Hewlett Packard Enterprise

HPE Network Automation Software

Software Version: 10.20
Windows® and Linux® operating systems

Multimaster Distributed System on Oracle Guide

Legal Notices

Warranty

The only warranties for HPE 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

Copyright Notice

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 installer bundle.

Documentation Updates

The title page of this document contains the following identifying information:

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

To check for recent updates or to verify that you are using the most recent edition of a document, go to: https://softwaresupport.hpe.com/group/softwaresupport/search-result?keyword=.

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

Support

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

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

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

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

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to **https://softwaresupport.hpe.com** and click **Register**.

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

Contents

Chapter 1: Getting Started

> **Note:** HPE Network Automation Software (NA) running in a Multimaster Distributed System environment does not support the following features:
>
> - The serial run mode that was added to many task types in NA 10.20
> - The Deploy Hotfix task, new in NA 10.20
> - Integration with HPE Network Node Manager i Software (NNMi)

This chapter contains the following topics:

- "Terminology" below
- "Overview" below
- "Architecture" on the next page
- "Is the Multimaster Distributed System Functionality Right for You?" on page 10
- "NA Internal Monitoring: Oracle Advanced Replication" on page 11
- "NA Internal Monitoring: Inter-NA Core Communication" on page 11

Terminology

The following terms are used throughout this guide:

- **NA Core** — A single NA Management Engine, associated services (Syslog and TFTP), and a single database. An NA Core can manage multiple Partitions.
- **Partition —** A set of devices with unique IP addresses. A Partition is managed by one (and only one) NA Core. Multiple Partitions can be managed by a single NA Core. For information about segmenting devices, see the *NA User Guide*.
- **NA Mesh** — Multiple NA Cores connected via multimaster replication.
- **Multimaster replication** — Replication where all data is kept consistent across all database servers.
- **Master Definition Site** — Oracle replication requires that one database in a multimaster replication NA Mesh hold the definition of the replication setup. This database is called the Master Definition Site.

Overview

The NA Distributed System on Oracle is a multimaster system where the data from each NA Core in an NA Mesh is accessible to all other NA Cores. This provides a comprehensive view of your data and allows for redundant data and failover in the event of a problem with a single NA Core. The multimaster features include both database data and certain file system data, such as software images and device driver packages. Keep in mind that software images and device driver packages are also replicated across the NA Mesh.

The following comprises the NA Distributed System:

- An NA Core and a Site:
  - A device is associated with a single Site.

  - A Site is associated with a single NA Core.

  - A task is associated with a specific NA Core.

- Multimaster replication on Oracle:
  - Conflict resolution in the event that the same piece of data is modified almost simultaneously on two separate NA Cores. This conflict is typically resolved using the latest timestamp method.

  - Replication monitoring and conflict notification is built into NA. You can manage Oracle replication conflicts and view replication job statuses from within the NA UI.

  - The NA scheduler is multi-core aware. You can schedule group tasks containing devices that are associated with different NA Cores. The system will run these tasks on the correct NA Core. You do not have to schedule tasks on the appropriate NA Core.

System setup requires a thorough understanding of Oracle and NA. Installation includes number of steps that must be performed on the various servers. In addition, certain network changes may be necessary to allow connections between the servers. Once setup is complete, you will need to partition your devices into Sites to ensure proper NA Core access to devices. For information about segmenting devices, see the *NA User Guide*.

Architecture

The following figure shows a typical two-NA Core Multimaster Distributed System installation. The installation enables failover from one NA Core to another, while the remaining NA Core continues to manage all devices.
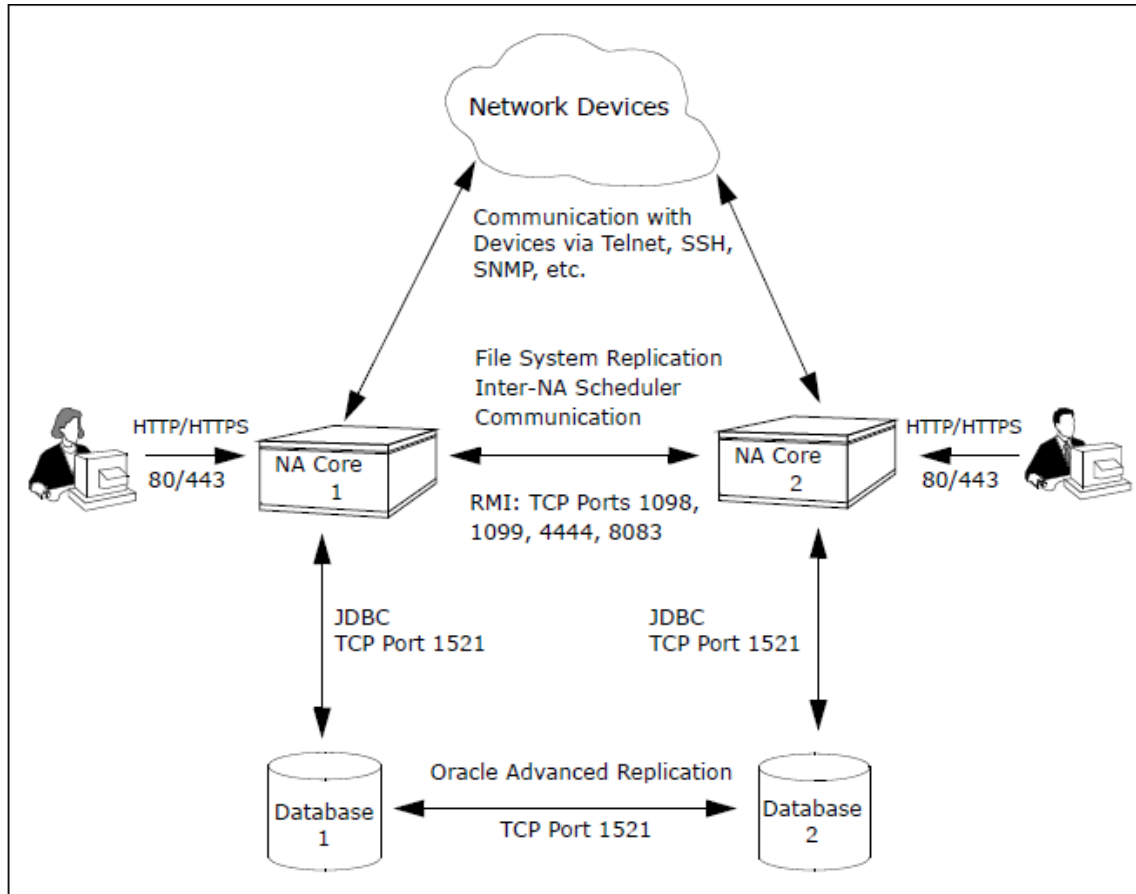
Keep in mind that some installations might have two separate network device collections that are inaccessible from the non-managing NA Core. These installations still provide data redundancy and scaling benefits, however.

Note that the database data transport uses standard Oracle Advanced Replication. For details on how replication works, see "System Administration" on page 37.

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

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

NA does not replicate NA Core specific options, including:

- TFTP server information
- TACACS server information
- Scripting language settings
- Hop Box definitions
- Driver directory
- Install directory
- Local Gateway information
- Active Directory settings
- Scheduler (concurrent task limits)

Is the Multimaster Distributed System Functionality Right for You?

Horizontal Scalability is the ability to combine multiple NA cores with a single database so that they work as a single logical unit to improve the performance of the overall system. Any NA console user sees the same information on each of the NA cores. NA distributes device tasks across the NA cores according to the Horizontal Scalability topology. Users do not need to know which NA core accesses a particular device. NA Horizontal Scalability can be configured in different ways, depending on how you use your system and where bottlenecks are apt to occur. For installation and configuration information, see the *NA Horizontal Scalability Guide*.

Horizontal Scalability is the recommended deployment architecture for NA. It provides the following benefits:

- NA core failover of tasks in the case that one NA core stops working or becomes unreachable
- Optional disaster recovery configuration for minimizing the business disruption should a significant event affect an entire data center.

Additionally, you can use database clustering and NA satellites with Horizontal Scalability to increase the reliability and reach of the NA management system.

A Multimaster Distributed System is intended for environments with needs beyond what Horizontal Scalability provides. These situations can include:

- Your organization requires that all NA cores always be actively used for device management.

  In this case, the network latency between the NA cores and the NA database servers must be low. Additionally, network down times must be infrequent and short.

- Users must access NA locally to avoid a higher latency WAN. These users manage local devices.

  In this case, the NA cores and database servers can be located in multiple data centers. Database replication is the only NA traffic on the WAN.

Tip: A Multimaster Distributed System deployment is most successful when there is a clear division of responsibility for device management. It is best when all NA work for a given device is initiated from only one NA core. NA work includes:

- Tasks
- Device modifications through the CLI, the NA console, or integrations with other products

Additionally, each database object (for example, a task) should be changed only on the NA core on which it was created.

This approach prevents conflicting changes in multiple NA databases and reduces the need for database conflict resolution.

NA Internal Monitoring: Oracle Advanced Replication

Oracle Advanced Replication can be configured to automatically handle certain types of conflicts. Where possible, NA uses Oracle's built-in conflict handling software. For example, for most tables with *update/update* conflicts, NA installs the default Oracle "latest timestamp wins" handler.

When it is not possible to use Oracle's default handler, NA either adds a stored procedure to handle the conflict (in particular, for *update/delete* conflicts), or adds additional code inside the application with logic that attempts to correct the error. If the error is uncorrected through that logic, NA alerts the NA administrator to the issue using NA' standard event notification emailing functionality. For information about configuring email notification, see the *NA User Guide*.

> **Note:** When a second NA Core is used as a standby for failover and not for normal operation, you should not encounter *update/update* conflicts.

For detailed information about NA monitors, see "Understanding Oracle Replication" on page 37.

In addition, in Oracle, NA monitors on a regular schedule the following on each NA Core for the database transport (i.e., replication):

- Push jobs that handle moving data from one NA Core to another. If the monitor check fails and the Push job is found to be broken, NA will attempt to restart it.
- Purge jobs that successfully clear pushed data from the push queue, and attempts to restart the Push queue each time the check fails.
- The deferred transaction queue for any transactions that could not be applied to the local database due to a conflict. NA applies additional application logic to attempt to correctly apply the transaction. If unsuccessful, NA alerts the NA administrator via email notification.
- Unrecoverable uniqueness conflicts. NA alerts the NA administrator in the event of unrecoverable uniqueness conflicts. Typically, there is no way to automatically correct these types of issues.
- Both the DEFTRAN and DEFLOB queue size. These queues hold pending data that needs to be pushed to other servers. Typically, if replication is working, but is beginning to backup, these queues will begin to grow. This can occur if the connection between databases is congested. NA alerts the NA administrator to this issue via email notification. The system will continue to push data. However, it could take a long time to catch up depending on network conditions.

> **Note:** For information about configuring email notification, see the *NA User Guide*.

NA Internal Monitoring: Inter-NA Core Communication

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

- RMI connectivity
- NA server timestamp differences

These monitors generate events during error conditions. The events can be emailed to the NA administrator or Oracle DBA using a standard event rule in NA. For information about configuring email notification, see the *NA User Guide*. For detailed information about NA system events, see "NA Generated Events" on page 41.

> **Note:** NA includes an example event rule that can be updated as necessary.

Chapter 2: Installation, Setup, and Upgrading

The chapter contains the following topics:

- "System Requirements" below
- "Prepare for Initial Replication" below
- "Oracle 11g: Set Up Initial Replication" on page 16
- "Oracle 11g: Adding a New NA Core to an Existing NA Mesh" on page 26
- "Removing an NA Core from an Existing NA Mesh" on page 29
- "Removing Replication Across the Entire NA Mesh" on page 31
- "Upgrade the NA Mesh" on page 32

For information about troubleshooting an NA Mesh, see "Troubleshooting" on page 53.

System Requirements

For information about database version and server hardware requirements for the NA Multimaster Distributed System on Oracle software, see the *NA Support Matrix*.

Prepare for Initial Replication

To properly install the Distributed System software, you must first complete:

- Device partitioning planning across NA Cores
- Network configuration planning for connectivity between NA servers and devices. For example, what network connectivity is required to support failover for device access? The ability of an NA Mesh to failover for device access depends in part on proper network setup to ensure access to devices. In some cases, you might not want to have failover work for complete device access, but instead have it ensure access to data while corrective action is taken to restore the network connectivity to the affected NA Core.
- Network configuration planning for connectivity and bandwidth between the different servers (NA and database) that comprise the NA Mesh and between the NA Management Engines and devices. The NA Cores in the NA Mesh will also need bandwidth between them equal to the bandwidth provided between a single NA server and its database in a single NA Core. For more information, see "Architecture" on page 8.
- Network configuration planning for bandwidth usage by the NA Mesh. The bandwidth required between an NA Management Engine and its database depends on the size of the deployment, including the number of devices, average configuration size, and the number of concurrent users. The bandwidth available between the Oracle databases should be as large as between the NA Management Engine and the database. For a deployment of 14,000 devices with an average 10K configurations that change once a week, the bandwidth requirements could conservatively be on the order of 1Mb/s. Peak bandwidth requirements could be higher, depending on usage.

- Oracle server setup planning. Keep in mind that the database properties required for replication can be set on initial database creation. You do not need to wait until replication setup to set these parameters.

