

HP Network Automation

Software Version: 9.11

Administration Guide

Document Release Date: October 2012
Software Release Date: October 2012



Legal Notices

Warranty

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

The information contained herein is subject to change without notice.

Restricted Rights Legend

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

Copyright Notices

© Copyright 2011–2012 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

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

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

UNIX® is a registered trademark of The Open Group.

Acknowledgements

ANTLR, Apache, Bouncy Castle, GNU, Jaxen, Jython, Netaphor, MetaStuff, Radius, Sleepcat, TanukiSoftware

Documentation Updates

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

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

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

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

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

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

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

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

Support

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

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

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

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

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

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

To find more information about access levels, go to:

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

To register for an HP Passport ID, go to:

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

Contents

1	About This Guide	7
	Revision History	7
2	Tuning NA Performance	9
	Disabling the JCS Disk Cache	9
	Tuning the NA Management Engine	10
	Task Scheduling	10
	Maximum Concurrent Tasks	10
	Maximum Data Source Pool Size	11
	Configuring the Java Virtual Machine	12
	Configuring MySQL for NA	14
	Configuring Oracle for NA	15
	Number of Database Connections from NA	15
	Size of the NA Tablespace	16
	Configuring SQL Server for NA	16
3	Extending the Number of Custom Enhanced Fields	17
4	Configuring the Default Setting of the Enforce Save Check Box for New Tasks	19
5	Limiting Device Password Access	21
	Enabling the Access-Limiting Functionality	21
	Using the Access-Limiting Functionality	22
6	Full Text Search of Configuration Text	23
	Enabling Full Text Search of Configuration Text	24
	Standalone NA Core Environment	24
	Horizontal Scalability Environment	24
	Multimaster Distributed System Environment	25
	Enabling Full Text Search on Oracle	25
	Enabling Full Text Search on Microsoft SQL Server	27
	Enabling the Contains (Full Text) Search Operator	28
	Using the Contains (Full Text) Search Operator	28
	Disabling the Contains (Full Text) Search Operator	29
	Disabling Full Text Search	30
	Command Line Reference	31
	Analyze the NA Database in Preparation for Index Generation	31
	Check the Status of the NA Database Index	32
	Generate an Index of the NA Database	32
	Remove the NA Database Index	32

Reindex the NA Database	32
7 Enabling Case-Insensitive Search (Oracle)	33
Affected Fields	33
Search Box	33
Search Criteria	33
Device Selector	34
Reports	34
Enabling Case-Insensitive Search of an Oracle Database.....	36
Disabling Case-Insensitive Search	37
8 Reclaiming Unused Space (Oracle Database)	39
We appreciate your feedback!	41

1 About This Guide

This guide contains a collection of information and best practices for administering HP Network Automation Software (NA). This guide is for an expert system administrator, network engineer, or HP support engineer with experience deploying and managing networks in large installations.

This guide assumes that you have already installed NA and that you are familiar with start-up configuration tasks. To learn more about these tasks, see the *NA Upgrade and Installation Guide* and the NA help.

HP updates this guide between product releases as new information becomes available. For information about retrieving an updated version of this document, see [Documentation Updates](#) on page 2.

Revision History

Table 1 lists the major changes for each new release of this document.

Table 1 Document Changes

Document Release Date	Description of Major Changes
January 2012 (9.10 Patch 2)	First publication for NA version 9.10.
October 2012 (9.11)	<p>Added the following chapters:</p> <ul style="list-style-type: none"> • Tuning NA Performance • Configuring the Default Setting of the Enforce Save Check Box for New Tasks • Limiting Device Password Access • Enabling Case-Insensitive Search (Oracle) <p>Removed the “Upgrading to Java JDK Version 6 Update 29” chapter because 9.11 supersedes this information.</p>

2 Tuning NA Performance

This chapter describes several ways to tune the performance of HP Network Automation Software (NA). It includes the following topics:

- [Disabling the JCS Disk Cache](#) on page 9
- [Tuning the NA Management Engine](#) on page 10
- [Configuring the Java Virtual Machine](#) on page 12
- [Configuring MySQL for NA](#) on page 14
- [Configuring Oracle for NA](#) on page 15
- [Configuring SQL Server for NA](#) on page 16

Disabling the JCS Disk Cache

To disable the JCS disk cache, follow these steps:

- 1 Change to the following directory:
 - *Windows*: <NA_HOME>\jre
 - *UNIX*: <NA_HOME>/jre
- 2 Back up the `cache.ccf` file to a location outside the <NA_HOME> directory.
- 3 In a text editor such as WordPad or vi, open the `cache.ccf` file.
- 4 Verify that the `jcs.default=` line does not include any settings. It should read exactly as follows:

```
jcs.default=
```

- 5 Search for the string `AVAILABLE AUXILIARY CACHES` to locate the following lines:

```
# AVAILABLE AUXILIARY CACHES
jcs.auxiliary.DC=org.apache.jcs.auxiliary.disk.indexed.IndexedDiskCacheFactory
jcs.auxiliary.DC.attributes=org.apache.jcs.auxiliary.disk.indexed.IndexedDiskCacheAttributes
jcs.auxiliary.DC.attributes.DiskPath=/tmp/cache
jcs.auxiliary.DC.attributes.MaxPurgatorySize=10000000
jcs.auxiliary.DC.attributes.MaxKeySize=1000000
jcs.auxiliary.DC.attributes.MaxRecycleBinSize=5000
jcs.auxiliary.DC.attributes.OptimizeAtRemoveCount=300000
jcs.auxiliary.DC.attributes.ShutdownSpoolTimeLimit=60
```

- 6 Comment out each of the lines located in [step 5](#) by inserting a number sign character (#) at the beginning of the line.
- 7 Save the `cache.ccf` file.

- 8 In an NA Horizontal Scalability environment, repeat [step 1](#) through [step 7](#) on each NA server.
- 9 On each NA server, restart all NA services.
 - **Windows:** Open the **Services** control panel. In the list of services, right-click **TrueControl ManagementEngine**, and then click **Restart**.
 - **UNIX:** Run the following command:

```
/etc/init.d/truecontrol restart
```

Tuning the NA Management Engine

This section describes recommended tuning of the NA Management Engine. If you update the maximum number of concurrent tasks, also update the maximum data source pool size and the number of connections from NA to the Oracle database.

