOME>\server\ext\jboss\server\default\deploy\db_ds.xml
     ```

   - *UNIX*:

     ```
     <NCM_HOME>/server/ext/jboss/server/default/deploy/db_ds.xml
     ```

   This file contains two lines defining the `JdbcUrl` attribute. For example:

   - Oracle:

     ```
     <attribute name="JdbcUrl">jdbc:oracle:thin:@db.example.com:
     1521:nadb</attribute>
     ```

   - SQL Server:

     ```
     <attribute name="JdbcUrl">jdbc:sqlserver://db.example.com:
     1433;DatabaseName=NCM;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 Oracle, also replace the database SID or service name (**nadb** in the example) with the database SID or service name for the NCM database in the disaster recovery location.

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

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

Task 4: Enable NCM in the Disaster Recovery Location

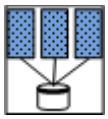

To enable NCM in the disaster recovery location, follow these steps:

1. If necessary, copy the NCM server files from the intermediate location to the correct locations on the NCM servers in the disaster recovery location. Maintain file ownership and permissions.
2. Activate the NCM cores in the disaster recovery location.
   a. On *one* NCM core in the disaster recovery location, start all NCM services.
      — *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Start**:

      TrueControl ManagementEngine

      TrueControl FTP Server

      TrueControl SWIM Server

      TrueControl Syslog Server

      TrueControl TFTP Server

      — *UNIX*: Run the following command:

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

   b. Using telnet or SSH, connect as a NCM administrator to the NCM proxy on that NCM core.
   c. Run the following command:

      ```
      list core
      ```

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

      ```
      core status -status normal -coreid <coreid>
      ```

This step changes the NCM core status in the NCM database in the disaster recovery location only. Unless reverse replication is running, the NCM database in the primary location still shows these cores as inactive.

3 If NCM is configured with tasks for a given device bound to only one core (the default Horizontal Scalability configuration), do the following:

   a On the *one* running NCM core in the disaster recovery location, restart all NCM 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
         ```

   b Log on to the NCM Console for the running NCM core in the disaster recovery location.

   c Update the site assignments. In the NCM Console, open the **Site Reassignment** page (**Admin > Distributed > Site Reassignment**), and then assign all partitions to NCM cores in the disaster recovery location.

4 Deactivate the NCM cores in the primary location.

   a Connect as a NCM administrator to the NCM proxy on the *one* running NCM core in the disaster recovery location.

   b Run the following command:

      ```
      list core
      ```

   c From the `list core` command output, determine the core IDs of the NCM cores in the primary location, and then run the following command for each core ID:

      ```
      core status -status standby -coreid <coreid>
      ```

This step changes the NCM core status in the NCM database in the disaster recovery location only. Unless reverse replication is running, the NCM database in the primary location still shows these cores as active.

5 Start (or restart) all NCM services on all NCM cores in the disaster recovery location.

- *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
  ```

The NCM services on the primary NCM cores must remain stopped.

6 Resume the tasks that were paused in Task 2, step 2 on page 21.

   a Log on to the NCM Console for one of the NCM 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.

Task 5: Finish Switchover

1 Configure any applications that integrate with NCM to connect to the working NCM cores in the disaster recovery location.

2 Notify users to connect to the NCM Console on the NCM cores in the disaster recovery location.

3 *Optional*. Configure the managed devices to send syslog messages to one NCM core in the disaster recovery location.

Without syslog messages, NCM will detect configuration changes at the next scheduled snapshot.

Figure 3 shows the results of switching over to the disaster recovery location.

The **Allow this core to run all tasks created on it locally** setting on the NCM cores in the primary location does not affect the distribution of tasks to the NCM cores in the disaster recovery location. The following CLI command can be used to move a task to a different NCM core:

```
mod task -id <Task ID> -coreid <Core ID>
```

Figure 3 After Switchover

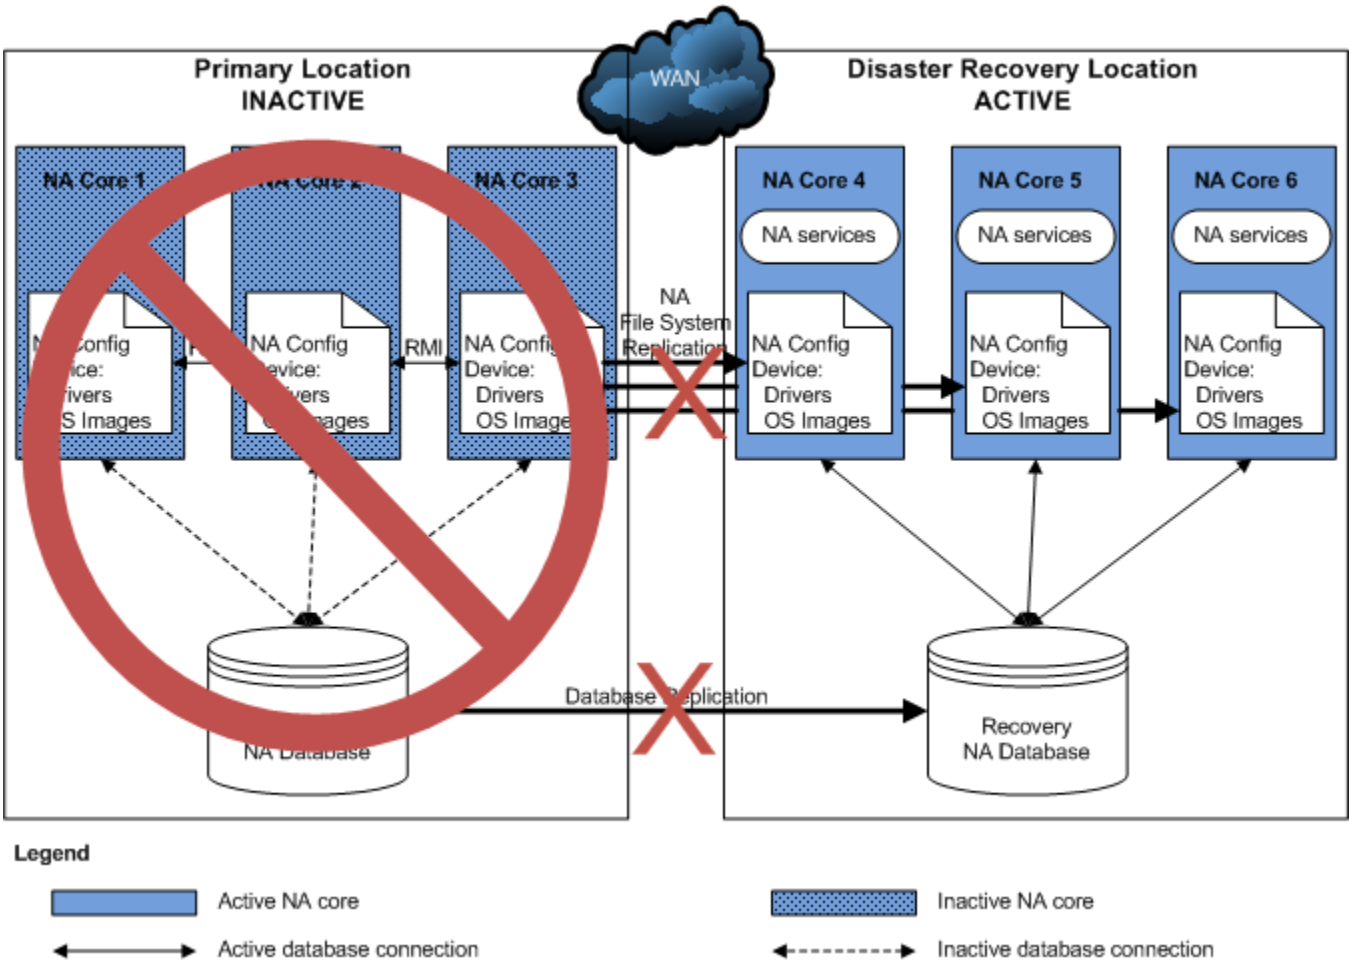

4 Switchback

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 chapter describes the following switchback scenarios:

- When the primary location again becomes available, an administrator can follow the procedure described in Switching Back to the Original Servers in the Primary Location on page 27 to switch the use of CiscoWorks Network Compliance Manager (NCM) back to the primary location.
- If the systems in the primary location cannot be recovered, configure new servers in the primary location or a new primary location, and then follow the process in Switching Back to Different NCM and Database Servers on page 33 to switch NCM to that location.
- Alternatively, you can run the disaster recovery location as the new primary location and configure a new disaster recovery location as described in Creating a New Disaster Recovery Location on page 42.

Switching Back to the Original Servers in the Primary Location

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

- Task 1: Disable NCM in the Disaster Recovery Location on page 27
- Task 2: Enable the Use of the Database in the Primary Location on page 28
- Task 3: Enable NCM in the Primary Location on page 30
- Task 4: Finish Switchback on page 31

Task 1: Disable NCM in the Disaster Recovery Location

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

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

1 Prepare to stop NCM in the disaster recovery location. Do the following:

   a Notify users to log out.

   b Log on to the NCM Console for one of the NCM cores in the disaster recovery location.

   c Pause tasks scheduled to start during the switchback process (Task 3, step 5 on page 31). Include time for the currently running tasks to complete. Also include time for database synchronization. For example:

      — On the Search for Task page (**Reports > Search For > Tasks**), for the Schedule Date field, set since to **`Until and anytime to 4 hours later`**.

      — 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 2, next.

2 Stop all NCM services on all NCM cores in the disaster recovery location.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Stop**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *UNIX*: Run the following command:

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

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

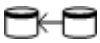

1 Synchronize the database in the primary location with the database in the disaster recovery location. Possible approaches include the following:

- Use database tools to copy the NCM database from the disaster recovery location to the primary location. Copy the NCM schema tables only.

  For example, you might follow this process:
  - Export the NCM database from the recovery database server in the disaster recovery location.
  - Wipe the NCM database from the database server in the primary location.
  - Import the NCM database to the database server in the primary location.
- 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.
  - If the transaction logs have overflowed, reverse replication becomes ineffective. In this case, disable reverse replication to the database in the primary location, and then use database tools to copy the NCM database from the disaster recovery location to the primary location. Copy the NCM schema tables only.
  - If the transaction logs are within bounds, wait for all database updates to replicate to the database in the primary location. After replication is complete, disable reverse replication to the database in the primary location.

2 Re-enable database replication from the primary location to the disaster recovery location.

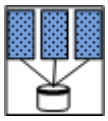

3 Connect the NCM cores in the disaster recovery location to the database in the primary location.

On each NCM server in the disaster recovery location, in a text editor, edit following file:

- *Windows*:

  `<NCM_HOME>\server\ext\jboss\server\default\deploy\db_ds.xml`

- *UNIX*:

  `<NCM_HOME>/server/ext/jboss/server/default/deploy/db_ds.xml`

This file contains two lines defining the `JdbcUrl` attribute. For example:

- Oracle:

  ```
  <attribute name="JdbcUrl">jdbc:oracle:thin:@db.example.com:
  1521:nadb</attribute>
  ```

- SQL Server:

  ```
  <attribute name="JdbcUrl">jdbc:sqlserver://db.example.com:
  1433;DatabaseName=NCM;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 primary location.

For Oracle, also replace the database SID or service name (**`nadb`** in the example) with the database SID or service name for the NCM database in the primary location.

For SQL Server, also replace the database name (**`NCM`** in the example) with the database name for the NCM database in the primary location.

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

4 Ensure that the NCM servers in the primary location are powered on.

5 If automatic starting of the NCM services is enabled, stop all NCM services on all NCM cores in the primary location.

- *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`

6 Verify that the NCM cores in the primary location are connected to the local database (the database in the primary location) as described in step 3.

Task 3: Enable NCM in the Primary Location

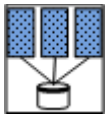

To enable NCM in the primary location, follow these steps:

1 If the configuration of the NCM cores in the disaster recovery location has changed since switchover, update the NCM core files on each NCM server in the primary location. Maintain file ownership and permissions. See Table 1 on page 17.

You will restart the NCM services in step 5 on page 31. You do not need to do so now.

2 Activate the NCM cores in the primary location.

a On *one* NCM core in the primary location, start all NCM services.