Keep in mind that during setup, data is dumped from an existing NA server. This could take time and use disk space as the database is copied. In addition, when data is copied from one database to another, there must be ample bandwidth between the two (or more) servers. It will also take time to import the exported database into the other NA Cores.

You can estimate the time it will take to copy data from server to server given the bandwidth between the servers. You can also calculate the disk space requirements for the export (and subsequent import) operations by knowing the size of your database. Note that if you want to export or import data from the same server as the database, the disk space requirement is twice the size of your database.

> **Note:** Estimating time for the import and export operations could be difficult. You should allocate a lengthy time frame to complete this work. In addition, the NA server(s) must be off during the export and import steps. A sufficiently long maintenance window should be planned for.

You will also have to ensure:

- Time synchronization setup for the NA servers in the NA Mesh
- Users are instructed to login to their "closest" NA Core
- Access to an Oracle DBA to support the NA Distributed System installation

To assist in planning, please note the following limitations and suggestions concerning the NA Distributed System:

- The NA database user name must be a custom value (not SYSTEM).
- NA currently does not support joining multiple existing NA installations into an NA Mesh. You can only create an NA Mesh from a single existing NA server, adding new NA Cores as appropriate. You can also create an NA Mesh from scratch.
- Users should not share logon accounts. Due to the replication system used to share data across NA Cores, two users should never use the same login name to connect to two different NA Cores at the same time. If they do so, the system will likely require additional work by the system administrator to ensure that the affected user's profile is properly synchronized across the NA Mesh.
- NA currently assumes that all servers in the Mesh (NA and database) share a single time zone.
- Future NA upgrades will take longer and require more downtime due to the need to both update the replication setup and to update all servers in the NA Mesh. NA does not support rolling upgrades where one part of the NA Mesh is running a version of NA while the rest of the NA Mesh is running a different version.
- The ReplicationScript.sql script updates the RN_CORE table. When the replication script runs, the RN_CORE table changes. There is no need for it to UPDATE and INSERT into the RN_CORE table on both Cores because replication is already running. The database changes will be pushed via replication to the second database. The contents of the RN_CORE tables on both databases should match whatever is set up in *ReplicationScript.sql*.

Planning

This section includes information for initial replication setup on a Master Definition Site (Master 1) and one additional Master Site (Master 2). NA will be offline during initial replication. It is recommended that you carefully read the following sections so as to streamline the replication process. A new database instance on Master 2 is manually created when running initial replication setup.

When initially setting up replication, complete the following table before starting. Be sure to save this information. You will need it each time you upgrade in the future.

Table 1 Distributed System Server Information

Information to Collect	NA Core 1	NA Core 2
Replication User (default: repadmin)		
Replication User Password		
Timezone Offset from UTC		Same for all servers in the NA mesh
Application Server Data	**NA Application Server 1**	**NA Application Server 2**
NA App Server FQDN		
NA App Server RMI Port		
Database Server Data	**Master 1 (Database Server 1)**	**Master 2 (Database Server 2)**
Database Server FQDN		
Oracle Listening Port		
Database SID		
Database User		Same as for Master 1
Database Password		Same as for Master 1

Protocols, Databases, and Ports

NA communicates with devices using a combination of the following protocols, databases, and ports. If you use a given protocol, NA requires access to the corresponding port. Specifically, if NA communicates with devices protected by firewalls, these ports must be open. For more information, see "Architecture" on page 8.

Protocol/Database/Port	From/To
NA server (running the Mgmt Engine, Syslog, TFTP) and network devices	
Telnet (port 23)	From the NA server to network devices.
SSH (port 22)	From the NA server to network devices.
TFTP (port 69/udp)	From network devices to the NA server.
Syslog (port 514/udp)	From network devices to the NA server.
SNMP (port 161/udp)	From the NA server to network devices.
Between the database servers	
Oracle (port 1521)	In a Distributed System configuration, the Oracle processes connect to

Protocol/Database/Port	From/To
	each other on port 1521.
Between the NA servers	
JNDI (ports 1098, 1099)	NA server to NA server. You can change this by editing the NA configuration files. Please contact Customer Support for assistance.
RMI (ports 4444 and 8083)	NA server to NA server. You can change this by editing the NA configuration files. Please contact Customer Support for assistance.
Between the NA server and the database server	
Oracle (port 1521)	From the NA server to an Oracle database server.
NA server and NA users	
HTTPS (port 443)	From the NA server to NA users. You can change this by editing the NA configuration files. Please contact Customer Support for assistance.
Telnet (port 23 - Windows or 8023 - Linux)	From the NA client to the NA server. This can be changed from the Administrative Settings option.
SSH (port 22 - Windows or 8022 - Linux)	From the NA client to the NA server. This can be changed from the Administrative Settings option.

Unpacking the Replication Script Installation Bundle

The setup files for the NA Multimaster Distributed System are the standard setup files for a normal NA installation, with the addition of the `OracleReplicationScriptToolBundle.zip` file. This file should be unzipped onto a standard NA core server. The setup files include a Java program that you run to generate a replication setup SQL script. This program uses the Java runtime that is available on the NA core server.

Oracle 11g: Set Up Initial Replication

Oracle 11g introduces a change that requires recreation of the database schema after exporting that schema for use on another database server. This procedure accommodates the new behavior of Oracle 11g.

NA can already be installed on one core server and database server (a single NA core). Alternatively, you can complete the first NA core installation as part of setting up initial replication.

This procedure uses the following identifiers:

- NA core servers: `NA1` and `NA2`
- Oracle database SIDs:
  - `SID1` on database server Master 1
  - `SID2` on database server Master 2
- NA database user for both databases: `nauser`

- Password for the NA database user: `napass`
- NA tablespace: `NAORA`
- Script to create the NA database user: `create_nauser.sql`
- NA schema export file: `NAEXP.dmp`

To use a different identifier, substitute the actual value for each instance of the example value within the procedure.

To set up initial replication between two NA cores, complete the following tasks:

Task 1: Verify Prerequisites (Oracle 11g)

Prepare two database servers and two core servers. For best performance, each core server-database server pair should be co-located. Additionally, the core server-core server link and the database server-database server link should each be a high-speed connection.

Verify that the following prerequisites have been met:

- The two database servers are running Oracle 11g with the following configuration:
  - Unique SIDs

    > **Note:** One SID *cannot* be a subset of the other SID, for example NARp and NARpBU. Instead, use SIDs that stand alone, for example NARp1 and NARp2.

  - The following parameters:
    - global_names = TRUE
    - open_links = at least 4 (increase this by 1 for each additional server in the NA Mesh beyond 2)
    - open_links_per_instance = at least 4 (same as open_links)
    - parallel_min_servers = 2
    - shared_pool_size = increase current value by 40m to support replication
    - job_queue_processes = at least 2

      > **Note:** The parameters must be set on each database in the NA Mesh. The parameters could be set up when the database is first created. If not, they must be modified at this point. You can view the current value of these parameters using Oracle Enterprise Manager or SQLPlus with the "SHOW PARAMETERS" or "SELECT NAME, VALUE FROM v$parameter" commands (where NAME = <parameter_name>, for example, global_names).

- The two core servers that will run NA have working hostnames. Note the following:

  - Each core server should have a high-speed connection to one of the database servers.

  - For core server hardware and operating system requirements, see the *NA Support Matrix*.

- It is recommended that the host names of both database servers and both core servers are in the `hosts` file on each core server. This file is located as follows:

  - *Windows*: `<Drive>:\Windows\System32\drivers\etc\hosts`

  - *Linux*: `/etc/hosts`

- All database servers and core servers are set to use the same time and time zone.

- If NA is already installed on one NA core, verify that the database user in place for this core has at least the permissions listed in "NA Database User Creation Script (Oracle 11g)" on page 26.

Task 2: (New Installations Only) Create the NA Schema on Master 1 (Oracle 11g)

If NA has not yet been installed, create the NA schema on the Master 1 database server by following these steps:

1. Log on to Master 1 as the SYS user, and then create a new tablespace dedicated to NA on Master 1.

   For example, create the NAORA tablespace for SID1 by using the following command:

   - *Windows*:

     ```
     CREATE TABLESPACE NAORA DATAFILE 'C:\APP\ADMINISTRATOR\ORADATA\SID1\
     NAORA.DBF' SIZE 200M REUSE AUTOEXTEND ON NEXT 50M;
     ```

   - *Linux*:

     ```
     CREATE TABLESPACE NAORA DATAFILE '/u01/app/oracle/oradata/SID1/
     NAORA.DBF' SIZE 200M REUSE AUTOEXTEND ON NEXT 50M;
     ```

2. Create a script for creating the NA database user as follows:

   a. Using a text editor such as WordPad or vi, copy the text in "NA Database User Creation Script (Oracle 11g)" on page 26 into a new text file, for example `create_user.sql`.

   b. Customize the script to your environment by making the following changes:

      - Replace `NAORA` with the tablespace name created in step 1.

      - Replace each instance of `nauser` with the desired NA database user name.

      - Replace `napass` with the desired password for the NA database user.

   c. Place a copy of the script on Master 1 and Master 2.

      - *Windows*: Place the file in `C:\`.

      - *Linux*: Place the file in the `$ORACLE_HOME/bin` directory, for example `/u01/app/oracle/product/11.2.0/dbhome_1/bin`.

3. On Master 1, create a new NA database user with the privileges listed in "NA Database User Creation Script (Oracle 11g)" on page 26 that uses the tablespace name created in step 1:

    a. Log on to Master 1 as the SYS user.

    b. At the `sqlplus` prompt, run the user creation script.

        For example: run: `SQL > @create_user.sql`

4. Install NA on the first core server (`NA1`).

   During installation, supply the following database credentials:

- For the database IP address or hostname, enter the hostname of Master 1.

- For the database port, enter the port of Master 1 that NA will use to communicate with the database.

- For the database SID, enter the SID on Master 1 (for example, `SID1`).

- For the database user name and password, enter the name and password of the NA database user created in step 3, for example, `nauser` and `napass`.

5. After NA installation is complete, log on to NA (as user `nauser`) to ensure that it works as expected.

6. Stop all NA services on NA1. See "Start, Stop, or Restart All NA Services" on page 57.

Task 3: Prepare Master 1 for Replication (Oracle 11g)

To create the database export and finish preparing the database on Master 1, follow these steps:

1. On Master 1, create a directory to receive the tablespace export. For example, `C:\NA_data1` or `/tmp/NA_data1`.

- Ensure that the database user, for example `nauser`, has read-write access to this directory.

- If NA was running prior to beginning this configuration procedure, ensure that this directory has more available disk space than the NA tablespace consumes.

2. Log on to Master 1 as the NA database user, and then associate the new directory with the datapump directory in preparation for the tablespace export. For example:

- *Windows*:

```
sqlplus nauser/napass@SID1
create DIRECTORY datapump AS 'C:\NA_data1';
exit
```

- *Linux*:

```
sqlplus nauser/napass@SID1
create DIRECTORY datapump AS '/tmp/NA_data1';
exit
```

3. Log on to Master 1 as the SYS user, and then grant the NA database user read-write access to the datapump directory. For example:

```
sqlplus / as sysdba
GRANT READ,WRITE ON DIRECTORY datapump to nauser;
exit
```

4. On Master 1, verify that `$ORACLE_SID` is set to the SID for the NA database (for example, `SID1`).

5. With the NA management engine on NA1 shut down, on Master 1 and outside of SQLPlus, use the Oracle Data Pump Utility to export the NA database schema. For example:

```
expdp nauser/napass TABLES=RN_% DIRECTORY=datapump DUMPFILE=NAEXP.dmp
LOGFILE=NAEXP.log
```

6. Examine the export log file, for example `NAEXP.log` in the `NA_data1` directory.

   - If the export was successful, continue to step 8 of this task.

   - If the export was not successful, resolve the problems causing the failure, and then repeat step 5 of this task.

7. If you did not run the user creation script to create the database user, identify this user's complete set of permissions. You will use this information to re-create the user in step 10 of this task.

8. With the NA management engine on NA1 shut down, log on to Master 1 as the SYS user, and then delete the NA database user while retaining the NA tablespace.

   For example, in the Oracle Enterprise Manager (preferred), delete `nauser`, or at the `sqlplus` prompt, enter the following command:

   ```
   DROP USER nauser CASCADE;
   ```

   > **Note:** Use the cascade delete option to completely delete the user and empty its tablespace.

9. Verify that the NA tablespace is now empty.

10. Log on to Master 1 as the SYS user, and then recreate the NA database user:

    - To use the user creation script, run commands similar to the following example:

      ```
      sqlplus / as sysdba
      run: SQL > @create_user.sql
      exit
      ```

    - Otherwise, recreate the user with the permissions collected in step 7 of this task.

    The newly-created NA database user connects to the existing NA tablespace. You might see an RMI error during the connection process. This error should clear after a few minutes, and the user-tablespace connection should occur without intervention.

11. On Master 1, outside of SQLPlus, use the Oracle Data Pump Utility to import the NA database schema for the NA database user. For example:

    ```
    impdp nauser/napass DIRECTORY=datapump TABLES=RN_% DUMPFILE=NAEXP.dmp
    LOGFILE=NAIMP1.log
    ```

    > **Note:** *Linux only*: Verify that the export file permissions are at least 755.