Task Scheduling

It is recommended that scheduled tasks be plan to run throughout the day to balance the use of NA server resources.

It is recommend that snapshot tasks occur after the work day ends to capture that day's changes.

Maximum Concurrent Tasks

The maximum number of concurrent tasks tunes the NA task functionality.

The recommended value for the maximum number of concurrent tasks depends on the size of the NA deployment, as described in “Tuning Settings” in the *NA Support Matrix*. A higher value is not necessarily better.

To set the maximum number of concurrent tasks, follow these steps:

- 1 Log on to the NA console as an NA administrator.
- 2 On the Administrative Settings - Server page (**Admin > Administrative Settings > Server**), under Tasks, set Max Concurrent Tasks to the value recommended in “Tuning Settings” in the *NA Support Matrix*, and then click **Save**.



After changing the maximum number of concurrent tasks, see [Maximum Data Source Pool Size](#) on page 11 and [Number of Database Connections from NA](#) on page 15.

Maximum Data Source Pool Size

If you change the Max Concurrent Tasks setting or the Max Concurrent Group Tasks setting or if the expected maximum number of concurrent users of the NA console changes considerably, update the maximum data source pool size configuration.

The correct maximum data source pool size is the sum of the following factors:

- The Max Concurrent Tasks setting
This value is listed under Tasks on the Administrative Settings - Server page.
- The Max Concurrent Group Tasks setting
This value is listed under Tasks on the Administrative Settings - Server page.
- The expected maximum number of concurrent NA users
This number depends on the way your company uses NA.
The All Users page (**Admin >Users**) lists all user accounts that can connect to NA.



- A buffer of 20

To set the maximum data source pool size configuration, follow these steps:

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


```
/etc/init.d/truecontrol stop
```
- 2 To set the maximum data source pool size value, do the following:
 - a Change to the following directory:
 - *Windows*: <NA_HOME>\server\ext\jboss\server\default\deploy
 - *UNIX*: <NA_HOME>/server/ext/jboss/server/default/deploy
 - b Back up the db-ds.xml file to a location outside the <NA_HOME> directory.
 - c In a text editor such as WordPad or vi, open the db-ds.xml file.
 - d Search for the string NASDataSource to locate the following lines:


```
<attribute name="DataSourceName">NASDataSource</attribute>
<attribute name="InitialPoolSize">0</attribute>
<attribute name="MinPoolSize">0</attribute>
<attribute name="MaxPoolSize">50</attribute>
```
 - e Set the MaxPoolSize attribute to the calculated value.

- f Search for the string `NASReportDataSource` to locate the following lines:


```
<attribute name="DataSourceName">NASReportDataSource</attribute>
<attribute name="InitialPoolSize">0</attribute>
<attribute name="MinPoolSize">0</attribute>
<attribute name="MaxPoolSize">50</attribute>
```
 - g Identify, but do *not* change, the value of the `MaxPoolSize` attribute for the NA report data source configuration.

The values of both maximum pool size attributes factor into the calculation of the number of available database connections.
 - h Save the `db-ds.xml` file.
- 3 In an NA Horizontal Scalability environment, repeat [step 2](#) on each NA server.
 - 4 On each NA server, 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 SWIM Server**
 - **TrueControl Syslog Server**
 - **TrueControl TFTP Server**
 - *UNIX*: Run the following command:


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

Configuring the Java Virtual Machine

The recommended configuration of the Java virtual machine (JVM) heap and young generation sizes depend on the size of the NA deployment, as described in “Tuning Settings” in the *NA Support Matrix*.



The JVM configuration is specified in megabytes.

To set the JVM heap and young generation size, follow these steps:

- 1 Change to the directory that contains the JVM configuration files:
 - *Windows*: `<HA_HOME>\server\ext\wrapper\conf`
 - *UNIX*: `<HA_HOME>/server/ext/wrapper/conf`
- 2 Back up the `jboss_wrapper.conf` file to a location outside the `<NA_HOME>` directory.
- 3 In a text editor such as WordPad or `vi`, open the `jboss_wrapper.conf` file.
- 4 Search for the string `initmemory` to locate the lines similar to the following lines:

```
# Initial Java Heap Size (in MB)
wrapper.java.initmemory=8192
# Maximum Java Heap Size (in MB)
wrapper.java.maxmemory=8192
```

- 5 Compare the values of the `wrapper.java.initmemory` and `wrapper.java.maxmemory` parameters to the minimums given for the initial and maximum Java heap size, respectively, in “Tuning Settings” in the *NA Support Matrix*.
 - If the values meet or exceed the recommendations, no action is required and you can stop here.
 - If the values are lower than the recommendations, continue with [step 6](#).
- 6 If necessary, set the `wrapper.java.initmemory` and `wrapper.java.maxmemory` parameters to the minimums given for the initial and maximum Java heap size, respectively, in “Tuning Settings” in the *NA Support Matrix*.
- 7 Set the young generation size as follows:
 - a To determine whether the young generation size has been set previously, search for the string `-Xmn`.
 - If this string is in the file, edit this line to set the recommended value for the young generation size in “Tuning Settings” in the *NA Support Matrix*.
For example:

```
wrapper.java.additional.3=-Xmn2730m
```
 - If this string is not in the file, add continue with [step b](#).
 - b Search for the string `Additional` to locate the Java Additional Parameters section.
 - c After the last uncommented line in this section, add the following line:

```
wrapper.java.additional.N=-XmnYGm
```
 - d In the newly added line, make the following substitutions:
 - Replace `N` with the next number in the sequence of uncommented `wrapper.java.additional` parameters.
For example, if the `wrapper.java.additional.11` parameter is uncommented and the `wrapper.java.additional.12` parameter is commented out with a number sign (`#`), set `N` to `12`.
 - Replace `YG` with the recommended value for the young generation size in “Tuning Settings” in the *NA Support Matrix*.
For example:

```
wrapper.java.additional.12=-Xmn2730m
```
- 8 Restart all NA services.
 - **Windows:** Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Restart**:
 - **TrueControl ManagementEngine**
 - **TrueControl FTP Server**
 - **TrueControl SWIM Server**
 - **TrueControl Syslog Server**
 - **TrueControl TFTP Server**
 - **UNIX:** Run the following command:

```
/etc/init.d/truecontrol restart
```

Configuring MySQL for NA

Restricting MySQL to a small number of concurrent threads can reduce NA performance. To avoid this problem, configure MySQL to use an infinite number of threads. This configuration varies across versions of MySQL. NA ships with MySQL 5.0.8, which interprets the value 20 for `innodb_thread_concurrency` as infinite.