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

TrueControl ManagementEngine

TrueControl FTP Server

TrueControl SWIM Server

TrueControl Syslog Server

TrueControl TFTP Server

— *UNIX*: Run the following command:

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

b Using telnet or SSH, connect as a NCM administrator to the NCM proxy on that NCM core.

c Run the following command:

```
list core
```

d From the `list core` command output, determine the core IDs of the NCM cores in the primary location, and then run the following command for each core ID:

```
core status -status normal -coreid <coreid>
```

This step changes the NCM core status in the NCM database in the primary location. Because replication is running, the NCM database in the disaster recovery location also shows these cores as active.

3 If NCM is configured with tasks for a given device bound to only one core (the default Horizontal Scalability configuration), do the following:

a On the *one* running NCMcore in the primary location, restart all NCM 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
```

b Log on to the NCM Console for the running NCM core in the primary location.

c Update the site assignments. In the NCM Console, open the **Site Reassignment** page (**Admin > Distributed > Site Reassignment**), and then assign all partitions to NCM cores in the primary location.

4 Deactivate the NCM cores in the disaster recovery location.

a Connect as a NCMadministrator to the NCM proxy on the *one* running NCM core in the primary location.

b Run the following command:

```
list core
```

c From the `list core` command output, determine the core IDs of the NCM cores in the disaster recovery location, and then run the following command for each core ID:

```
core status -status standby -coreid <coreid>
```

This step changes the NCM core status in the NCM database in the primary location. Because replication is running, the NCM database in the disaster recovery location also shows these cores as inactive.

5 Start (or restart) all NCM services for all NCM cores in the primary location.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Start**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *UNIX*: Run the following command:

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

6 Resume the tasks that were paused in Task 1, step 1 on page 27.

a Log on to the NCM Console

a for one of the NCM 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.

Task 4: Finish Switchback

1 Configure any applications that integrate with NCM to connect to the working NCM cores in the primary location.

2 Notify users to connect to the NCM Console on the NCM 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 NCM files to an intermediate server, to conserve resources in the disaster recovery location, power down the NCM servers in the disaster recovery location.

5 *Optional*. Configure the managed devices to send syslog messages to one NCM core in the primary location and to no NCM cores in the disaster recovery location.

Figure 4 shows the results of switching back to the primary location.

The **Allow this core to run all tasks created on it locally** setting on the NCM cores in the primary location does not affect the distribution of tasks to the NCM cores in the disaster recovery location. The following CLI command can be used to move a task to a different NCM core:

```
mod task -id <Task ID> -coreid <Core ID>
```

Figure 4 After Switchback

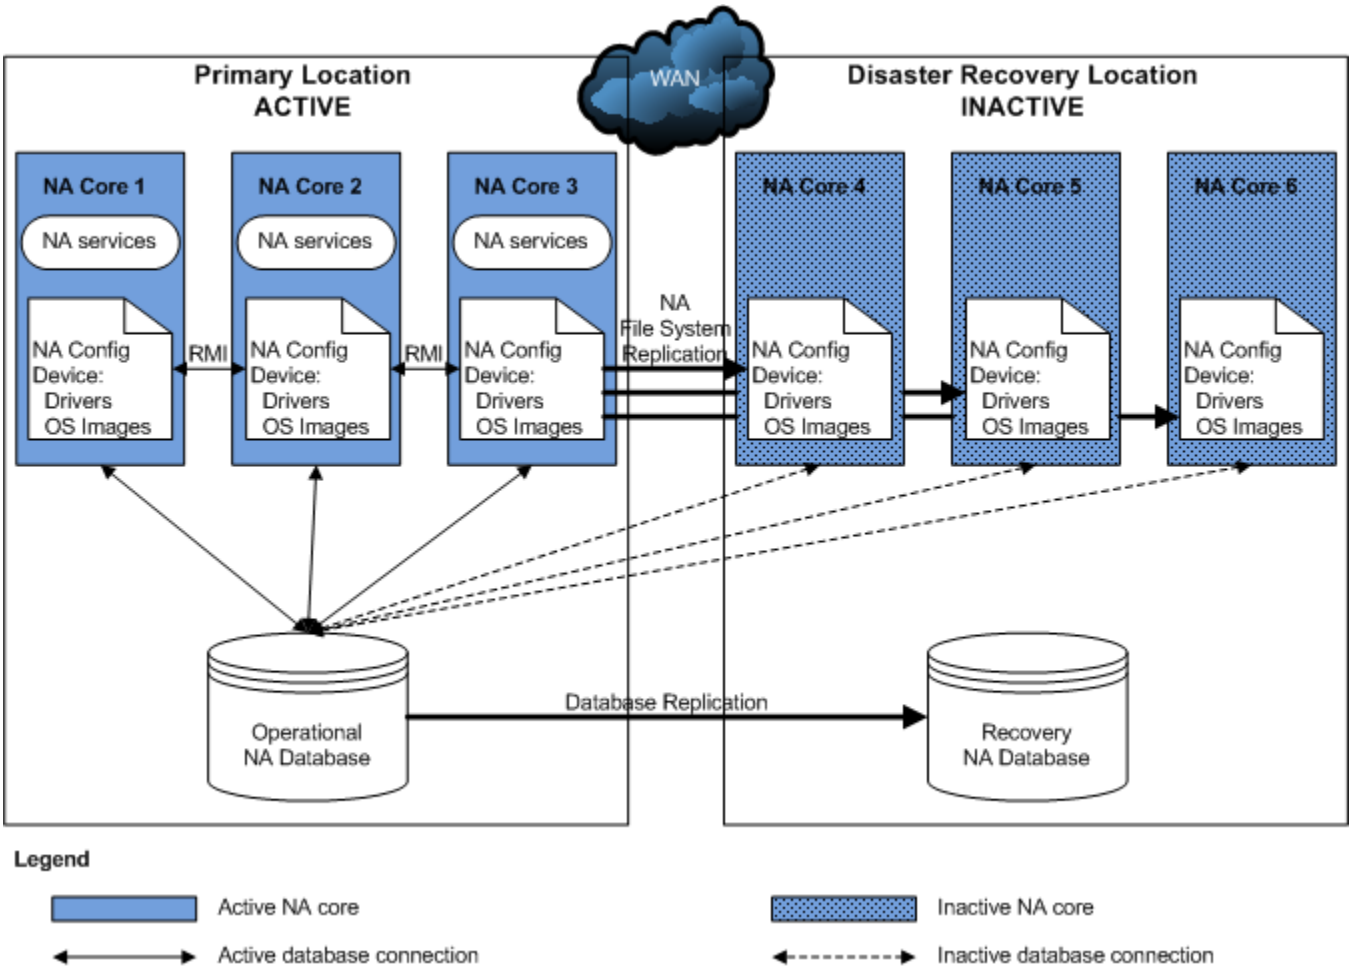

Switching Back to Different NCM and Database Servers

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 NCM cores could be running on newly-provisioned servers. Additionally, the NCM 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 NCM 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 NCM 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 NCM cores are NCM cores in the original primary location data center that are still available for switchback.
- The new NCM cores are NCM cores running on newly-provisioned servers in the primary location.
- The disaster recovery location is data center currently hosting the existing NCM 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 NCM in the Disaster Recovery Location on page 33
- Task 2: Enable the Use of the Database in the Primary Location on page 34
- Task 3: Enable NCM in the Primary Location on page 36
- Task 4: Finish Switchback on page 40

Task 1: Disable NCM in the Disaster Recovery Location

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

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

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

      ```
      list core
      ```

   c. Note the core ID for each unavailable NCM server.

2 Prepare to stop NCM in the disaster recovery location. Do the following:

a Notify users to log out.

b Log on to the NCM Console for one of the NCM cores in the disaster recovery location.

c Pause tasks scheduled to start during the switchback process (Task 3, step 9 on page 40). Include time for the currently running tasks to complete. Also include time for database synchronization. For example:

— On the Search for Task page (**Reports > Search For > Tasks**), for the Schedule Date field, set since to **`Until`** and **anytime** to **`4 hours later`**.

— 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 NCM services on all NCM cores in the disaster recovery location.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Stop**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *UNIX*: Run the following command:

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

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

1 If the database server is newly-provisioned, create the NCM schema on that database server.

Note the following:

- The NCM database user in the disaster recovery location must have the same name and permissions as the NCM database user in the primary location.
- Copy the NCM schema tables only.
- For Oracle or Oracle RAC, the SID or service name should be different between the two database servers.

  One SID *cannot* be a subset of the other SID, for example NCMRp and NCMRpBU. Instead, use SIDs that stand alone, for example NCMRp1 and NCMRp2.
- For Microsoft SQL Server, the database name should be different between the two database servers.

For example, you might follow this process:

a Install the database software.

b Create a database user with the same name and permissions as the NCM database user for the database in the disaster recovery location.

2 Synchronize the database in the primary location with the database in the disaster recovery location. Possible approaches include the following:

- Use database tools to copy the NCM database from the disaster recovery location to the primary location. Copy the NCM schema tables only.

  For example, you might follow this process:

  — Export the NCM database from the recovery database server in the disaster recovery location.

  — Wipe the NCM database from the database server in the primary location.

  — Import the NCM database to the database server in the primary location.

- 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.

  — If the transaction logs have overflowed, reverse replication becomes ineffective. In this case, disable reverse replication to the database in the primary location, and then use database tools to copy the NCM database from the disaster recovery location to the primary location. Copy the NCM schema tables only.

  — If the transaction logs are within bounds, wait for all database updates to replicate to the database in the primary location. After replication is complete, disable reverse replication to the database in the primary location.

3 Configure database replication as follows:

- If the NCM database server is newly-provisioned, update database replication as follows:

  — On the disaster recovery location database server, remove the configuration for database replication from the original primary location database server.

  — Configure one-way database replication from the database in the primary location to the recovery database in the disaster recovery location.

- If the NCM database server is *not* newly-provisioned, re-enable database replication from the primary location to the disaster recovery location.

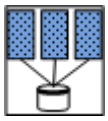

4 Connect the NCM cores in the disaster recovery location to the database in the primary location.

On each NCM server in the disaster recovery location, in a text editor, edit following file:

- *Windows*:

  `<NCM_HOME>\server\ext\jboss\server\default\deploy\db_ds.xml`

- *UNIX*:

  `<NCM_HOME>/server/ext/jboss/server/default/deploy/db_ds.xml`

This file contains two lines defining the `JdbcUrl` attribute. For example:

- Oracle:

  ```
  <attribute name="JdbcUrl">jdbc:oracle:thin:@db.example.com:
  1521:nadb</attribute>
  ```

- SQL Server:

  ```
  <attribute name="JdbcUrl">jdbc:sqlserver://db.example.com:
  1433;DatabaseName=NCM;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 primary location.

For Oracle, also replace the database SID or service name (**nadb** in the example) with the database SID or service name for the NCM database in the primary location.

For SQL Server, also replace the database name (**NCM** in the example) with the database name for the NCM database in the primary location.

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

5 Ensure that all existing NCM servers in the primary location are powered on.

6 If automatic starting of the NCM services is enabled for the existing NCM cores, stop all NCM services on all NCM cores in the primary location.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Stop**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *UNIX*: Run the following command:

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

7 Connect the existing NCM cores in the primary location to the local database (the database in the primary location) as follows:

- If the NCM database server is newly-provisioned, update the `db_ds.xml` file as described in step 4.
- If the NCM database server is *not* newly-provisioned, verify the `db_ds.xml` file as described in step 4.

8 In the primary location database, remove the unavailable NCM servers from the NCM database.

a Locate the list of core IDs for the unavailable NCM servers, as determined in Task 1, step 1 on page 33.

b For each unavailable NCM server in the original primary location, delete the entry for that NCM server from the RN_CORE table of the NCM database.

For example on Oracle:

```
DELETE FROM RN_CORE WHERE CoreID = <coreid>;
COMMIT;
```

For example on SQL Server:

```
DELETE FROM RN_CORE WHERE CoreID = <coreid>;
```

Task 3: Enable NCM in the Primary Location

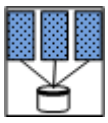

To enable NCM in the primary location, follow these steps:

1 If the configuration of the NCM cores in the disaster recovery location has changed since switchover, update the NCM core files on each existing NCM server in the primary location. See Table 1 on page 17.

You will restart the NCM services in step 9 on page 40. You do not need to do so now.