12. Examine the import log file, for example `NAIMP1.log` in the `NA_data1` directory.

    - If the import was successful, continue to Task 4: Prepare Master 2 for Replication (Oracle 11g).

    - If the import was not successful, do the following:
      - Resolve the problems causing the failure.

        Verify that the permissions for the user created in step 10 of this task match the user permissions collected in step 7 of this task.

- Drop the NA database user as described in step 8 of this task.
- Repeat step 11 of this task.

Task 4: Prepare Master 2 for Replication (Oracle 11g)

To import the Master 1 database onto Master 2, follow these steps:

1. Log on to Master 2 as the SYS user, and then create the NA tablespace with the same name and parameters as the NA tablespace on Master 1. For example:

   - *Windows*:

   ```
   sqlplus nauser/napass@SID2
   CREATE TABLESPACE NAORA DATAFILE 'C:\APP\ADMINISTRATOR\ORADATA\SID2\
   NAORA.DBF' SIZE 200M REUSE AUTOEXTEND ON NEXT 50M;
   exit
   ```

   - *Linux*:

   ```
   sqlplus nauser/napass@SID2
   CREATE TABLESPACE NAORA DATAFILE '/u01/app/oracle/oradata/SID2/NAORA.DBF' SIZE
   200M REUSE AUTOEXTEND ON NEXT 50M;
   exit
   ```

2. Log on to Master 2 as the SYS user, and then create the NA database user:

   - To use the user creation script, copy the script to Master 2 and then run commands similar to the following example:

   ```
   sqlplus / as sysdba
   run: SQL > @create_user.sql
   exit
   ```

   - Otherwise, create the user with the permissions collected in step 7.

   The newly-created NA database user connects to the existing NA tablespace. You might see an RMI error during the connection process. This error should clear after a few minutes, and the user-tablespace connection should occur without intervention.

3. On Master 2, create a directory to hold the tablespace export. For example, `C:\NA_data2` or `/tmp/NA_data2`.

   Ensure that the database user (for example, `nauser`) has read-write access to this directory.

4. Copy the tablespace export file (for example, `NAEXP.dmp`) to the new directory on Master 2.

   Ensure that the export file owner and group are those for the Oracle user for the operating system.

5. Log on to Master 2 as the NA database user, and then associate the new directory with the datapump directory in preparation for the tablespace import. For example:

   - *Windows*:

   ```
   sqlplus nauser/napass@SID2
   create DIRECTORY datapump AS 'C:\NA_data2';
   exit
   ```

- *Linux*:

```
sqlplus nauser/napass@SID2
create DIRECTORY datapump AS '/tmp/NA_data2';
exit
```

6. Log on to Master 2 as the SYS user, and then grant the NA database user read-write access to the datapump directory. For example:

```
sqlplus / as sysdba
GRANT READ,WRITE ON DIRECTORY datapump to nauser;
exit
```

7. On Master 2, verify that $ORACLE_SID is set to the SID for the NA database (for example, SID2).

8. On Master 2, outside of SQLPlus, use the Oracle Data Pump Utility to import the NA database schema for the NA database user. For example:

```
impdp nauser/napass DIRECTORY=datapump TABLES=RN_% DUMPFILE=NAEXP.dmp
LOGFILE=NAIMP2.log
```

> **Note:** *Linux only*: Verify that the export file permissions are at least 755.

9. Examine the import log file, for example NAIMP2.log in the NA_data2 directory.

   - If the import was successful, continue to Task 5: Configure Replication (Oracle 11g).

   - If the import was not successful, do the following:
     - Resolve the problems causing the failure.
     - Drop the NA database user as described in step 8 of Task 3: Prepare Master 1 for Replication (Oracle 11g)
     - Repeat step 8 of this task.

Task 5: Configure Replication (Oracle 11g)

To configure replication between the Master 1 and Master 2 databases, follow these steps:

1. On Master 1, add the Master 2 database information to the local TNSNAMES.ora. Verify that the Master 1 database information also exists in the TNSNAMES.ora. If the Master 1 database information does not exist, add it in the same format as the other entries.

   > **Note:** You might also need to make modifications to the LISTENER.ora and SQLNET.ora files depending on how the Oracle servers are configured. In particular, in certain environments the NAMES.DEFAULT_DOMAIN property needs to be commented out in SQLNET.ora. For more information, contact your Oracle DBA.

2. On Master 2, add the Master 1 database information to the local TNSNAMES.ora. Verify that the Master 2 database information also exists in the TNSNAMES.ora. If the Master 2 database information does not exist, add it in the same format as the other entries.

   > **Note:** You might also need to make modifications to the LISTENER.ora and SQLNET.ora files depending on how the Oracle servers are configured. In particular, in certain environments the

> NAMES.DEFAULT_DOMAIN property needs to be commented out in SQLNET.ora. For more information, contact your Oracle DBA.

3. On NA1, unpack the OracleReplicationScriptToolBundle.zip file to a known location, for example the NA_data directory used earlier in this procedure.

> **Tip:** On Windows systems, unpack the replication bundle to a location with low security, for example, C:\tmp.

4. From the known location of step 3, customize the ReplicationScriptTool.properties file for your environment:

    a. In a text editor such as WordPad or vi, open the ReplicationScriptTool.properties file.

    b. Search for and replace the text REPLACEME with the appropriate value for that variable.

        ○ Use the fully-qualified hostnames to identify the database and core servers. For example:

        ```
        db.server.0=Master1.example.com
        db.port.0=1521
        db.name.0=SID1
        db.username.0=nauer
        db.password.0=napass
        app.server.0=NA1.example.com
        app.rmiport.0=1099
        app.corename.0=Core1

        db.server.1=Master2.example.com
        db.port.1=1521
        db.name.1=SID2
        db.username.1=nauser
        db.password.1=napass
        app.server.1=NA2.example.com
        app.rmiport.1=1099
        app.corename.1=Core 2
        ```

    c. Set timezone_offset to the value that matches the time zone setting for all servers in the distributed system.

    d. Verify that mode is set to initial.

5. From the known location of step 3, run java -version to verify that Java is version 1.7 or later.

6. From the known location of step 3, run the ReplicationScriptTool:

    ```
    java -classpath . ReplicationScriptTool
    ```

    This command creates a script for starting replication. By default, the script is named ReplicationTool.sql.

    > **Tip:** The following error message can mean that Java does not have permission to run the replication tool:
    >
    > ```
    > Could not find the main class: ReplicationScriptTool.
    > ```
    >
    > Move all files from the replication bundle to a directory with lower security, and then run the java command from that directory.

7. On Master 1, ensure that SQLPlus that can access all NA Cores in the NA Mesh. To verify this, connect to one database with SQLPlus and issue the `CONNECT <USER>/<PASSWORD>@<SID>` command for each different SID in the Mesh. This is necessary because the script issues CONNECT commands to change which database it is running procedures against.

8. Copy the replication script created in step 6 (for example, `ReplicationScript.sql`) from the known location to the server identified in step 7.

   - *Windows*: Place the file in `C:\`.

   - *Linux*: Place the file in the `$ORACLE_HOME/bin` directory, for example `/u01/app/oracle/product/11.2.0/dbhome_1/bin`.

   > **Note:** Make sure that files named `*_spool.txt` do not exist in the directory. When run, the script will log its actions to a file with that suffix.

9. On Master 1, open two SQLPLus windows, each connected as the NA database user (for example, `nauser`).

10. On Master 2, open a SQLPlus window connected as the NA database user (for example, `nauser`).

11. In one SQLPlus window on Master 1, run the replication script. For example:

    `@ReplicationScript.sql`

12. Follow the instructions provided by the script at several points as it runs. This process involves running commands in the second SQLPlus window on Master 1 to verify that certain operations necessary for later steps in the script have completed.

    At the first and second pauses, verify the configuration for both of the databases by running the same command in each of the following windows:

    - The second SQLPlus window on Master 1.

    - The SQLPlus window on Master 2.

    At the third pause, run the command in the second SQLPlus window on Master 1 only. When the `SELECT STATUS,REQUEST,ONAME FROM DBA_REPCATLOG;` statement returns "0 rows selected," replication setup is almost completed. Finish the process in the NA database user window.

    > **Note:** If you see errors during replication setup, before performing initial replication setup again, see "Removing Replication Across the Entire NA Mesh" on page 31.

    > **Tip:** Depending on database size, it can take 30 to 60 minutes to see the "0 rows selected" results.

13. After the replication script completes, close both SQLPlus windows to ensure that all actions are logged to the spool file.

14. Verify that the replication is working correctly. Make sure:

    - The spool file generated by running the replication SQL script contains no errors.

    - The RN_CORE table in each database contains the appropriate list of servers in the NA Mesh. To check the RN_CORE table, using SQLPlus, enter the following command on each server:

```
SELECT * FROM RN_CORE;
```

- Verify (using the Oracle Enterprise Manager) that each database in the NA Mesh has replication set up, is in the `Normal` state, and has purge, push, and repcat_admin jobs installed correctly. Contact your Oracle DBA for assistance.

- Verify that the replication of data is working correctly. To verify that replication is working between all NA Cores in the NA Mesh, enter the following commands:

```
CONNECT <USER>/<PASSWORD>@<MASTER1_SID>;
UPDATE RN_CORE SET COMMENTS = 'Test from sid1 to sid2' where CoreID=1;
COMMIT;
```

wait 10 to 30 seconds

```
CONNECT <USER>/<PASSWORD>@<MASTER2_SID>;
SELECT Comments from RN_CORE where CoreID=1;
```

Enter this from every NA Core in the NA Mesh. Check the values each time on all of the other NA Cores to ensure that the links between each database in the system are working correctly.

- The following command must return 'no rows selected' and not 'ready':

```
SELECT STATUS,REQUEST,ONAME FROM DBA_REPCATLOG;
```

15. Install NA on the second core server (NA2).

    During installation, connect NA to the database on Master 2. When prompted for the database credentials, choose: use existing database.

Task 6: Configure the NA Core Servers for Replication

1. Stop all NA services on NA2. See "Start, Stop, or Restart All NA Services" on page 57.

2. Copy the `distributed.rcx` file from the known location of step 3 of Task 5: Configure Replication (Oracle 11g) to the following directory on both core servers, NA1 and NA2:

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

   - *Linux*: `<NA_HOME>/jre`

3. Start all NA services on both NA core servers, NA1 and NA2. See "Start, Stop, or Restart All NA Services" on page 57.

4. Connect to the NA console on both core servers, NA1 and NA2, (as user `nauser`), and ensure that there are no issues.

   > **Tip:** An RMI error message at this point can mean that NA is unable to identify the localhost. For more information, see "RMI Errors in the NA Console" on page 53.

5. Verify the NA installation:

   a. Edit an object on one NA server (for example, a Comments field for a device).

   b. Wait a minute and then verify that the updated comment exists on the second server.

c. Check the status of the Oracle Distributed Monitor in the NA console to ensure that no problems are being reported. (For information, see "Distributed Monitor Results Page" on page 45.) It could take up to five minutes for this monitor to initially run.

NA Database User Creation Script (Oracle 11g)

```
/*Create a new user with a default tablespace*/
CREATE USER nauser IDENTIFIED BY napass DEFAULT TABLESPACE NAORA;

/*grant Role equivalent to a system user (only choose public if you choose to provide
basic privileges)*/
GRANT DBA to nauser WITH ADMIN OPTION;

/*grant permissions for normal NA operation*/
GRANT CREATE SEQUENCE, CREATE SESSION, CREATE TABLE, CREATE PROCEDURE, SELECT ANY
DICTIONARY, CONNECT, CREATE VIEW TO nauser;
GRANT EXECUTE ON SYS.DBMS_REPCAT_ADMIN TO nauser;
GRANT EXECUTE ON SYS.DBMS_DEFER_SYS TO nauser;
GRANT CREATE PUBLIC DATABASE LINK TO nauser;
GRANT CREATE DATABASE LINK TO nauser;
GRANT CREATE USER TO nauser;
GRANT CREATE ANY PROCEDURE TO nauser;

/*For normal operation of a NA Core in a replication environment:*/
GRANT EXECUTE ON SYS.DBMS_REPUTIL TO nauser;
GRANT SELECT ON SYS.DBA_CONSTRAINTS TO nauser;
GRANT SELECT ON SYS.DBA_CONS_COLUMNS TO nauser;
GRANT SELECT ON SYS.DBA_REPCOLUMN TO nauser;

/*For replication deletion:*/
GRANT DROP USER TO nauser;
GRANT DROP PUBLIC DATABASE LINK TO nauser;
GRANT DROP ANY PROCEDURE TO nauser;

/*grant permissions for system level*/
GRANT CREATE MATERIALIZED VIEW, GLOBAL QUERY REWRITE, SELECT ANY TABLE to nauser;
GRANT UNLIMITED TABLESPACE to nauser WITH ADMIN OPTION;
commit;
```

Oracle 11g: Adding a New NA Core to an Existing NA Mesh

Master 1 is the existing Master Definition Site and Master 3 is the new NA Core being added.

To add an NA core to an existing NA mesh, follow these steps:

1. System Setup:

   - A current NA 10.20 database on Master 1.

   - NA core server of NA Core 1 connected to the database on Master 1.