To set the number of concurrent threads in the MySQL configuration, follow these steps:

- 1 Locate the file to modify.

The global configuration file is commonly located as follows:

- *Windows*: `<Drive>:\MySQL\my.ini`
- *UNIX*: `/etc/my.cnf`

- 2 Back up the file identified in [step 1](#).

- 3 Open the configuration file in a text editor, such as WordPad or vi.

- 4 To set infinite thread concurrency, edit the configuration file as follows:

- a Locate the line that contains the string `innodb_thread_concurrency=`.
- b Set infinite thread concurrency by updating this line to read:

```
innodb_thread_concurrency=20
```

- 5 Save the configuration file.

- 6 Stop all NA services:

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

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

- 7 Restart MySQL:

- *Windows*: Open the Services control panel. In the list of services, right-click the MySQL service, and then click **Restart**.
- *UNIX*: Run the following command:

```
/etc/init.d/mysql restart
```

- 8 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 SWIM Server**
- **TrueControl Syslog Server**
- **TrueControl TFTP Server**

- *UNIX*: Run the following command:

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

Configuring Oracle for NA

This section describes known tuning of Oracle for NA.

Number of Database Connections from NA

The number of database connections is the total number of connections that NA can make to the database at any moment. This number depends primarily on the NA configuration for the maximum number of concurrent tasks.

If you change the maximum data source pool size, update the Oracle database configuration for the number of database connections.

Additionally, the following errors indicate the need to update the configuration for the number of database connections:

- This task did not complete. Connections could not be acquired from the underlying database!
- This task did not complete. An SQLException was provoked by the following failure: `com.mchange.v2.resourcepool.ResourcePoolException: A ResourcePool cannot acquire a new resource -- the factory or source appears to be down.`
- This task did not complete. Can't find CustomScript
Find failed: `java.sql.SQLException: Connections could not be acquired from the underlying database!`

For an Oracle database, the value of the `processes` parameter sets the number of database connections. The value of the `processes` parameter should be greater than or equal to the sum of the following factors:

- For *each* active NA core, the value of the maximum pool size attribute for the NA data source configuration
- For *each* active NA core, the value of the maximum pool size attribute for the NA report data source configuration
- For *each* active NA core, a buffer of 50



If the active NA cores in an NA Horizontal Scalability environment are configured identically, the calculation in this step is the same as multiplying the result of the calculation for one NA core by the number of active NA cores in the NA Horizontal Scalability environment.

According to the Oracle documentation, the values of the `sessions` and `transactions` parameters are relative to the value of the `processes` parameter. If the value of the `processes` parameter needs to be changed, the values of the `sessions` and `transactions` parameters should also be updated.

Size of the NA Tablespace

The following error suggests that the NA tablespace does not have sufficient space for its contents:

```
The system could not save the data for device id 50851 - An SQLException was provoked
by the following failure: com.mchange.v2.resourcepool.ResourcePoolException: A
ResourcePool cannot acquire a new resource -- the factory or source appears to be down.
Contact Technical Support. (Reference stack trace ID 1690)"
```

Report this message to the database administrator (DBA), and suggest that the DBA evaluate the free space of the NA tablespace.

Also see [Reclaiming Unused Space \(Oracle Database\)](#) on page 39.

Configuring SQL Server for NA

At this time, there is no recommended tuning for Microsoft SQL Server with NA.

3 Extending the Number of Custom Enhanced Fields

In the NA console, you can configure up to 31 custom data fields each for the Device Details page and the Device Interfaces page. These fields are available as follows:

- Six fields can be configured on the Admin > Custom Data Setup page.
- 25 fields can be configured on the Admin > Enhanced Custom Fields Setup page (when the Enable Enhanced Custom Fields check box is selected on the Admin > Administrative Settings > User Interface page).

To extend the available number of enhanced custom fields for the Device Details page, the Device Interfaces page, or both pages, follow these steps:

- 1 Back up the `adjustable_options.rcx` file to a location outside the `<NA_HOME>` directory.
- 2 In a text editor, such as Word or vi, edit the `adjustable_options.rcx` file as follows:
 - To extend the number of enhanced custom fields for the Device Details page, add the following line:

```
<option name=" metadata/field_limit/RN_DEVICE">100</option>
```

- To extend the number of enhanced custom fields for the Device Interfaces page, add the following line:

```
<option name=" metadata/field_limit/RN_DEVICE_PORT">100</option>
```



To restrict the number of available enhanced custom fields, replace 100 with a smaller value. (Specifying a larger value has the same effect as the leaving the value at 100.)

- 3 Save the `adjustable_options.rcx` file.
- 4 Reload the `.rcx` settings by doing one of the following:
 - Run the `reload server options` command from the NA proxy.
 - Restart the NA management engine.

4 Configuring the Default Setting of the Enforce Save Check Box for New Tasks

As of NA 9.11, for many NA device tasks, the Enforce Save task option specifies whether NA should overwrite the startup configuration with the current running configuration at the completion of the task. The setting applies to only those devices that support a startup configuration.