2 Verify that you have the *NCM Horizontal Scalability Guide* for NCM 1.8 or later.

3 In the primary location, install the new NCM cores as additional NCM cores connected (through Horizontal Scalability) to the database in the primary location.

For information, see "Adding Additional NCM Cores" in the *NCM Horizontal Scalability Guide* for NCM 1.8 or later.

Complete the script edits to account for having removed one or more NCM cores from the Horizontal Scalability environment.

4 Configure the new NCM cores as follows:

a Stop all NCM services on the new NCM cores.

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

TrueControl ManagementEngine

TrueControl FTP Server

TrueControl SWIM Server

TrueControl Syslog Server

TrueControl TFTP Server

— *UNIX*: Run the following command:

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

b Put into place any modified files from those listed in Table 1 on page 17. Options include:

— Copy the files from one of the NCM servers in the disaster recovery location.

— Retrieve the files from a backup.

c Start all NCM services on the new NCM cores.

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

TrueControl ManagementEngine

TrueControl FTP Server

TrueControl SWIM Server

TrueControl Syslog Server

TrueControl TFTP Server

— *UNIX*: Run the following command:

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

d *Optional*. Perform any NCM core-specific tuning. Restart the NCM services as needed.

5 If the NCM environment includes NCM 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 NCM core. During installation, configure each core gateway in the primary location as follows:

— Use the same Gateway Crypto Data file as for the core gateways in the original primary location.

— Assign the same realm name, typically Default Realm, to each core gateway.

For information, see the *NCM Satellite Guide*.

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

For each core gateway that remains from the original primary location configuration and is not installed on a NCM server, connect that core gateway with a NCM server as described in "Configuring NCM to Communicate with the Core Gateway" of the *NCM Satellite Guide*.

c For each satellite in the gateway mesh, update the satellite configuration to enable communication with a core gateway in the primary location.

Edit the remote gateway configuration file:

`<gateway_install_dir>/opswgw-<gateway_name>/opswgw.properties`
(The default value of `<gateway_install_dir>` is `/etc/opt/opsware`.)

In the remote gateway configuration file, 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
```

d Restart each remote core gateway.

6 Activate the existing NCM cores in the primary location.

a Connect as a NCM administrator to the NCM proxy on a NCM core running on a newly-provisioned NCM server.

b Run the following command:

```
list core
```

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

```
core status -status normal -coreid <coreid>
```

This step changes the NCM core status in the NCM database in the primary location. Because replication is running, the NCM database in the disaster recovery location also shows these cores as active.

7 If NCM is configured with tasks for a given device bound to only one core (the default Horizontal Scalability configuration), do the following:

a On *one* running NCM core in the primary location, restart all NCM 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
```

b Log on to the NCM Console for the NCM core that was just restarted in the primary location.

c Update the site assignments. In the NCM Console, open the **Site Reassignment** page (**Admin > Distributed > Site Reassignment**), and then assign all partitions to NCM cores in the primary location.

8 Deactivate the NCM cores in the disaster recovery location.

a Connect as a NCM administrator to the NCM proxy on the NCM core that was recently restarted in the primary location.

b Run the following command:

```
list core
```

c From the `list core` command output, determine the core IDs of the NCM cores in the disaster recovery location, and then run the following command for each core ID:

```
core status -status standby -coreid <coreid>
```

This step changes the NCM core status in the NCM database in the primary location. Because replication is running, the NCM database in the disaster recovery location also shows these cores as inactive.

9 Start (or restart) all NCM services for all NCM cores in the primary location.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Start**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *UNIX*: Run the following command:

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

10 Resume the tasks that were paused in Task 1, step 2 on page 34.

a Log on to the NCM Console for one of the NCM 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.

Task 4: Finish Switchback

1 Configure any applications that integrate with NCM to connect to the working NCM cores in the primary location.

2 Notify users to connect to the NCM Console on the NCM cores in the primary location.

3 Configure replication of the NCM file system from one of the active NCM cores to the inactive NCM cores (or an intermediate server) as described in Files to Synchronize Across NCM Cores on page 16.

4 *Optional*. If you chose to synchronize the NCM files to an intermediate server, to conserve resources in the disaster recovery location, power down the NCM servers in the disaster recovery location.