   - A database with no data or NA schema setup on Master 3. The database identifier (SID) needs to be different from the existing databases in the NA Mesh.

   - A network connection from the Master 1 servers to the Master 3 servers (and vice-versa) that enables ports 1521 (Oracle only), 1098, 8083, 1099, and 4444 (or appropriate variations) to be connected between these servers. Port 1521 must be open between the Oracle servers. The remaining ports must be open between the NA servers.

   - The *ReplicationScriptTool bundle* files installed on a Java-capable system.

   - All database servers and core servers are set to use the same time and time zone.

2. Obtain the following information:

   - Login/password/database identifier for NA into the two databases (on Master 1 and Master 3).

   - The replication password for the existing repadmin user.

   - NA server hostname, NA server RMI listening port, database hostname, and database listening port for Master 3.

   - The time zone offset (integer from UTC) for the entire NA Mesh. This must be a constant across the NA Mesh.

   - The existing replication group name.

3. Set (or verify) the following database parameters for Master 3:

   - global_names = TRUE

   - open_links = at least 5 (increase this by 1 for each additional server in the NA Mesh beyond 2)

   - open_links_per_instance = at least 5 (same as open_links)

   - parallel_min_servers = 2

   - shared_pool_size = increase current value by 40m to support replication

   - job_queue_processes = at least 2.

   Note: The parameters must be set on each database in the NA Mesh. The parameters could be set up when the database is first created. If not, they must be modified at this point. You can view the current value of these parameters using Oracle Enterprise Manager or SQLPlus with the "SHOW PARAMETERS" or "SELECT NAME, VALUE FROM v$parameter" commands (where NAME = <parameter_name>, for example, global_names).

4. Stop all NA services on all NA core servers. See "Start, Stop, or Restart All NA Services" on page 57.

5. As the **repadmin** user, enter the following command using SQLPlus on Master 1 to quiesce the replication group:

   ```
   EXEC DBMS_REPCAT.SUSPEND_MASTER_ACTIVITY('<replication_group_name>');
   ```

   where <replication_group_name> is from Step 2.

6. Wait for the replication group to quiesce. Do this by executing the `SELECT GNAME,STATUS FROM DBA_REPGROUP;` command to check on the status. Wait until the status is listed as QUIESCED.

7. Create the Master 1 database export and finish preparing the database on Master 1 as described in "Task 3: Prepare Master 1 for Replication (Oracle 11g)" on page 19.

8. Import the Master 1 database onto Master 3 as described in "Task 4: Prepare Master 2 for Replication (Oracle 11g)" on page 21.

   Substitute "Master 3" for each reference to "Master 2."

9. Configure replication between the Master 1 and Master 3 databases as described in "Task 5: Configure Replication (Oracle 11g)" on page 22 and "Task 6: Configure the NA Core Servers for Replication" on page 25.

   Note the following differences:

   - Substitute "Master 3" for each reference to "Master 2."

   - For step 4, if you still have the `ReplicationScriptTool.properties` file from setting up initial replication you can use it. Otherwise, use the file from the most recent unpacking of the `OracleReplicationScriptToolBundle.zip` file.

     ○ Verify or edit the server information for the initial core and second master site.

     ○ Uncomment the Additional Master site configuration area, and replace the text REPLACEME with the appropriate value for that variable. Use the fully-qualified hostnames to identify the database and core servers. For example:

       ```
       db.server.1=Master3.example.com
       db.port.1=1521
       db.name.1=SID3
       db.username.1=nauser
       db.password.1=napass
       app.server.1=NA3.example.com
       app.rmiport.1=1099
       app.corename.1=Core 3
       ```

     ○ Set the mode to add_server.

     ○ For step 13, the spool file contains the following acceptable warning that occurs after the DBMS_REPCAT.SUSPEND_MASTER_ACTIVITY command:

       ```
       'ORA-23326: object group <yourgroup>.<your_replication_ group> is quiesced'
       ```

   - Carry out step 3 and step 4 on all NA core servers in the NA mesh.

Removing an NA Core from an Existing NA Mesh

> **Caution:** After you remove an NA Core from an NA Mesh, the data in that NA Core's database will no longer remain in sync with the data in the rest of the NA Mesh. The data cannot be merged back into the NA Mesh.

To remove an NA Core from an NA Mesh:

1. If the NA core you want to remove is the Master Definition Site, change the Master Definition Site to be another NA Core as described in .

2. Install the *ReplicationScriptTool bundle* files on a Java-capable system.

3. Ensure that all NA devices belong to Sites on NA Cores that are not going to be removed.

4. Modify all Sites to point to an NA Core that is not going to be removed. (Alternatively, you can remove those Sites.)

5. Turn off the NA server from the NA Core that is being removed.

6. Delete the RN_CORE entry that was removed using the following commands using SQLPlus:

```
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;
```

> **Tip:** To identify the Core ID of the NA Core to delete, enter the `SELECT*FROM RN_CORE` command.

7. Update the variables for all masters in the *ReplicationScriptTool.properties* file. Ensure that you modify all "REPLACEME" text entries in the *ReplicationScriptTool.properties* file with the appropriate values for your environment.

> **Note:** The master being deleted must be the last master entry and the Master Definition Site must be the first entry. In addition, update the replication user, password, and replication group name in the *ReplicationScriptTool.properties* file if needed. Make sure the "mode" property is set to "delete_server".

8. Run `java -version` to verify that Java is version 1.7 or later.

9. Run the ReplicationScriptTool via the following command:

```
java -classpath . ReplicationScriptTool
```

10. Ensure that you have a server and directory (typically on one of the database systems) where you can run a copy of SQLPlus that can access all NA Cores in the NA Mesh. To verify this, connect to one database with SQLPlus and issue the `CONNECT <USER>/<PASSWORD>@<SID>` command for each different SID in the NA Mesh. This is necessary because the script issues CONNECT commands to change which database it is running procedures against.

11. Copy the newly generated SQL file (typically named *ReplicationScript.sql* depending on what value was set in the *ReplicationScriptTool.properties* file), from the ReplicationScriptTool working directory to the

server and directory where you plan to run SQLPlus.

> **Note:** Make sure that files named `*_spool.txt` do not exist in that directory. When run, the script will log its actions to a file of that suffix.

12. Ensure that you have two SQLPLus windows open before running the script: One connected as the NA user and one connected as the SYSTEM user.

13. Run the replication script copied in the previous step in the SQLPlus window connected as the NA user using `@ReplicationScript.sql`.

14. Follow the instructions provided by the script at several points as it runs. This will involve running other commands in the second SQLPlus window to verify that certain operations necessary for later steps in the script have completed.

15. Exit from SQLPlus once the script completes. This ensures that all actions are logged to the spool file.

16. Verify that the replication is working correctly. Make sure:

    - The spool file generated by running the replication SQL script contains no errors.

    - The RN_CORE table on each database contains the appropriate list of servers in the NA Mesh. To check the RN_CORE table, using SQLPlus, enter the following command on each server:

      "`SELECT * FROM RN_CORE;`"

    - Verify (using Oracle's Enterprise Manager) that each database in the NA Mesh has replication set up, is in the "Normal" state, and has purge, push, and repcat_admin jobs installed correctly. Consult your Oracle DBA for assistance.

    - Verify that the replication of data is working correctly. To verify that replication is working between all NA Cores in the NA Mesh, enter the following commands:

      ```
      CONNECT <USER>/<PASSWORD>@<MASTER1_SID>;
      UPDATE RN_CORE SET COMMENTS = 'Test from sid1 to sid2' where CoreID=1;
      COMMIT;
      ```

      wait 10 to 30 seconds

      ```
      CONNECT <USER>/<PASSWORD>@<MASTER2_SID>;
      SELECT Comments from RN_CORE where CoreID=1;
      ```

      You should enter this from every NA Core in the NA Mesh. Be sure to check the values each time on all of the other NA Cores to ensure that the links between each database in the system are working correctly.

17. Uninstall the NA server from the NA Core that was removed.

18. Restart the NA servers in the NA Mesh.

Changing the Master Definition Site

It may be necessary to change the Master Definition Site in the event of failure of that Site or when you want to remove an NA Core from the NA Mesh. To change the Master Definition Site:

1. If the old Master Definition Site is accessible (for example in the case where it is being removed from the NA Mesh), as the "repadmin" user on the existing Master Definition Site, enter the following command using SQLPlus:

```
BEGIN
    DBMS_REPCAT.RELOCATE_MASTERDEF (
        gname => '<REPLICATION GROUP NAME>',
          old_masterdef => '<OLD MASTERDEF SID>',
          new_masterdef => '<NEW MASTERDEF SID>',
          notify_masters => TRUE,
 include_old_masterdef => TRUE);
END;
/
COMMIT;
```

2. If the old Master Definition Site is not accessible (for example in the case of failure of the Master Definition database server), as the "repadmin" user on the existing Master Definition Site, enter the following command using SQLPlus on any other NA Core in the NA Mesh:

```
BEGIN
    DBMS_REPCAT.RELOCATE_MASTERDEF (
        gname => '<REPLICATION GROUP NAME>',
          old_masterdef => '<OLD MASTERDEF SID>',
          new_masterdef => '<NEW MASTERDEF SID>',
          notify_masters => TRUE,
 include_old_masterdef => FALSE);
END
/
COMMIT;
```

3. For both cases, as the NA database login user, modify the RN_CORE table entries:

```
UPDATE RN_CORE SET IsMasterDef = 0 WHERE CoreID = <ID OF OLD MASTERDEF>;
COMMIT;

UPDATE RN_CORE SET IsMasterDef = 1 WHERE CoreID = <ID OF NEW MASTERDEF>;
COMMIT;
```

Removing Replication Across the Entire NA Mesh

Caution: All NA servers participating in the NA Mesh will now operate as independent NA installations. The data cannot be merged once they are no longer in sync.

To remove replication across the entire NA Mesh:

1. Make sure that the *ReplicationScriptTool bundle* files are installed on a Java-capable machine.

2. Turn off the NA server on all NA Cores.

3. Update the variables for all masters in the *ReplicationScriptTool.properties* file. Ensure that you modify all "REPLACEME" text entries in the *ReplicationScriptTool.properties* file with the appropriate values for your environment.

> **Note:** The Master Definition Site must be the first entry. In addition, update the replication user, password, and replication group name in that file if needed. Make sure the "mode" property is set to "delete".

4. Run `java -version` to verify that Java is version 1.7 or later.

5. Run the ReplicationScriptTool via the following command:

   `java -classpath . ReplicationScriptTool`

6. Ensure that you have a server and directory (typically on one of the database systems) where you can run a copy of SQLPlus that can access all NA Cores in the NA Mesh. To verify this, connect to one database with SQLPlus and issue the `CONNECT <USER>/<PASSWORD>@<SID>` command for each different SID in the NA Mesh. This is necessary because the script issues CONNECT commands to change which database it is running procedures against.

7. Copy the newly generated SQL file (typically named *ReplicationScript.sql* depending on what value was set in the *ReplicationScriptTool.properties* file), from the ReplicationScriptTool working directory to the server and directory where you plan to run SQLPlus.

> **Note:** Make sure that files named "*_spool.txt" do not exist in that directory. When run, the script will log its actions to a file of that suffix.

8. Ensure that you have two SQLPLus windows open before running the script: One connected as the NA user and one connected as the SYSTEM user.

9. Run the replication script copied in the previous step in the SQLPlus window connected as the NA user using `@ReplicationScript.sql`.

10. Follow the instructions provided by the script at several points as it runs. This will involve running other commands in the second SQLPlus window to verify that certain operations necessary for later steps in the script have completed.

11. Exit SQLPlus after the script completes (this ensures that all actions are logged to the spool file).

12. Verify that the replication has been removed correctly. Make sure:

    - The spool file generated by running the replication SQL script contains no errors.

    - Verify (using Oracle's Enterprise Manager) that each database in the NA Mesh does not have replication setup. Consult your Oracle DBA for assistance.

13. For each server, ensure that devices are deleted or associated with Sites that map to the local NA Core as needed.

14. Remove the `distributed.rcx` file from each NA core server.

15. Restart the NA servers.

16. For each database, remove all entries from RN_CORE except for the local NA Core system's entry.

17. Restart each NA server to enable the NA Core changes to take effect.

Upgrade the NA Mesh

During the NA upgrade process, the NA cores must be completely offline. You will be upgrading the:

- NA core servers
- NA database schema
- Replication among the database servers

Note the following:

- The NA mesh upgrade path is as follows:

  - NA 10.1x to NA 10.20

  Enter this path at your current version of NA.

  > **Note: Note:** For upgrading from NA 10.0x, the mesh upgrade path is as follows:
  >
  > NA 10.0x to NA 10.10 to NA10.20

- At each step on the NA mesh upgrade path, upgrade all NA cores in the NA mesh before moving to the next step in the upgrade path.

For upgrade instructions, see the following topics:

- "Upgrading from 10.0x to NA 10.10" below
- "Upgrading from NA 10.1x to NA 10.20" below

Upgrading from 10.0x to NA 10.10

You must explicitly upgrade all NA core servers in the NA Multimaster Distributed System environment to NA 10.10 before upgrading to NA 10.20. For more information about upgrading from NA 10.0x to NA 10.10, and see the *NA Multimaster Distributed System on Oracle Guide (10.10)*.