Per the additional information for FIX QCCR1B105606 in the patch readme, the default value of this setting is configurable per task type in the `adjustable_options.rcx` file. For each device task, the `adjustable_options.rcx` file contains an option in the following format:

```
<option name="DeviceInteraction/EnforceConfigurationSave/task_name">setting</option>
```

Possible values for *task_name* are:

- Take Snapshot
- Discover Driver
- Run ICMP Test
- Deploy Passwords
- Deploy Config
- Configure Syslog
- Run Command Script
- Run Diagnostics
- Synchronize Startup and Running
- Update Device Software
- Backup Device Software
- Reboot Device
- Run Device Script
- Delete ACLs
- VLAN Task
- Port Scan
- Add Device Context
- Remove Device Context
- OS Analysis
- Provision Device
- Batch Insert ACL Line
- Batch Remove ACL Line

Possible values for *setting* are:

- `true`—The Enforce Save field is visible for this task type and defaults to selected (overwrite the startup configuration). The user running the task can override the default setting by clearing the Enforce Save check box.
- `false`—The Enforce Save field is visible for this task type and defaults to cleared (do not change the startup configuration). The user running the task can override the default setting by selecting the Enforce Save check box.
- `disabled`—The Enforce Save field is not visible for this task type. The task will never attempt to overwrite the startup configuration with the running configuration.

To change the default setting of the Enforce Save check box for a specific device task type, follow these steps:

- 1 Back up the `adjustable_options.rcx` file to a location outside the `<NA_HOME>` directory.
- 2 In the `adjustable_options.rcx` file, locate the following line for the task that you want to change:

```
<option name="DeviceInteraction/EnforceConfigurationSave/task_name">setting</option>
```

- 3 In the `adjustable_options.rcx` file, edit the *setting* value.
- 4 Save the `adjustable_options.rcx` file.
- 5 Restart all NA services.
 - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Restart**:
 - **TrueControl ManagementEngine**
 - **TrueControl FTP Server**
 - **TrueControl SWIM Server**
 - **TrueControl Syslog Server**
 - **TrueControl TFTP Server**
 - *UNIX*: Run the following command:


```
/etc/init.d/truecontrol restart
```



The change take effect for new tasks only.

5 Limiting Device Password Access

As of NA 9.11, NA provides functionality for limiting some NA users to their AAA credentials when connecting to devices from NA. The affected users cannot use the device-specific passwords and the password rule passwords that are stored in NA.

Enabling the Access-Limiting Functionality

To enable the functionality for limiting some NA users to their AAA credentials when connecting to devices from NA, follow these steps:

- 1 Create the UseStoredDevicePasswords custom data field.
 - a In the NA console, open the Custom Data Setup page (**Admin > Custom Data Setup**).
 - b In the list, select **Users**.
 - c Under Users, locate an unused field configuration, and then select the check box to enable that field.
 - d For the enabled field, do the following:
 - Set API Name to: **UseStoredDevicePasswords**
 - Set Display Name to: **User can use device passwords stored in NA**
 - Select the **Limit to** check box, and then enter the limits as: **Yes, No**
 - e Click **Save**.
- 2 Force the use of AAA credentials for all task types.
 - a In the NA console, open the Administrative Settings – Device Access page (**Admin > Administrative Settings > Device Access**).
 - b Under Task Credentials, locate Allow User AAA Credentials, and then select all check boxes for the task types.
 - c Also under Task Credentials, locate Enable Password Option, and then select the Try only user's password for enable option.
 - d Click **Save**.
- 3 Enable the UseStoredDevicePasswords custom data field.
 - a Back up the `adjustable_options.rcx` file to a location outside the `<NA_HOME>` directory.
 - b In the `adjustable_options.rcx` file, add the following line:

```
<option name="Users/
CanUsePasswordsCustomField">UseStoredDevicePasswords</option>
```
 - c Save the `adjustable_options.rcx` file.

- d Reload the .rcx settings by restarting the NA services.
 - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Restart**:
 - TrueControl ManagementEngine**
 - TrueControl FTP Server**
 - TrueControl SWIM Server**
 - TrueControl Syslog Server**
 - TrueControl TFTP Server**
 - *UNIX*: Run the following command:

```
/etc/init.d/truecontrol restart
```

Using the Access-Limiting Functionality

To limit some NA users to their AAA credentials when connecting to devices from NA, follow these steps:

- 1 In the NA console, open the All Users page (**Admin > Users**).
- 2 In the row for the user to receive limited access, click **Edit**.
- 3 On the Edit User page, locate the User can use device passwords stored in NA setting, and then select **No**.
- 4 Click **Save**.



This change takes effect the next time this user logs on to the NA console.

For any device task this user runs from the NA console or the NA proxy, NA uses that user's AAA credentials when connecting to the target devices. In the NA console, the Device Credentials Options section of the task creation, editing, and scheduling pages includes only the pre-selected Use task owner's AAA credentials option for the affected user. The user cannot change this configuration.

6 Full Text Search of Configuration Text

As of version 9.10.01, HP Network Automation Software (NA) supports a “contains (full text)” search of Configuration Text. After full text search is enabled, faster configuration text search is available for the following report options:

- Reports > Search For > Devices > Configuration Text > contains (full text)
- Reports > Search For > Configurations > Configuration Text > contains (full text)
- Reports > Search For > Device Templates > Configuration Text > contains (full text)
- Reports > Advanced Search > Search Criteria > Configuration Text > contains (full text)

Additionally, you can create a dynamic group based on the results of a Search Criteria > Configuration Text > contains (full text) search.

Similarly, these searches also support searching for configuration text that “does not contain (full text).” The search is case insensitive for the “contains (full text)” and “does not contain (full text)” operators.

The “contains (full text)” search is an indexed search and requires that the database is enabled for full text search.

Because the “contains (full text)” search is indexed, it returns results faster than does the “contains” search. However, the “contains (full text)” search supports fewer options than does the “contains” search.

➤ This feature is not supported on MySQL.

This topic contains the following sections:

- [Enabling Full Text Search of Configuration Text](#) on page 24
- [Using the Contains \(Full Text\) Search Operator](#) on page 28
- [Disabling the Contains \(Full Text\) Search Operator](#) on page 29
- [Disabling Full Text Search](#) on page 30
- [Command Line Reference](#) on page 31

➤ The commands in this documentation are available as of NA 9.10.02.

Enabling Full Text Search of Configuration Text

The procedure for enabling full text search of NA configuration text depends on the NA environment.