5 *Optional*. Configure the managed devices to send syslog messages to one NCM core in the primary location and to no NCM cores in the disaster recovery location.

6 *Optional*. For the NCM cores on the newly-provisioned NCM servers, consider the risk that if the primary location is not accessible during switchover to the disaster recovery location, the NCM cores in the primary location cannot be deactivated. In this case, two NCM cores (one each in the primary and disaster recovery locations) might run the same task. To mitigate this risk, on each NCM server in the primary location, disable automatic starting of the NCM services.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, click **Properties**, and then set Startup Type to **Disabled**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *Linux*:

  ```
  mv /etc/rc.d/rc5.d/S99truecontrol /etc/S99truecontrol
  ```
- *Solaris*:

  ```
  mv /etc/rc2.d/S99truecontrol /etc/S99truecontrol
  ```

Figure 4 on page 32 shows the results of switching back to the new primary location.

The **Allow this core to run all tasks created on it locally** setting on the NCM cores in the primary location does not affect the distribution of tasks to the NCM cores in the disaster recovery location. The following CLI command can be used to move a task to a different NCM core:

```
mod task -id <Task ID> -coreid <Core ID>
```

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 NCM deployment. This location is the original disaster recovery location.
- The new disaster recovery location is the data center that will receive the new NCM 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 NCM cores in the primary location cannot be deactivated. In this case, two NCM cores (one each in the primary and disaster recovery locations) might run the same task. To mitigate this risk, on each NCM server in the primary location, disable automatic starting of the NCM services.
   - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, click **Properties**, and then set Startup Type to **Disabled**:
     - **TrueControl ManagementEngine**
     - **TrueControl FTP Server**
     - **TrueControl SWIM Server**
     - **TrueControl Syslog Server**
     - **TrueControl TFTP Server**
   - *Linux*:

     ```
     mv /etc/rc.d/rc5.d/S99truecontrol /etc/S99truecontrol
     ```
   - *Solaris*:

     ```
     mv /etc/rc2.d/S99truecontrol /etc/S99truecontrol
     ```
2. Determine the core IDs of the NCM cores in the original primary location that are no longer available.
   a. Connect as a NCM administrator to the NCM proxy on a NCM core running in the disaster recovery location.
   b. Run the following command:

      ```
      list core
      ```
   c. Note the core ID for each unavailable NCM server.
3. *Optional*. Perform any tuning needed to prepare the NCM cores in the new primary location for daily use. Restart the NCM services as needed.
4. *Optional*. Configure the managed devices to send syslog messages to one NCM core in the new primary location.

5 Prepare to stop NCM in the new primary location.

a Notify users to log out.

b Log on to the NCM Console for one of the NCM cores in the new primary location.

c Pause tasks scheduled to start during the disaster recovery configuration process (until Task 4, step 1 on page 15). Include time for the currently running tasks to complete. Also include time for database synchronization. For example:

— On the Search for Task page (**Reports > Search For > Tasks**), for the Schedule Date field, set since to **`Until and anytime`** to **`4 hours later`**.

— 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 NCM services on all NCM cores.

- *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Stop**:
  - **TrueControl ManagementEngine**
  - **TrueControl FTP Server**
  - **TrueControl SWIM Server**
  - **TrueControl Syslog Server**
  - **TrueControl TFTP Server**
- *UNIX*: Run the following command:

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

7 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 NCM servers from the NCM database.

a Locate the list of core IDs for the unavailable NCM servers, as determined in step 2 on page 42.

b For each unavailable NCM server in the original primary location, delete the entry for that NCM server from the RN_CORE table of the NCM database.

For example on Oracle:

```
DELETE FROM RN_CORE WHERE CoreID = <coreid>;
COMMIT;
```

For example on SQL Server:

```
DELETE FROM RN_CORE WHERE CoreID = <coreid>;
```

9 Remove any remaining configuration for replication of the NCM file system from one of the NCM cores in the original primary location to the original disaster recovery NCM cores (or an intermediate server).

10 Verify that you have the *NCM Horizontal Scalability Guide* for NCM 1.8 or later.

11 Beginning with Task 2: Configure Database Replication on page 13 and through the end of Chapter 2, NCM Disaster Recovery Initial Setup, complete the process for setting up NCM for disaster recovery.