Upgrading from NA 10.1x to NA 10.20

Before you upgrade, be sure that:

- A current NA 10.00 or NA 10.10 database is installed on all Masters.
- The *ReplicationScriptTool* bundle files for NA 10.20 are installed on a Java-capable system (on one of the NA core servers).
- You have the following information:

  - Login/password/DatabaseIdentifier (SID) for NA into all databases in the NA mesh

  - Replication password for the existing repadmin user

  - DatabaseIdentifier (SID), NA core server hostname, NA server RMI listening port, database server hostname, and database listening port for all Masters

  - The time zone offset (integer from UTC) for the entire NA mesh (This must be a constant across the NA mesh.)

  - The existing replication group name

To upgrade from NA 10.00 or NA 10.10 to NA10.20, follow these steps:

1. Stop all NA services on all NA core servers. See "Start, Stop, or Restart All NA Services" on page 57.

2. Update the variables for all Masters in the *ReplicationScriptTool.properties* file. Ensure that you modify all "REPLACEME" text entries in the *ReplicationScriptTool.properties* file with the appropriate values for your environment.

   > **Note:** All servers should be listed in the *ReplicationScriptTool.properties* file. Set the `mode` property to **upgrade_from_10_0**.

3. Run `java -version` to verify that Java is version 1.7 or later.

4. Run the ReplicationScriptTool via the following command:

   `java -classpath . ReplicationScriptTool`

5. Ensure that you have a server and a directory (typically on one of the database systems), where you can run a copy of SQLPlus that can access all NA cores in the NA mesh. To verify this, connect to one database with SQLPlus and issue the CONNECT <USER>/<PASSWORD>@<SID> command for each unique SID in the NA mesh. This is necessary because the script issues CONNECT commands to change which database it is running procedures against.

6. Copy the newly generated SQL file (typically named *ReplicationScript.sql,* depending on what value was set to in the *ReplicationScriptTool.properties* file), from the ReplicationScriptTool working directory to the server and directory where you plan to run SQLPlus.

   > **Note:** Make sure that files named `*_spool.txt` do not exist in the directory. When run, the script will log its actions to a file with that suffix.

7. Ensure that you have two SQLPLus windows open before running the script:

   - In one window, connect as the NA database user (for example, `nauser`).

   - In the other window, connect as the SYS user.

8. Run the replication script copied in the previous step in the SQLPlus window connected as the NA user using @ReplicationScript.sql.

9. Follow the instructions provided by the script at several points as it runs. This will include running other commands in the second SQLPlus window to verify that certain operations necessary for later steps in the script have completed.

   > **Note:** If you skipped or misconfigured any of the above steps, and as a result see errors during replication setup, before performing the upgrade, see "Removing Replication Across the Entire NA Mesh" on page 31.

10. After the script completes, exit from SQLPlus to ensure that all actions are logged to the spool file.

11. Verify that the replication is working correctly. Make sure:

    - The spool file generated by running the replication SQL script contains no errors, other than the following warning: 'ORA-23326: object group <yourgroup>.<your_replication_ group> is quiesced' that occurs after the DBMS_REPCAT.SUSPEND_MASTER_ACTIVITY command.

- Verify (using Oracle's Enterprise Manager) that each database in the NA Mesh has replication set up and is in the "Normal" state, and has purge, push, and repcat_admin jobs installed correctly. Consult your Oracle DBA for assistance.

- Verify that the replication of data is working correctly. To verify that replication is working between all NA Cores in the NA Mesh, enter the following commands:

  ```
  CONNECT <USER>/<PASSWORD>@<MASTER1_SID>;
  UPDATE RN_CORE SET COMMENTS = 'Test from sid1 to sid2' where CoreID=1;
  COMMIT;
  ```

  wait 10 to 30 seconds

  ```
  CONNECT <USER>/<PASSWORD>@<MASTER2_SID>;
  SELECT Comments from RN_CORE where CoreID=1;
  ```

  You should enter this from every NA Core in the NA Mesh. Be sure to check the values each time on all of the other NA Cores to ensure that the links between each database in the system are working correctly.

12. Upgrade each NA core server as described in the *Upgrading to NA 10.20 from a Different System* chapter or the *Upgrading to NA 10.20 on the Same System* chapter in version *10.20* of the *NA Installation and Upgrade Guide*, as appropriate.

    Run the upgrade procedures in parallel. That is, complete step 1 on each NA core server before initiating step 2 on any NA core server, and so forth. Follow the referenced procedure through the step to run the NA10.20 Service Pack Installer.

    Note the following:

    - To prevent database access, on each NA core server, after the NA 10.20 Service Pack Installer runs, stop all NA services. See "Start, Stop, or Restart All NA Services" on page 57.

    - The NA database schema upgrade has already been handled by running the replication script in step 7.

    - Do *not* upgrade Oracle while upgrading NA. If you need to upgrade Oracle, complete this procedure, and then upgrade Oracle as described in "Restoring NA Replication After Upgrading to Oracle 11g" on page 60.

13. If you moved NA to different core servers in the previous step, update the RN_CORE table on one NA core with the fully-qualified domain names of the new NA core servers. For example:

    ```
    CONNECT <USER>/<PASSWORD>@<MASTER1_SID>;
    update RN_CORE set COREHOSTNAME = 'NA1.example.com' where CoreID = 1;
    update RN_CORE set COREHOSTNAME = 'NA2.example.com' where CoreID = 2;
    ```

14. Start all the NA services on the master core server and make sure that you can login to NA on the master core. See "Start, Stop, or Restart All NA Services" on page 57.

15. Start the NA services on all the other new cores.

16. Complete the upgrade procedure from the step to run the Data Pruning task.

17. Temporarily enable diagnostic check points to populate the Most Recent Diagnostic column in the NA database with accurate data.

    a. Log on to the NA console as an NA administrator.

    b. On the Configuration Mgmt page (**Admin > Administrative Settings > Configuration Mgmt**) enable the **Force a save of the latest diagnostics.** option under Diagnostics.

    c. Create a new Run Diagnostics task with the following configuration:

- Set **Applies to** to the entire device inventory.
- In the **Diagnostics to Run** field, select all diagnostics.

    d. After the Run Diagnostics task completes, disable the **Force a save of the latest diagnostics** option on the Configuration Mgmt page.

Complete this step on only one NA core. This step is necessary in Multimaster Distributed System environments only.

18. Verify the NA installation:

    a. Edit an object on one NA core (for example, a Comments field for a device).

    b. Wait a minute and then verify that the updated comment exists on the second server.

    c. Check the status of the Oracle Distributed Monitor in the NA console to ensure that no problems are being reported. (For information, see "Distributed Error List" on page 46.)

Chapter 3: System Administration

This chapter contains the following topics:

-
-
-
-
-

Getting Started

In general, an NA server that is part of a Distributed System NA Mesh should be transparent to users. However, there are a number of operations that the system administrator may need to do to keep the Distributed System NA Mesh functioning properly.

Understanding Oracle Replication

Oracle replication uses three types of jobs:

- **Transaction Push Job** — Each NA Core will have a Transaction Push job for each of the other NA Cores in the NA Mesh. For example, if you have two NA Cores, each NA Core has one Transaction Push job. If you have three NA Cores, each NA Core has two Transaction Push jobs for the two destination NA Cores. These jobs are identified in Oracle's Enterprise Manager or from DBA_JOBS by the `DBMS_DEFER_SYS.PUSH()` command string.
- **Transaction Purge Job** — Each NA Core will have one Transaction Purge job for cleaning out successfully pushed transactions. These jobs can be identified in Oracle's Enterprise Manager or from DBA_JOBS by the `DBMS_DEFER_SYS.PURGE ()` command string.
- **Replication Administration Job** — Each NA Core will have one Replication Administration job. These jobs are not used in normal NA operation. It is used if the replication setup is changed. These jobs can be identified in Oracle's Enterprise Manager or from DBA_JOBS by the `DO_DEFERRED_REPCAT_ADMIN()` command string.

NA automatically monitors all three jobs and attempts to restart any job that fails. You can monitor the jobs using Oracle's Enterprise Manager or by viewing DBA_JOBS. Using SQLPlus, enter: `SELECT * FROM DBA_JOBS;`

To view specific job IDs and if they are broken, enter: `SELECT JOB, BROKEN, WHAT FROM DBA_JOBS;`

On Oracle, NA monitors on a regular schedule the following on each NA Core for the database transport (i.e., replication):

- Push jobs that handle moving data from one NA Core to another. If the monitor check fails and the Push job is found to be broken, NA will attempt to restart it.
- Purge jobs that successfully clear pushed data from the push queue, and attempts to restart the Push queue each time the check fails.

- The deferred transaction queue for any transactions that could not be applied to the local database due to a conflict. NA applies additional application logic to attempt to correctly apply the transaction. If unsuccessful, NA alerts the NA administrator via email notification.

- Unrecoverable uniqueness conflicts. NA alerts the NA administrator in the event of unrecoverable uniqueness conflicts. Typically, there is no way to automatically correct these types of issues.

- Both the DEFTRAN and DEFLOB queue size. These queues hold pending data that needs to be pushed to other servers. Typically, if replication is working, but is beginning to backup, these queues will begin to grow. This can occur if the connection between databases is congested. NA alerts the NA administrator to this issue via email notification. The system will continue to push data. However, it could take a long time to catch up depending on network conditions.

> **Note:** For information about configuring email notification, see the *NA User Guide*.

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

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

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

> **Note:** NA includes an example event rule that can be updated as necessary.

How NA Removes In-Memory and Database Information

There could be a delay in removing or cancelling a task because NA must communicate with the appropriate owning NA Core to perform this operation. To ensure proper removal of all in-memory and database information, and to avoid replication conflicts due to matching timestamps, NA active tasks are automatically deleted on the NA Core with which they are associated. An active task is any task that does not have the "Succeeded," "Failed," "Duplicate", "Skipped", or "Warning" status.

Keep in mind that deleted tasks could be displayed in task lists for a few moments while the replication process pushes the deletes to other NA Cores in the system. In addition, if the NA Core from which the task originated is not accessible, the delete will fail.

Regenerating Replication

After enabling or disabling case-insensitive search, regenerate the database replication to ensure that the index information is synchronized.

To regenerate database replication

1. On NA1, unpack the `OracleReplicationScriptToolBundle.zip` file to a known location, for example the `NA_data` directory used earlier in this procedure.

> **Tip:** On Windows systems, unpack the replication bundle to a location with low security, for example, `C:\tmp`.

2. From the known location of step 1, customize the `ReplicationScriptTool.properties` file for your environment:

   a. In a text editor such as WordPad or vi, open the `ReplicationScriptTool.properties` file.

   b. Search for and replace the text `REPLACEME` with the appropriate value for that variable.

      ○ Use the fully-qualified hostnames to identify the database and core servers. For example:

      ```
      db.server.0=Master1.example.com
      db.port.0=1521
      db.name.0=SID1
      db.username.0=nauer
      db.password.0=napass
      app.server.0=NA1.example.com
      app.rmiport.0=1099
      app.corename.0=Core1

      db.server.1=Master2.example.com
      db.port.1=1521
      db.name.1=SID2
      db.username.1=nauser
      db.password.1=napass
      app.server.1=NA2.example.com
      app.rmiport.1=1099
      app.corename.1=Core 2
      ```

   c. Set `timezone_offset` to the value that matches the time zone setting for all servers in the distributed system.

   d. Set `mode` to `regenerate_replication`.

3. From the known location of step 1, run `java -version` to verify that Java is version 1.7 or later.

4. From the known location of step 1, run the ReplicationScriptTool:

   `java -classpath . ReplicationScriptTool`

   This command creates a script for starting replication. By default, the script is named `ReplicationTool.sql`.

   > **Tip:** The following error message can mean that Java does not have permission to run the replication tool:
   >
   > `Could not find the main class: ReplicationScriptTool.`
   >
   > Move all files from the replication bundle to a directory with lower security, and then run the `java` command from that directory.

5. On Master 1, ensure that SQLPlus that can access all NA Cores in the NA Mesh. To verify this, connect to one database with SQLPlus and issue the `CONNECT <USER>/<PASSWORD>@<SID>` command for each different SID in the Mesh. This is necessary because the script issues CONNECT commands to change which database it is running procedures against.

6. Copy the replication script created in step 4 (for example, `ReplicationScript.sql`) from the known

location to the server identified in step 5.

- *Windows*: Place the file in C:\.

- *Linux*: Place the file in the $ORACLE_HOME/bin directory, for example
  /u01/app/oracle/product/11.2.0/dbhome_1/bin.

> **Note:** Make sure that files named *_spool.txt do not exist in the directory. When run, the script will log its actions to a file with that suffix.

7. On Master 1, open two SQLPLus windows, each connected as the NA database user (for example, nauser).

8. On Master 2, open a SQLPlus window connected as the NA database user (for example, nauser).

9. In one SQLPlus window on Master 1, run the replication script. For example:

   @ReplicationScript.sql

10. Follow the instructions provided by the script at several points as it runs. This process involves running commands in the second SQLPlus window on Master 1 to verify that certain operations necessary for later steps in the script have completed.

    At the first and second pauses, verify the configuration for both of the databases by running the same command in each of the following windows:

    - The second SQLPlus window on Master 1.

    - The SQLPlus window on Master 2.