Standalone NA Core Environment

In a standalone NA Core environment, do the following:

- 1 On the NA server, enable full text search of the database. Follow the steps appropriate to the database type:
 - [Enabling Full Text Search on Oracle](#) on page 25
 - [Enabling Full Text Search on Microsoft SQL Server](#) on page 27
- 2 On the NA server, enable the “contains (full text)” search operator as described in [Enabling the Contains \(Full Text\) Search Operator](#) on page 28.

Horizontal Scalability Environment

In a Horizontal Scalability environment, do the following:

- 1 On *one* NA server, enable full text search of the database. Follow the steps appropriate to the database type:
 - [Enabling Full Text Search on Oracle](#) on page 25
 - [Enabling Full Text Search on Microsoft SQL Server](#) on page 27
- 2 On *one* NA server, enable the “contains (full text)” search operator as described in [Enabling the Contains \(Full Text\) Search Operator](#) on page 28.

Use the `reload server options` command to reload the `.rcx` settings. This command pushes the change to the other NA servers in the Horizontal Scalability environment.

- ▶ Restarting the NA services does *not* push the change to the other NA servers in the Horizontal Scalability environment. If you use this approach to reload the `.rcx` settings, also modify the `.rcx` file on each of the other NA servers.



Use a similar approach for disabling full text search of NA configuration text.

Multimaster Distributed System Environment

In a Multimaster Distributed System environment, do the following:

- 1 On *each* NA server, enable full text search of the database.

Run the enablement procedures in parallel. That is, complete step 1 on each NA server before initiating step 2 on any NA server, and so forth.

Follow the steps appropriate to the database type:

- [Enabling Full Text Search on Oracle](#) on page 25
- [Enabling Full Text Search on Microsoft SQL Server](#) on page 27

- 2 On *one* NA server, enable the “contains (full text)” search operator as described in [Enabling the Contains \(Full Text\) Search Operator](#) on page 28.

Use the `reload server options` command to reload the `.rcx` settings. This command pushes the change to the other NA servers in the Multimaster Distributed System environment.

- ▶ Restarting the NA services does *not* push the change to the other NA servers in the Multimaster Distributed System environment. If you use this approach to reload the `.rcx` settings, also modify the `.rcx` file on each of the other NA servers.

- ▶ Use a similar approach for disabling full text search of NA configuration text.

Enabling Full Text Search on Oracle

Full-text search accesses an index of the text records in the database. The initial index generation requires available time and disk space.

- ▶ If Oracle Text is not yet enabled, also plan for database downtime.

NA maintains the full text index by incrementally indexing new configurations added during snapshot tasks and by removing the index entries of deleted configurations.

- ▶ Note the following:

- Because index generation is CPU-intensive, NA tasks might run slower than normal during the process of enabling full text search.
- Do not restart the NA services while index generation is in progress.

To enable full text search on an Oracle database, follow these steps:

- 1 Verify that Oracle Text is enabled and has the required privileges and space:
 - a Log on to the NA proxy with the credentials used to install NA.
 - b Run the following command:

```
run checkdb -resolver analyzesearchindex -verbose
```

- c Examine the output of the `analyzesearchindex` command.
 - If Oracle Text is not enabled, see [Analyze the NA Database in Preparation for Index Generation](#) on page 31 for suggested Oracle resources.
 - If Oracle Text is enabled, verify that the approximate additional space required for the index generation process is available on the database server.
- 2 *Optional* (available for 9.10.03 or later). Consider the approximate time required for the index generation process. The `analyze` command calculates time based on the use of a single thread. You can reduce this time by using multiple threads while generating the index. To configure NA to use multiple parallel threads while generating the full text index, make the following edit:
 - a Back up the `adjustable_options.rcx` file to a location outside the `<NA_HOME>` directory.
 - b In the `adjustable_options.rcx` file, add the following line:


```
<option name="fulltextsearch/parallel">T</option>
```


T is the number of parallel threads. Possible values range from 1 to one less than the number of database server cores.
 - c Save the `adjustable_options.rcx` file.
 - d Reload the `.rcx` settings by doing one of the following:
 - Run the `reload server options` command from the NA proxy.
 - Restart the NA services.
- 3 In the NA console, delay any Take Snapshot tasks that are scheduled to start before the end of the approximate time required for index generation to complete.
- 4 Generate the full text index:
 - a From the NA proxy, run the following command:


```
run checkdb -resolver addsearchindex -verbose
```
 - b Examine the output of the `addsearchindex` command.
 - The expected status is COMPLETE & VALID.
 - If the status is IN PROGRESS, wait for index generation to complete.
 - If the status is INVALID, remove the index with the `run checkdb -resolver removesearchindex -verbose` command, and then repeat [step a](#).



You can close the command prompt window during index generation. In this case, run the following command to determine the status of the index generation:

```
run checkdb -resolver statusindex -verbose
```

- 5 Edit the `adjustable_options.rcx` file as described in [Enabling the Contains \(Full Text\) Search Operator](#) on page 28.
 - 6 In the NA console, examine the status of recent Take Snapshot tasks. Rerun any that failed.
-  On an Oracle database, the log file contains the following error for any Take Snapshot tasks that were running during the generation of the full text index. You can ignore this error:

```
java.sql.SQLException: ORA-29861: domain index is marked LOADING/  
FAILED/UNUSABLE
```

Enabling Full Text Search on Microsoft SQL Server

Full-text search accesses an index of the text records in the database. The initial index generation requires available time and disk space.



If the SQL Server Full Text Search service is not installed and enabled, also plan for database downtime.

NA maintains the full text index by incrementally indexing new configurations added during snapshot tasks and by removing the index entries of deleted configurations.



Note the following:

- Because index generation is CPU-intensive, NA tasks might run slower than normal during the process of enabling full text search.
- Do not restart the NA services while index generation is in progress.

To enable full text search on a Microsoft SQL Server database, follow these steps:

1 Verify that the SQL Server Full Text Search service is installed and enabled by using one of the following approaches:

- Using the SQL Server configuration manager, determine whether the SQL Server Full Text Search service is installed and started.
- Using SQL Management Studio, run the following select statement:

```
SELECT fulltextserviceproperty('IsFulltextInstalled');
```

This statement returns 1 if the database is full-text enabled and 0 if not.

2 On SQL Server 2005, remove the SQL Server noise words as follows:

- a Change to the `$SQL_Server_Install_Path\Microsoft SQL Server\MSSQL.1\MSSQL\FTDATA\` directory.
- b Back up the `noiseENU.txt` file.
- c Delete all entries in the `noiseENU.txt` file to leave an empty file.

For more information about editing noise words, see the “Noise Words” topic in the MSDN library:

<http://msdn.microsoft.com/en-us/library/ms142551%28SQL.90%29.aspx>



On SQL Server 2008, by default no noise words are enabled.

3 Generate the full text index:

- a Log on to the NA proxy with the credentials used to install NA.
- b Run the following command:

```
run checkdb -resolver addsearchindex -verbose
```



On SQL Server, this command returns immediately and starts full-text indexing. Wait some time before you start using the new search. In the verbose output, verify that this run did not generate any SQL exceptions.

- 4 Determine the status of the index generation by running the following command:


```
run checkdb -resolver statusindex -verbose
```

 - The expected status is COMPLETE & VALID.
 - If the status is IN PROGRESS, wait for index generation to complete.
 - If the status is INVALID, remove the index with the `run checkdb -resolver removesearchindex -verbose` command, and then repeat [step 3](#).
- 5 Edit the `adjustable_options.rcx` file as described in [Enabling the Contains \(Full Text\) Search Operator](#) on page 28.

Enabling the Contains (Full Text) Search Operator

After the NA database is enabled for full text search, enable the “contains (full text)” search operator in the NA console as follows:

- 1 Back up the `adjustable_options.rcx` file to a location outside the `<NA_HOME>` directory.
- 2 In the `adjustable_options.rcx` file, add the following line:


```
<option name="fulltextsearch/enabled">true</option>
```
- 3 Save the `adjustable_options.rcx` file.
- 4 Reload the `.rcx` settings by doing one of the following:
 - Run the `reload server options` command from the NA proxy.
 - Restart the NA services.

Using the Contains (Full Text) Search Operator

The “contains (full text)” operator supports the following search types:

- Search for a single word (a sequence of characters containing no spaces).

Valid examples:

 - interface
 - telnet
 - snmp
- Search for a single word using the wildcard.
 - The asterisk character (*) is the only supported wildcard; this wildcard must be at the end of the search phrase.
 - Valid examples:
 - interf*
 - tel*
- Search for an IP address (same as searching for a single word).

The asterisk character (*) is the only supported wildcard; this wildcard must be at the end of the search phrase.

- Search for an IPv4 address is supported and treated as single word.
Valid examples:
 - 10.11.12.13 results in exact match
 - 10.11.12.* results in all addresses that start with 10.11.12
 - 10.11.* results in all addresses that start with 10.11
- Search for a phrase (a sequence of characters containing one or more white spaces):
Valid examples:
 - set vlan
 - set vpn name
- Search for a phrase using the wildcard.
 - The asterisk character (*) is the only supported wildcard; this wildcard must represent one or more complete words within of the search phrase.
 - The * must NOT be at the end of the search phrase.
 - Use white space on either side of the *. If white space is absent, NA treats the search as a word using the wildcard.
 - Valid examples:
 - set * name
 - telnet * table * settings



Logical operators (AND, OR, NOT, ACCUM, EQUIV) are not supported nor considered as part of search string.

Disabling the Contains (Full Text) Search Operator

To temporarily disable the “contains (full text)” search operator in the NA console, follow these steps:

- 1 If any dynamic groups are configured to use the “contains (full text)” or “does not contain (full text)” operator, edit or delete these dynamic group configurations.
- 2 Back up the `adjustable_options.rcx` file to a location outside the `<NA_HOME>` directory.
- 3 In the `adjustable_options.rcx` file, locate the following line:


```
<option name="fulltextsearch/enabled">true</option>
```
- 4 Edit the located line to set the `fulltextsearch/enabled` option to false:


```
<option name="fulltextsearch/enabled">false</option>
```

Alternatively, delete this option from the file.
- 5 Save the `adjustable_options.rcx` file.
- 6 Reload the `.rcx` settings by doing one of the following:
 - Run the `reload server options` command from the NA proxy.
 - Restart the NA services.

Disabling Full Text Search

To permanently disable the “contains (full text)” search operator in the NA console and to remove the full text index from the database, follow these steps:

- 1 If any dynamic groups are configured to use the “contains (full text)” or “does not contain (full text)” operator, edit or delete these dynamic group configurations.
- 2 Disable the feature in NA:
 - a Back up the `adjustable_options.rcx` file to a location outside of the `<NA_HOME>` directory.
 - b In the `adjustable_options.rcx` file, locate the following line:

```
<option name="fulltextsearch/enabled">true</option>
```
 - c Edit the located line to set the `fulltextsearch/enabled` option to false:

```
<option name="fulltextsearch/enabled">false</option>
```

Alternatively, delete this option from the file.
 - d Save the `adjustable_options.rcx` file.
 - e Reload the `.rcx` settings by doing one of the following:
 - Run the `reload server options` command from the NA proxy.
 - Restart the NA services.
- 3 Remove the full text index:
 - a Log on to the NA proxy with the credentials used to install NA.
 - b Run the following command:

```
run checkdb -resolver removesearchindex -verbose
```

Command Line Reference

This section describes the NA-provided commands available for working with a full text search index on an Oracle database. Run these commands from the NA proxy.

This content is the complete documentation for these commands. No command-line help is available for them.

Analyze the NA Database in Preparation for Index Generation

To analyze the NA database in preparation for full-text index generation, run the following command:

```
run checkdb -resolver analyzesearchindex -verbose
```

If Oracle Text is enabled and no full-text index is present, this command determines the total configuration text present in the NA database. It then outputs the following estimates related to full-text index generation:

- Approximate additional space required for the index generation process
- Approximate time required for index generation to complete

Ensure that the database enough space for index as estimated by this CLI command.

If Oracle Text is not enabled, engage the Oracle database administrator to change the configuration. For information about enabling Oracle Text, see “Administering Oracle Text” in the *Oracle Text Application Developer’s Guide*.