    At the third pause, run the command in the second SQLPlus window on Master 1 only. When the SELECT STATUS,REQUEST,ONAME FROM DBA_REPCATLOG; statement returns "0 rows selected," replication setup is almost completed. Finish the process in the NA database user window.

    > **Note:** If you see errors during replication setup, before performing initial replication setup again, see "Removing Replication Across the Entire NA Mesh" on page 31.

    > **Tip:** Depending on database size, it can take 30 to 60 minutes to see the "0 rows selected" results.

11. After the replication script completes, close both SQLPlus windows to ensure that all actions are logged to the spool file.

12. Verify that the replication is working correctly. Make sure:

    - The spool file generated by running the replication SQL script contains no errors.

    - The RN_CORE table in each database contains the appropriate list of servers in the NA Mesh. To check the RN_CORE table, using SQLPlus, enter the following command on each server:

      SELECT * FROM RN_CORE;

    - Verify (using the Oracle Enterprise Manager) that each database in the NA Mesh has replication set up, is in the Normal state, and has purge, push, and repcat_admin jobs installed correctly. Contact your Oracle DBA for assistance.

    - Verify that the replication of data is working correctly. To verify that replication is working between all

NA Cores in the NA Mesh, enter the following commands:

```
CONNECT <USER>/<PASSWORD>@<MASTER1_SID>;
UPDATE RN_CORE SET COMMENTS = 'Test from sid1 to sid2' where CoreID=1;
COMMIT;
```

wait 10 to 30 seconds

```
CONNECT <USER>/<PASSWORD>@<MASTER2_SID>;
SELECT Comments from RN_CORE where CoreID=1;
```

Enter this from every NA Core in the NA Mesh. Check the values each time on all of the other NA Cores to ensure that the links between each database in the system are working correctly.

- The following command must return 'no rows selected' and not 'ready':

```
SELECT STATUS,REQUEST,ONAME FROM DBA_REPCATLOG;
```

Suspending Replication for Database Maintenance

To suspend replication because you are bringing down one or both database servers for maintenance, complete the following procedure as user repadmin from the masterdef database site:

1. Stop all NA services all NA core servers. See .
2. `CONNECT repadmin/<PWD>@<MASTER1_SID>;`
3. `EXEC DBMS_REPCAT.SUSPEND_MASTER_ACTIVITY(<REPGROUP>);`

NA Generated Events

By default, NA generates system events. Event rules can alert you to certain error conditions requiring attention. Each event is listed below, along with an explanation and required action to be taken. You should examine the default "Distributed System" event rule to ensure all of the events are included in the event rule and that the event rule is configured to send the email notification to the appropriate administrator.

Note: For information about configuring email notification, see the *NA User Guide*.

Broken Replication Job

Event format:

```
Job ID: <ID>
 Job Type: <push to <host>, purge, or repcat admin>
 Database Host: <hostname>
 Database SID: <SID>
 Stalled: <true or false>
```

Broken jobs can occur because an Oracle server cannot push transactions to another server due to that server being down or in the event of a network connectivity loss. Replication jobs can also become broken

because they stall. Stalled replication jobs represent a hang in an internal Oracle process. This event is sent only once for each broken replication job (unless the NA server is restarted, in which case the event will be resent if the job is still broken). If a stalled replication job is detected, it is best to restart the Oracle server.

For broken replication jobs that are not stalled, the course of action depends on the job type. If the job type is "purge" or "repcat admin," consult with your Oracle DBA. Failure of these jobs represents an internal problem or resource limit in the Oracle server that will need to be investigated.

For "push to <host>" job types, the typical failure will be loss of connectivity to the other host, either through a network problem, outage of the Oracle server, or outage of the database host system itself. NA will restart this job automatically once the connectivity problem has been remedied.

To troubleshoot this problem, make sure:

1. The host, to which the server cannot connect, is up and running.
2. The Oracle instance on that host is running.
3. From a command line enter `ping <host>` to ensure that network connectivity exists between servers.
4. From a command line enter `tnsping <OTHER_SID>` to ensure that Oracle is accepting connections.
5. Ensure that database links are working:

    - Using SQLPlus, enter `SELECT CURRENT_TIMESTAMP FROM DUAL@<OTHER_SID>;`

    - Using SQLPlus, enter `SELECT DBLINK FROM DBA_REPSITES;`

Failures of any of these steps will point to corrective actions needed.

Fixed Replication Job

Event format:

```
Job ID: <ID>
 Job Type: <push to host, purge, or repcat admin>
 Database Host: <hostname>
 Database SID: <SID>
```

This event is sent when the system detects that a previously broken job is now working correctly. This could happen when network connectivity is restored between two database servers, allowing data to be replicated. No action is required when this happens.

> **Note:** This event is only sent when the correction occurs while NA is running. If NA is stopped while the problem is corrected, when NA is restarted no event is sent.

Replication Errors

Event format:

```
TransactionID: <id>
 Call Number: <n>
 Table Name: <TABLE_NAME>
 Transaction Type: UPDATE
 NA Source Host: <hostname>
 DB Destination: <SID>
```

NA sends this event when it encounters a replication error that it cannot resolve on its own. In general, NA attempts to resolve conflicts automatically. Certain conflict types cannot be resolved using the built-in Oracle methods, but can be resolved by NA after they occur. In particular, UPDATE conflicts due to out-of-order transactions, causing a primary key or foreign key constraint violation, can be resolved by re-executing the transaction after a short wait.

Additionally, conflicts due to UPDATES of items that have already been deleted can by resolved by removing the failed transaction. These two actions are taken by NA during its regular replication monitoring run, which happens at several minute intervals. If these corrective actions do not fix the problem, an event is sent.

Typically, to fix an UPDATE error requires editing the object in question to set its values to those expected by the transaction, and then re-executing the transaction (also known as retrying the error). To do this:

1. Open the View Distributed Error page and click the Details option in the Actions column. All of the database attributes of the affected object are displayed. See "Distributed Conflict List" on page 47.

2. Locate the object in the NA system (or simply update the database directly using SQLPlus or some other tool), and change all of its attributes to those listed in the Original Value column on the View Distributed Error Details page.

3. Retry the error.

Replication errors can also occur by disconnecting from the NA proxy, immediately after connecting to it. To prevent such errors, follow these steps:

1. On each NA core, add the following line to the `adjustable_options.rcx` file:

   ```
   <option name="proxy/mm_log_offset">1000</option>
   ```

   This setting adds 1000ms (1 second) to the end date.

2. Reload the `.rcx` settings by doing one of the following:

   - On the User Interface page (**Admin** > **Administrative Settings** > **User Interface**) of the NA console, click **Save**.

   - From the NA proxy, run the `reload server options` command.

   - Restart the NA management engine.

For information about `.rcx` files, see Working with .rcx Files.

Uniqueness Conflict

Event format:

```
Server: <host>
Object: <Object>
Conflicting Value: <value>
Refer to the replication documentation for instructions on correcting this conflict.
```

This event is sent when NA detects a conflict in a uniqueness constraint. You will receive an event per NA Core, since the conflicts are local to each NA Core. To correct a naming conflict, go to one NA Core and update the names for the affected objects. Both the renamed <NAME>.<SID> and <NAME> should be edited to force an update on the other NA Cores.

To correct a rule priority conflict, go to the Device Password Rule Priority Reset page click the Reset Priority button. See "Device Password Rule Priority Reset Page" on page 50.

If this does not solve the problem, you will need to manually edit the rules on each NA Core, setting the priority order correctly and verifying existence of correct rules. When finished, return to the Device Password Rule Priority Reset page and click the Reset Priority button.

Time Synchronization Warning

Event format:

```
Time difference: <N> seconds
 Local Core: <hostname>
 Remote Core: <hostname>
```

NA replication conflict resolution depends on a latest timestamp method. To work correctly, this requires different NA servers' clocks to differ by only a small amount. To correct this problem, make sure that the time is synchronized on the NA server systems across the NA Mesh.

RMI Error

Event format:

```
Local Core: <hostname>
 Remote Core: <hostname>
 Error: <Exception text>
```

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

1. The host, to which the server cannot connect, is up and running.

2. The NA instance on that host is running.

3. From a command line, enter `ping <host>` to ensure that network connectivity exists between servers.

4. From a command line, enter telnet <host> to port 1099 (or whatever your RMI listen port is set to) to ensure that RMI connections are being accepted. If working correctly, you should get back some data that includes the text string "java.rmi.MarshalledObject".

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

Deferred Transactions Exceed Threshold

Event format:

```
Exceeded number of deferred transactions in threshold.
```

This event occurs when NA determines that the deferred transaction queue, as reported by the DEFTRAN view, exceeds a certain threshold. This typically means that the NA Core reporting this event cannot push transactions quickly enough to other NA Cores.

To troubleshoot this issue:

1. Verify that network connectivity exists between the database servers in the NA Mesh.

2. Verify that the network connections between the database servers in the NA Mesh are not experiencing bandwidth congestion or other network issues.

Correcting these issues should enable the databases to catch up with their transaction pushes. In addition, reducing the amount of tasks or other operations on the NA Cores can help the system recover more quickly.

Deferred LOBs Exceed Threshold

Event format:

```
Exceeded number of deferred LOBs in threshold.
```

This event is identical to Deferred Transactions Exceed Threshold event, except the view being monitored is the DEFLOB view.

Using the NA Distributed System Pages

When you install the Distributed System software, the NA user interface includes specific Distributed System pages to help you monitor and administer the system.

Distributed Monitor Results Page

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

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

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

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

- **Scheduled Jobs** — Oracle contains a job scheduler. Replication requires certain jobs be setup using that job scheduler. In particular, there are jobs on each NA Core database that:

  - Push transactions

  - Purge previously pushed transactions from the log

  - Handle administrative commands for replication

- **Error Transactions** — Error transactions are transaction that have been pushed from a remote server and for some reason were not able to execute correctly on the local server.

- **Uniqueness Conflicts** — Certain NA database constraints restrict columns to unique values. In a distributed environment, these constraints can be violated when updates are made on two different NA Cores where the unique column is set to the same value. These conditions are captured by the Replication Conflict Resolution System and logged. NA cannot automatically resolve these conflicts. They must be resolved manually.

- **Local NA Core Definition** — The local NA Core must be able to determine which entry in the RN_CORE

table it is. If the "The *local core for this system is undefined.*" error message is displayed, the CoreHostname property needs to be updated for the NA Core. This can be done using the Edit Core page. For information, see "Edit Core Page" on page 49.

> **Note:** When this condition occurs, the NA Management Engine's log file will contain the following text: "*Fatal error - could not assign local core.*"

The CoreHostname value can be either the DNS, *etc/hosts* value, or an IP address. If you are using an NA server with multiple IP addresses, you might need to tell NA which IP address to use. This is done by adding the following setting to the *distributed.rcx* file:

```
<option name="distributed/NA_server_local_ip">A.B.C.D</option>
```

> **Note:** The *distributed.rcx* file is located in the same location as the *appserver.rcx* file. For the location of the files, see "Oracle 11g: Adding a New NA Core to an Existing NA Mesh" on page 26

The value A.B.C.D should be replaced with the appropriate NAT IP address for the NA server and should match the RN_CORE table's CoreHostname value for that NA Core.

Distributed Error List

The Distributed Error List page displays distributed transaction errors on the current NA Core. This information is useful for finding error conditions that require administrative action. Keep in mind that all conflicts are displayed until they are:

- Manually resolved
- Automatically resolved by NA
- Manually deleted without resolution (not recommended)

To open the Distributed Error List, on the menu bar under Admin select Distributed and click Error List. The Distributed Error List opens.

Distributed Error List Page Fields

Field	Description
Transaction ID	The ID of the transaction that was placed in the deferred error queue due to the conflict.
Call No.	A transaction can update more than one row in the database. Each row will have a separate call number. For example, for a transaction that updates seven rows, there will be seven entries with the same Transaction ID, but different call numbers.
Package Name	The name of the table on which the conflict occurred.
Proc Name	The type of transaction, for example UPDATE or DELETE.
Source	The database from which the transaction originated.
Status	Statuses can include: • new — This is a new error. The system has not attempted to deal with it yet. Keep

Field	Description
	in mind that certain types of errors require waiting for other transactions to complete before they can be retried. • alert_sent — The system has sent an alert. • to_retry — The system has noted this error and is waiting before retrying. • old — The system noted the error and retried it. However, the retry did not work.
Actions	You can select the following options: • Detail — Opens the View Distributed Error page, where you can view details about the transaction in that row. • Delete — Deletes the error from the database. • Retry — Re-execute the transaction.

View Distributed Error Page

The View Distributed Error page displays the detail for a particular distributed transaction error. This information is useful for providing the details necessary to correct the error.

To open the View Distributed Error page:

1. On the menu bar under Admin select Distributed and click Error List. The Distributed Error List opens.
2. In the Actions column, click the Detail option. The View Distributed Error page opens. The Details section displays the column(s) from the table named in the Package Name field that contains conflicting data.

Distributed Conflict List

The Distributed Conflict List page displays the uniqueness constraint conflict list. This provides information about uniqueness conflicts that will need to be manually corrected to ensure that the databases in the NA Mesh are in sync.

To open the Distributed Conflict List, on the menu bar under Admin select Distributed and click Conflict List. The Distributed Conflict List opens.