Another information source is the Oracle MetaLink document collection, for which you must have a MetaLink account with Oracle. Documents of interest include the following:

- 280713.1: Manual installation, deinstallation of Oracle Text 10gR1
- 979705.1: Manual installation, deinstallation of Oracle Text 10gR2
- 579601.1: Manual installation, deinstallation and verification of Oracle Text 11gR1
- 970473.1: Manual installation, deinstallation and verification of Oracle Text 11gR2

The index configuration process requires available disk space of 50% to 200% of the configuration text size. Actual space requirements depend on the database contents.

The index configuration process is resource-intensive. Actual time depends on database hardware and configuration as well as the volume of text to be indexed.

For more information, see “Frequently Asked Questions About Indexing Performance” in the *Oracle Text Application Developer’s Guide*.

Check the Status of the NA Database Index

To determine the status of the NA database full-text index, run the following command:

```
run checkdb -resolver statusindex -verbose
```

This command returns the current status of the full-text index for the NA database. Possible values are:

- **COMPLETE & VALID:** Indexing has completed. The index is valid.
- **IN PROGRESS:** Indexing is in progress.
- **NOT PRESENT:** No full-text index is present. Searches based on the “contains (full text)” or “does not contain (full text)” search operator will fail.
- **INVALID:** The full-text index is corrupted. Remove the index, and then add it again.

Generate an Index of the NA Database

To generate a full-text index (if no index exists) for the NA database, run the following command:

```
run checkdb -resolver addsearchindex -verbose
```

Follow the recommendations from the `run checkdb -resolver analyzesearchindex -verbose` command to understand the disk space and time required for this index generation process.

To track the status of the index generation, do one of the following:

- Run the `run checkdb -resolver statusindex -verbose` command.
- Watch the NA logs with the troubleshooting option `feature/proxy` set to debug.

Remove the NA Database Index

To remove the full-text index (if present) from the NA database, run the following command:

```
run checkdb -resolver removesearchindex -verbose
```

Searches based on the “contains (full text)” or “does not contain (full text)” search operator will fail.

Reindex the NA Database

If search queries are very slow, the full-text index might be fragmented. To regenerate the full-text index of the NA database, run the following command:

```
run checkdb -resolver reindex -verbose
```


7 Enabling Case-Insensitive Search (Oracle)

As of version 9.11, HP Network Automation Software (NA) supports case-insensitive searches of many objects in the NA database on Oracle. (The MySQL and Microsoft SQL Server database searches are already case-insensitive.)

This topic contains the following sections:

- [Affected Fields](#) on page 33
- [Enabling Case-Insensitive Search of an Oracle Database](#) on page 36
- [Disabling Case-Insensitive Search](#) on page 37

Affected Fields

When enabled, case-insensitive search is available for most text fields in the NA console, as described here. Additionally, the command-line interface is case-insensitive for device hostname.

Search Box

The IP or Hostname search box follows the case-sensitivity configuration.

Search Criteria

The Search Criteria field is available for the following functions:

- Defining a dynamic device group on the New Group and Edit Group pages.
- Defining a dynamic policy scope on the New Policy and Edit Policy pages.
- Creating a custom search on the Advanced Search page.

With an Oracle database, case-insensitive search is not available for the following fields:

- ACL Application
- ACL Configuration
- Comments
- Configuration Text with the contains and does not contain operators. (The contains (full text) and does not contain (full text) operators are always case-insensitive.)

All other fields follow the case-sensitivity configuration.

Device Selector

For the New Task and Rerun Task pages, the Filter box on the device selector follows the case-sensitivity configuration.

Reports

Table 2 lists the report fields that can be searched on a case-insensitive basis when the case-insensitive search feature is enabled.

Table 2 Case Sensitivity of Report Search Fields

Search Type	Case-Insensitive Fields	Case-Sensitive Fields	
Device	<ul style="list-style-type: none"> • Host Name • Device Vendor • Device Model • FQDN • Access Methods • Device Location • Serial Number • Asset Tag • Device Software Version • Device Firmware Version • Device Description • Password Rule • ACL ID • ACL Handle 	<ul style="list-style-type: none"> • ACL Type • Module Slot • Module Description • Module Model • Module Serial • Module Firmware Version • Module Hardware Revision • ROM Version • Service Type • Custom Service Type • VTP Domain Name • VTP Operating Mode 	<ul style="list-style-type: none"> • Comments • Configuration Text • ACL Configuration • ACL Application
Interface	<ul style="list-style-type: none"> • Port Name • Port Type • Port Status • Running Port State • Description • Configured Duplex • Configured Speed • Negotiated Duplex 	<ul style="list-style-type: none"> • Negotiated Speed • VLAN Name • Host Name • Module Slot • Module Description • Module Model • Module Serial • Module Firmware Version 	
Module	<ul style="list-style-type: none"> • Host Name • Module Slot • Module Description • Module Model 	<ul style="list-style-type: none"> • Module Serial • Module Firmware Version • Module Hardware Revision 	<ul style="list-style-type: none"> • Comments
Policy	<ul style="list-style-type: none"> • Policy Name • CVE 	<ul style="list-style-type: none"> • Vendor URL • Solution URL 	<ul style="list-style-type: none"> • Solution
Policy, Rule, and Compliance	<ul style="list-style-type: none"> • Host Name 	<ul style="list-style-type: none"> • CVE 	

Table 2 Case Sensitivity of Report Search Fields (cont'd)