Distributed Conflict List Page Fields

Field	Description
Object	The database table on which the conflict occurred.
Site Name	The core in the NA Mesh for which the conflict occurred (actually the database SID).
Value	The value of the column for which there was a uniqueness conflict. Keep in mind that NA does not automatically resolve uniqueness conflicts. NA simply appends the name of the NA Core from which the transaction originated to the conflicting value so that the transaction can complete.

Field	Description
Create Date	The Date the conflict occurred.
Status	Statuses include: • new — This is a new conflict. The system has not yet alerted the system administrator. • alert_sent — The system has sent an alert.
Actions	You can select the following options: • Detail — Opens the View Distributed Conflict page, where you can view details on an individual uniqueness constraint. See "View Distributed Conflict Page" below. • Delete — Deletes the conflict from the database.

View Distributed Conflict Page

The View Distributed Conflict page provides details on a specific uniqueness constraint.

To open the View Distributed Conflict page:

1. On the menu bar under Admin select Distributed and click Conflict List. The Distributed Conflict List opens.
2. In the Actions column, click the Detail option. The View Distributed Conflict page opens.

Table Name is the name of the database table on which the uniqueness conflict occurred. Object IDs are the primary keys of the objects in the database.

If a key has only one column, it is displayed in Object ID 1 field. Some keys are formed from multiple columns. Consequently, those values are displayed in the Object ID 1, Object ID 2, and Object ID 3 fields.

Site Reassignment Page

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

To open the Site Reassignment page, on the menu bar under Admin select Distributed and click Site Reassignment. The Site Reassignment opens. You can select NA Cores from the drop-down menu.

List Cores Page

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

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

List Cores Page Fields

Field	Description
Name	The NA Core's name.
Core Hostname	The hostname of the NA Core's NA server.
Is Master?	Is the NA Core the Master Definition? (Yes or No)
Timezone Offset	The time zone offset of the actual NA Core server.
Status	Currently, there is only Normal status.
Realm	The default Realm for the NA Core.
Actions	You can select the following option: • Edit — Open the Edit Core page. See "Edit Core Page" below.

Edit Core Page

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

To open the Edit Core page:

1. On the menu bar under Admin select Distributed and click Core List. The List Cores page opens.
2. In the Actions column, click the Edit option. The Edit Core page opens.

You can complete the following fields:

- Name — Enter the NA Core name.
- Database Identifier — Enter the Database Identifier. This is the name of the NA Core's database as it appears in the *tnsnames.ora* file. The Database Identifier is corresponds to Oracle's SID (System Identifier). This information is needed to make connections to a particular Oracle instance on a server. Multiple instances may be running on any given server, but each will have different SIDs.
- Core Hostname - Enter the hostname of this NA Core's server.
- RMI Port — Enter the RMI port. RMI (Remote Method Invocation) is Java's remote procedure call protocol. The distributed system makes RMI calls between NA servers in the NA Mesh to transfer information about scheduled tasks, system settings, software images, and so on.
- Database hostname — Enter the Database hostname.
- Database Port — Enter the port on the database server with which NA communicates with the database.
- Timezone Offset — Select a Timezone offset from the drop-down menu.
- Replication Admin User — Enter the name of the Replication Admin user. The Replication Admin user is created and used by the Oracle database to manage replication.
- Replication Password — Enter the Replication Admin user's password.
- Confirm Replication Password — Re-enter the Replication Admin user's password.
- Comments — Add any comments about the NA Core.
- Realm Name — Enter the Realm in which the NA Core resides. For information about segmenting devices, see the *NA User Guide*.

Device Password Rule Priority Reset Page

The Device Password Rule Priority Reset page enables you to reset device password rule priorities in the event that a uniqueness constraint conflict occurs for those objects.

To open the Reset Password Priority page, on the menu bar under Admin select Distributed and click Device Password Rule Priority Reset. The Device Password Rule Priority Reset page opens.

Click the Reset Priority button to reset the device password rule priorities.

Renew Configuration Options Page

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

To open the Renew Configuration Options page, on the menu bar under Admin select Distributed and click Renew Configuration Options. The Renew Configurations Options page opens.

Click the Renew Config Options button to ensure that all options on the NA Core are in sync with the rest of the NA Mesh.

Chapter 4: Failover and Recovery

This chapter contains the following topics:

- "Failover" below
- "Recovery" below

Failover

When the network has been configured to failover, if an NA Core fails, users can continue to access all data in the system using a different NA Core. All Sites that were originally managed by the failed NA Core can be pointed to a new NA Core using the Site Reassignment page. For information, see "Site Reassignment Page" on page 48. Procedures for system recovery will vary depending on how the remote server failed.

Recovery

There are three basic recovery scenarios:

- Loss of network connectivity
- Loss of an NA server
- Loss of a Database server

Loss of Network Connectivity

In the case of lost network connectivity, failover occurred due to network issues. No problems occurred with the NA server or with the Oracle database server. Recovery consists of the following steps:

1. Resolve the network issues.
2. Reset Sites that had been reassigned back to their original NA Core. This can be accomplished in NA. See "Site Reassignment Page" on page 48.
3. If any drivers have been added to the system during the outage, click the "Reload Drivers" button on the Start/Stop Services page. This action reloads the driver files and pushes them across to other NA Cores in the NA Mesh. This action should be performed on the NA server where the drivers were added. For information, see the *NA User Guide*.
4. If any system settings have been modified during the outage, use the "Renew Config Options" page to make sure options are synced across the NA Mesh. See "Renew Configuration Options Page" on page 50.
5. If any NA Cores have lost connectivity for a long period of time, restart the NA Core server that lost connectivity after data sync so as to reload certain Site data and avoid exception errors due to obsolete data.

Once the network issues are resolved, the system should recover as replication syncs data between the databases.

Loss of an NA Server

In cases where the NA server suffers a failure that requires re-installation of the server, recovery consists of the following steps:

1. During NA installation, select the "Use existing database" option. In addition, the database selected should be the one the failed server was previously using.

2. Add the *distributed.rcx* file from the *ReplicationScriptToolBundle* to the directory where the *appserver.rcx* file resides.

3. Re-add any NA Core-specific special case options for patches and support issues.

4. Restart NA.

5. Reset Sites that had been re-assigned back to their original NA Core. This can be accomplished in NA. See "Site Reassignment Page" on page 48.

6. If any drivers have been added to the system during the outage, click the "Reload Drivers" button on the Start/Stop Services page. This action reloads the driver files and pushes them across to other NA Cores in the NA Mesh. This action should be performed on the NA server where the drivers were added. For information, see the *NA User Guide*.

7. If any system settings have been modified during the outage, use the "Renew Config Options" page to make sure options are synced across the NA Mesh. See "Renew Configuration Options Page" on page 50.

8. Edit the original NA Core to modify any parameters that may be different (perhaps the installation happened on a new server with a different hostname).

9. Copy the software images repository from a good NA Core to the recovered NA Core.

Loss of a Database Server

In the case of a lost database server, the NA server is still running but cannot access the database. The database will need to be rebuilt and replication setup again on the database. Recovery consists of the following steps:

1. Pause or delete any tasks that appear to be stuck waiting or pending and not running because they are associated with the lost server. You can perform this action on another NA server in the NA Mesh.

2. If the server that was lost is the Master Definition server for replication, the procedure described for relocating the Master Definition server should performed. See "Changing the Master Definition Site" on page 30.

3. For instructions on adding a new NA Core, see "Oracle 11g: Adding a New NA Core to an Existing NA Mesh" on page 26.

4. Reset Sites that had been reassigned back to their original NA Core. This can be accomplished in NA. See "Site Reassignment Page" on page 48.

5. In NA, edit the NA Core that was recovered to make sure all information is correct for the new setup. For information, see the *NA User Guide*.

Chapter 5: Troubleshooting

This chapter contains the following topics:

RMI Errors in the NA Console

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

RMI Error

Event format:

```
Local Core: <hostname>
Remote Core: <hostname>
Error: <Exception text>
```

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

1. The host that the server cannot connect to is up and running.

2. The NA instance on that host is running.

3. From a command line, enter ping `<host>` to ensure that network connectivity exists between servers.

4. From a command line, enter `telnet <host> to port 1099` (or whatever your RMI listen port is set to) to ensure that RMI connections are being accepted. If working correctly, you should get back some data that includes the text string `java.rmi.MarshalledObject`.

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

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

> **Note:** This solution is for static IP environments only.

1. In a text editor such as WordPad or vi, open the following file:

   - *Windows*: `<Drive>:\Windows\System32\drivers\etc\hosts`

   - *Linux*: `/etc/hosts`

2. Set the localhost line to read:

   `127.0.0.1 localhost`

3. For each NA core server, add a line in the following format:

   <xx.xx.xx.xx> *<NA.example.com>* *<NA>*

   - Replace *<xx.xx.xx.xx>* with the IP address of the NA core server.

   - Replace *<NA.example.com>* with the fully-qualified domain name of the NA core server.

   - Replace *<NA>* with the short hostname of the NA core server.

4. Repeat step 3 until the `hosts` file includes all NA core servers in the distributed system environment.

5. To use the updated hosts information, restart the NA core server.

Oracle Replication Setup

The Oracle replication setup process is complicated. The following guidelines help prevent problems with replication setup.

- Ensure that the *tnsnames.ora* file is setup correctly on each system. You can verify correct setup by running SQLPlus on each system to make certain that you can connect to other databases. For example, if you have two systems (A and B), run SQLPlus on A to connect to the database on B using the following command:
  `sqlplus USER/PWD@DATABASE_B`.

- It may be necessary to remove or comment out the `NAMES.DEFAULT_DOMAIN` property entry from the *sqlnet.ora* file for the appropriate connections to work. Consult with your Oracle DBA if you have questions.

- If the replication setup process fails at any step, it is recommended that you follow the process for removing replication from the entire NA Mesh (see "Removing Replication Across the Entire NA Mesh" on page 31), and then follow the replication setup sets again (see "Oracle 11g: Set Up Initial Replication" on page 16).

Replication Tools and Commands

The following tools and commands can help to isolate problems with replication setup.

To ensure that database links are working:

- Using SQLPlus, enter: `SELECT CURRENT_TIMESTAMP FROM DUAL@<OTHER_SID>;`

- Using SQLPlus, enter: `SELECT DBLINK FROM DBA_REPSITES;`

- From a command line, enter: `tnsping <OTHER_SID>`

- From a command line, enter: `ping <HOST>`

To ensure the replication group is working:

- Using SQLPlus, enter: `SELECT GNAME,STATUS FROM DBA_REPGROUP;`

To ensure there are no errors for replication admin requests:

- Using SQLPlus, enter: `SELECT STATUS,REQUEST,ONAME FROM DBA_REPCATLOG;`

Common Errors

Some errors can occur during replication setup. Ones that have been reported are noted below, with suggested corrective action. In some cases, it might be necessary to remove replication across the NA Mesh and re-setup replication. For information, see "Removing Replication Across the Entire NA Mesh" on page 31.

Error

```
SQL> EXEC DBMS_REPCAT.RESUME_MASTER_ACTIVITY(gname=>' <REPGROUP_NAME>');BEGIN DBMS_
REPCAT.RESUME_MASTER_ACTIVITY (gname=>'<REPGROUP_NAME>'); END;
*
ERROR at line 1:
ORA-23419: regenerate replication support before resuming master activity
...
```

Corrective Action

An internal Oracle error has prevented replication support from being generated correctly for one or more objects. To correct this, it will be necessary to remove and re-setup replication for the NA Mesh.

Error

```
ORA-23374: object group "PUBLIC"."REPLNA" already exists
ORA-06512: at "SYS.DBMS_SYS_ERROR", line 95
ORA-06512: at "SYS.DBMS_REPCAT_MAS", line 798
ORA-06512: at "SYS.DBMS_REPCAT", line 548
```

Corrective Action

An internal Oracle error has prevented replication support from being generated correctly for one or more objects. To correct this, it will be necessary to remove and re-setup replication for the NA Mesh. For information, see "Removing Replication Across the Entire NA Mesh" on page 31.

Error

```
BEGIN DBMS_REPCAT.DROP_MASTER_REPOBJECT('SYSTEM', 'RN_DEVICE_VIEW','TABLE',FALSE)
; END;
*
ERROR at line 1:
ORA-00060: deadlock detected while waiting for resource
...
```

Corrective Action

Rerun the deletion script you were running when this occurred. You may see other errors due to certain actions already having been done, but that is expected.

Error

```
BEGIN DBMS_REPCAT.DROP_MASTER_REPOBJECT('SYSTEM',' RN_DEVICE_VIEW','TABLE',FALSE)
; END;
*
ERROR at line 1:
ORA-23308: object SYSTEM.RN_DEVICE_VIEW does not exist or is invalid
...
```

Corrective Action

Rerun the deletion script you were running when this occurred. You may see other errors due to certain actions already having been done, but that is expected.

Appendix A: Common Procedures

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

- "Start, Stop, or Restart All NA Services" below
- "Disable All NA Services" on the next page
- "Working with .rcx Files" on the next page

Start, Stop, or Restart All NA Services

Stopping the NA services before changing the NA configuration prevents conflicting data from being stored in the NA database. Some procedures call for restarting the NA services to read the updated configuration.

To start all NA 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 Syslog Server**
  - **TrueControl TFTP Server**
  - **TrueControl SA Client**

- *Linux*: Run the following command:

  /etc/init.d/truecontrol start

To stop all NA services

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

- *Linux*: Run the following command:

  /etc/init.d/truecontrol stop

To restart all NA services

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

    - **TrueControl ManagementEngine**

    - **TrueControl FTP Server**

    - **TrueControl Syslog Server**

    - **TrueControl TFTP Server**

    - **TrueControl SA Client**

- *Linux*: Run the following command:

    `/etc/init.d/truecontrol restart`

Disable All NA Services

Some procedures call for disabling automatic startup of the NA services on system boot.

To disable all NA 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 Syslog Server**

    - **TrueControl TFTP Server**

    - **TrueControl SA Client**

- *Linux*:

    `mv /etc/rc.d/rc5.d/S99truecontrol /etc/S99truecontrol`

Working with .rcx Files

The HPE Network Automation Software (NA) property files use the `.rcx` extension. NA reads `.rcx` files in reverse alphabetical order. If a given setting is in multiple `.rcx` files, NA uses the last-read value. Thus, the settings in the `adjustable_options.rcx` file take precedence over the settings in the other `.rcx` files installed with NA.

> **Note:** At startup, NA reads *all* files in the `jre` directory and interprets their contents for NA configuration options. For this reason, save all backup copies of `.rcx` files outside the root NA directory.

In Horizontal Scalability environments, NA shares the actual values of most settings, not the `.rcx` files, across the NA cores. When a setting is modified on one NA core, that setting is replicated to the other NA cores. If an NA core is not operational during the change replication, that NA core does not receive the change. In that case, at a later time, use the Admin > Distributed > Renew Configuration Options page to push changes to other NA cores.

> **Tip:** The distributed system options section of the `appserver.rcx` file lists the settings that are specific to one NA core and are not shared across the NA cores.

Some configuration changes require `.rcx` file modifications. The `.rcx` files are located in the following directory:

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

> **Caution:** Always edit `.rcx` files with care. These files use XML format. If a `.rcx` file change results in invalid XML, the NA console might not start correctly.

> **Tip:** It is recommended to make all configuration changes in the `adjustable_options.rcx` file. NA patch installations and product upgrades might overwrite any of the other NA-installed `.rcx` files.

The general procedure for changing `.rcx` files is as follows:

1. Back up the `.rcx` file to a location outside the `<NA_HOME>` directory.

   (NA reads all `.rcx` files within the NA directory structure.)

2. Add new content or update existing content as described in the instructions.

3. Save the `.rcx` file.

4. Reload the `.rcx` settings by doing *one* of the following:

   - In the NA console, on the Admin > Administrative Settings > User Interface page, click **Save**.

   - Run the `reload server options` command from the NA proxy.

   - Restart the NA services.

> **Tip:** Some changes do not take effect until the NA services have been restarted.

Appendix B: Restoring NA Replication After Upgrading to Oracle 11g

This appendix describes the process for upgrading the database of an HPE Network Automation Software (NA) Multimaster Distributed System from Oracle 10g to Oracle 11g, including how to restore NA replication after the upgrade. At a high level, the process involves the following steps:

1. Remove replication, which creates standalone NA cores.
2. Upgrade the database.
3. Set up replication using the database from only one of the NA cores.

> **Note:** If you also plan to upgrade the NA software, do not mix that upgrade with the database upgrade. Upgrade NA to version NA10.20, and then verify that replication works correctly before upgrading the database.

Prerequisites

Before upgrading the Oracle database in an NA Multimaster Distributed System, verify that the following prerequisites have been met:

- The existing NA database user name is a custom value (not SYSTEM). If the NA schema is currently owned by the SYSTEM user, consider doing a fresh installation of the NA Multimaster Distributed System on Oracle 11g.
- NA database replication is correctly configured and working with Oracle 10g for all cores in the NA mesh.
- The ReplicationScriptTool bundle files are installed on a Java-capable system.
- The information listed in Table 1 is available. If the NA mesh is larger than two cores, also collect the information for the additional cores.

Database Upgrade Process

To upgrade the database from Oracle 10g to Oracle 11g within an NA mesh, follow this process:

1. Stop all NA services on all NA core servers. See "Start, Stop, or Restart All NA Services" on page 57.
2. Remove all replication from the NA mesh.
   a. If the NA mesh consists of three or more cores, remove the additional cores as described in "Removing an NA Core from an Existing NA Mesh" on page 29.

   Note the following:
   - On each NA core server, move the `distributed.rcx` file to a location outside of `<NA_HOME>`.
   - Do not uninstall the NA application.
   b. Remove replication from the final two cores in the NA mesh as described in "Removing Replication

Across the Entire NA Mesh" on page 31.

Note the following:

- On each NA core server, move the `distributed.rcx` file to a location outside of `<NA_HOME>`.
- Do not uninstall the NA application.

3. Verify that each NA core behaves as a standalone instance of NA:

Enter the following commands:

```
CONNECT <USER>/<PASSWORD>@<MASTER1_SID>;
UPDATE RN_CORE SET COMMENTS = 'Test2 from sid1' where CoreID=1;
COMMIT;
```

wait 10 to 30 seconds

```
CONNECT <USER>/<PASSWORD>@<MASTER2_SID>;
SELECT Comments from RN_CORE where CoreID=1;
```

Enter this from every NA Core in the NA Mesh. Check the values each time on all of the other NA Cores to ensure that no replication exists among the databases that were formerly in the system.

4. Upgrade each database server to a supported version of Oracle 11g.

5. Prepare Master 1 for replication as described in "Task 3: Prepare Master 1 for Replication (Oracle 11g)" on page 19.

6. Prepare Master 2 for replication as follows:

   a. With the NA management engine on NA2 shut down, log on to Master 2 as the SYS user, and then delete the NA database user while retaining the NA tablespace.

   For example, in the Oracle Enterprise Manager (preferred), delete `nauser`, or at the `sqlplus` prompt, enter the following command:

   ```
   DROP USER nauser CASCADE;
   ```

   > **Note:** Use the cascade delete option to completely delete the user and empty its tablespace.

   b. On Master 2, verify that the NA tablespace is now empty.

   c. Beginning with step 2, complete the steps of "Task 4: Prepare Master 2 for Replication (Oracle 11g)" on page 21.

   > **Note:** The NA tablespace exists from the prior replication setup. Do not create it at this time.

7. Prepare additional database servers by repeating step 6 for each database server that was previously in the NA mesh.

8. Configure replication between Master 1 and Master 2 as described in "Task 5: Configure Replication (Oracle 11g)" on page 22.

In the `ReplicationScriptTool.properties` file, `mode` should be set to `initial`.

> **Note:** Do *not* install the NA application.

9. With the NA management engines on NA1 and NA2 shut down, on each of NA1 and NA2, move the `distributed.rcx` file from the known location outside of `<NA_HOME>` to the following directory:

- *Windows*: <NA_HOME>\jre

- *Linux*: <NA_HOME>/jre

10. Start all NA services on both NA core servers, NA1 and NA2. See "Start, Stop, or Restart All NA Services" on page 57.

11. Connect to the NA console on both core servers, NA1 and NA2, (as user nauser), and ensure that there are no issues.

> **Tip:** An RMI error message at this point can mean that NA is unable to identify the localhost. For more information, see "RMI Errors in the NA Console" on page 53.

12. Verify the NA installation:

    a. Edit an object on one NA server (for example, a Comments field for a device).

    b. Wait a minute and then verify that the updated comment exists on the second server.

    c. Check the status of the Oracle Distributed Monitor in the NA console to ensure that no problems are being reported. (For information, see "Distributed Monitor Results Page" on page 45.) It could take up to five minutes for this monitor to initially run.

13. For each additional database server that was previously in the NA mesh, configure replication as described in "Task 5: Configure Replication (Oracle 11g)" on page 22.

    Note the following differences:

    - Substitute "Master 3" for each reference to "Master 2."

    - For step 4, use the ReplicationScriptTool.properties file from setting up replication between Master 1 and Master 2.

        ○ Verify the server information for the initial core and second master site.

        ○ Uncomment the Additional Master site configuration area, and replace the text REPLACEME with the appropriate value for that variable. Use the fully-qualified hostnames to identify the database and core servers. For example:

        ```
        db.server.1=Master3.example.com
        db.port.1=1521
        db.name.1=SID3
        db.username.1=nauser
        db.password.1=napass
        app.server.1=NA3.example.com
        app.rmiport.1=1099
        app.corename.1=Core 3
        ```

        ○ Set the mode to add_server.

        ○ For step 13, the spool file contains the following acceptable warning that occurs after the DBMS_REPCAT.SUSPEND_MASTER_ACTIVITY command:

        ```
        'ORA-23326: object group <yourgroup>.<your_replication_ group> is quiesced'
        ```

14. With the NA management engine on the newly added NA core servers shut down, on each newly added NA core server, move the distributed.rcx file from the known location outside of <NA_HOME> to the following directory:

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

- *Linux*: `<NA_HOME>/jre`

15. Start all NA services on each newly added NA core server. See "Start, Stop, or Restart All NA Services" on page 57.

16. Connect to the NA console on each newly added core server (as user `nauser`), and ensure that there are no issues.

> **Tip:** An RMI error message at this point can mean that NA is unable to identify the localhost. For more information, see "RMI Errors in the NA Console" on page 53.

17. Verify the NA installation:

    a. Edit an object on one NA server (for example, a Comments field for a device).

    b. Wait a minute and then verify that the updated comment exists on the other servers in the NA mesh.

    c. Check the status of the Oracle Distributed Monitor in the NA console to ensure that no problems are being reported. (For information, see "Distributed Monitor Results Page" on page 45.) It could take up to five minutes for this monitor to initially run.

18. Optional. Add new cores (up to the supported limit) to the NA mesh as described in "Oracle 11g: Adding a New NA Core to an Existing NA Mesh" on page 26.

Appendix C: Extracting Data from the NA Database

This appendix includes the steps for extracting information about the Multimaster Mesh from the NA database.

> **Note:** This information is also available on the appropriate Edit Core page. For information, see "Edit Core Page" on page 49.

To extract this data from the NA database, run the following SQL queries against the NA database. The query will return multiple rows, one row for each database core.

For the NA core server hostname, run the following:

```
SELECT CoreHostname FROM RN_CORE;
```

For the NA core server RMI listening port, run the following:

```
SELECT CoreRMIPort FROM RN_CORE;
```

For the Database Identifier (SID), run the following:

```
SELECT DatabaseIdentifier FROM RN_CORE;
```

For the NA database server hostname, run the following:

```
SELECT DatabaseHostname FROM RN_CORE;
```

For the NA database server listening port, run the following:

```
SELECT DatabasePort FROM RN_CORE;
```

For the Timezone offset, run the following:

```
SELECT TimezoneOffset FROM RN_CORE;
```

For the Replication User Name, run the following:

```
SELECT ReplicationAdminUser FROM RN_CORE;
```

For example, in a two NA Core Multimaster Mesh, run the following queries:

```
SELECT CoreHostname FROM RN_CORE;
SELECT CoreRMIPort FROM RN_CORE;
SELECT DatabaseIdentifier FROM RN_CORE;
SELECT DatabaseHostname FROM RN_CORE;
SELECT DatabasePort FROM RN_CORE;
SELECT TimezoneOffset FROM RN_CORE;
SELECT ReplicationAdminUser FROM RN_CORE;
```

You will receive the following:

```
SQL> SELECT CoreHostname FROM RN_CORE;
COREHOSTNAME
------------------------------
10.255.132.177
10.255.136.203
SQL> SELECT CoreRMIPort FROM RN_CORE;
CORERMIPORT
------------------------------
1099
1099
SQL> SELECT DatabaseIdentifier FROM RN_CORE;
DATABASEIDENTIFIER
------------------------------
core1
core2
SQL> SELECT DatabaseHostname FROM RN_CORE;
DATABASEHOSTNAME
------------------------------
10.255.55.30
10.255.56.30
SQL> SELECT DatabasePort FROM RN_CORE;
DATABASEPORT
------------------------------
1521
1521
SQL> SELECT TimezoneOffset FROM RN_CORE;
TIMEZONEOFFSET
------------------------------
-8
-8
SQL> SELECT ReplicationAdminUser FROM RN_CORE;
REPLICATIONADMINUSER
------------------------------
repadmin
repadmin
```

From this data, you can start filling out the *ReplicationScriptTool.properties* file as follows:

```
# This .0 server will be the master definition site, and is the default
# initial core installed.
db.server.0=10.255.55.30
db.port.0=1521
db.name.0=core1
db.username.0=admin
db.password.0=rendition
app.server.0=10.255.132.177
app.rmiport.0=1099
app.corename.0=Core 1

# Second Master site
db.server.1=10.255.56.30
db.port.1=1521
db.name.1=core2
db.username.1=admin
db.password.1=rendition
app.server.1=10.255.136.203
app.rmiport.1=1099
app.corename.1=Core 2

# From UTC
timezone_offset=-8
replication_user=repadmin
replication_password=<YOU WILL NEED TO GET THIS FROM YOUR DBA>
replication_group_name=<YOU WILL NEED TO GET THIS FROM YOUR DBA>
```

We appreciate your feedback!

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

Feedback on Multimaster Distributed System on Oracle Guide, July 2016 (Network Automation Software 10.20)

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.