Search Type	Case-Insensitive Fields	Case-Sensitive Fields
Configuration	<ul style="list-style-type: none"> • Host Name • Changed By 	<ul style="list-style-type: none"> • Comments • Configuration Text
Diagnostic	<ul style="list-style-type: none"> • Host Name 	<ul style="list-style-type: none"> • Diagnostic Text
Task	<ul style="list-style-type: none"> • Task Name • Host Name • Scheduled By 	<ul style="list-style-type: none"> • Comments • Result
Session	<ul style="list-style-type: none"> • Host Name • Created By 	<ul style="list-style-type: none"> • Session Data
Event	<ul style="list-style-type: none"> • Added By • Host Name 	<ul style="list-style-type: none"> • Description
User	<ul style="list-style-type: none"> • First Name • Last Name • User Name • Email Address • AAA User Name • Comments 	
ACL	<ul style="list-style-type: none"> • Host Name • ACL ID • ACL Handle • ACL Type • Changed By 	<ul style="list-style-type: none"> • ACL Configuration • ACL Application • Comments
MAC Address	<ul style="list-style-type: none"> • Host Name • Port Name • Port Description • VLAN 	
IP Address	<ul style="list-style-type: none"> • Host Name • Port Name • Port Description • VLAN • Associated MAC 	
VLAN	<ul style="list-style-type: none"> • Host Name • VLAN Name • VLAN Type • VLAN Description • Private VLAN 	
Device Template	<ul style="list-style-type: none"> • Template Name • Device Vendor • Device Model • Device Description 	<ul style="list-style-type: none"> • Comments • Configuration Text
Single Search	<ul style="list-style-type: none"> • Added By • Host Name • Description 	

Enabling Case-Insensitive Search of an Oracle Database

For an Oracle database, case-insensitive search accesses a case-insensitive index of the text records in the database for each field in the query.

In a Horizontal Scalability environment, enable case-insensitive searching on *one* NA server.

In a Multimaster Distributed System environment, enable case-insensitive searching on *each* NA server.

To enable case-insensitive search of NA with an Oracle database, follow these steps to generate the case-insensitive indexes:

- 1 Connect to the NA proxy with the credentials used to install NA.
- 2 Run the following command:

```
mod oracletcaseinsensitive -option enable
```
- 3 In a Horizontal Scalability environment, synchronize the database configuration by restarting NA on all other NA servers connected to the NA database:
 - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Restart**:
 - **TrueControl ManagementEngine**
 - **TrueControl FTP Server**
 - **TrueControl SWIM Server**
 - **TrueControl Syslog Server**
 - **TrueControl TFTP Server**
 - *UNIX*: Run the following command:

```
/etc/init.d/truecontrol restart
```

Disabling Case-Insensitive Search

In a Horizontal Scalability environment, disable case-insensitive searching on *one* NA server.

In a Multimaster Distributed System environment, disable case-insensitive searching on *each* NA server.

To permanently disable case-insensitive search of NA with an Oracle database and to remove the case-insensitive indexes from the database, follow these steps:

- 1 If any dynamic groups are configured with case-insensitive search criteria, edit or delete these dynamic group configurations.
- 2 If any policies are configured with case-insensitive search criteria, edit or delete these policy configurations.
- 3 Remove the case-insensitive indexes:
 - a Connect to the NA proxy with the credentials used to install NA.
 - b Run the following command:

```
mod oraclecaseinsensitive -option disable
```

- 4 In a Horizontal Scalability environment, synchronize the database configuration by restarting NA on all other NA servers connected to the NA database:
 - *Windows*: Open the **Services** control panel. In the list of services, right-click each of the following services, and then click **Restart**:
 - **TrueControl ManagementEngine**
 - **TrueControl FTP Server**
 - **TrueControl SWIM Server**
 - **TrueControl Syslog Server**
 - **TrueControl TFTP Server**
 - *UNIX*: Run the following command:

```
/etc/init.d/truecontrol restart
```


8 Reclaiming Unused Space (Oracle Database)

Database maintenance often involves deleting data chunks within a database table, which results in free space inside the table. New records added after this maintenance populates the free space inside the table first, so the new records can be spread across several physical locations within the table. This fragmentation degrades database performance by extending data access times.

HP Network Automation Software (NA) pruning tasks can cause database table fragmentation. This section identifies one way to defragment an Oracle database tablespace. This procedure can be performed while the database is online.



This documentation describes one approach to this database administration task. Read the prerequisites to determine whether this approach applies to your situation. For other approaches and more detailed information, see the documentation for your database type and version.

Tablespace defragmentation can be run against all tables in the NA schema. [Table 3](#) lists the NA database tables and the associated LOB columns that are most frequently affected by fragmentation.

Table 3 NA Database Tables Frequently Affected by Fragmentation

Table Name	Target LOB Columns
RN_DEVICE_ACCESS_LOG	<ul style="list-style-type: none"> • ChangeEventData • Comments
RN_DEVICE_DATA	<ul style="list-style-type: none"> • DataBlock • Comments
RN_DEVICE_TOPOLOGY_DATA	
RN_DIAGNOSTIC_DATA	<ul style="list-style-type: none"> • DataBlock • Comments
RN_EVENT	<ul style="list-style-type: none"> • EventText • EventData
RN_EVENT_MESSAGE	<ul style="list-style-type: none"> • MessageBody
RN_SCHEDULE_TASK	<ul style="list-style-type: none"> • Comments • Result • TaskData

To defragment an Oracle database tablespace, follow these steps:

- 1 Verify that the tablespace meets the following prerequisites:
 - The tablespace must be set with automatic segment space management (ASSM).
 - The disk space available to the redo log must be sufficiently large relative to the size of the tablespace.
- 2 Enter the SQL*Plus command-line interface as the SYSDBA user.
- 3 Use the Oracle Segment Advisor to determine whether defragmentation is needed. Either check the results of the Automatic Segment Advisor or run the Segment Advisor manually.

For more information, see “Using the Segment Advisor” in the *Oracle Database Administrator’s Guide*.

- 4 For each table that requires defragmentation, do the following:
 - a Enable row movement by running the following command:

```
ALTER TABLE <table_name> ROW MOVEMENT;
```
 - b Reclaim unused rows by running the following command:

```
ALTER TABLE <table_name> SHRINK SPACE;
```
 - c Reclaim unused LOB columns by running the following command:

```
ALTER TABLE <table_name> MODIFY LOB (<lob_column_name>) (SHRINK SPACE);
```



Alternatively, reclaim unused rows and columns with one command as follows:

```
ALTER TABLE <table_name> SHRINK SPACE CASCADE;
```

This CASCADE command replaces [step b](#) and [step c](#).

We appreciate your feedback!

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

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

Product name and version: NA 9.11

Document title: *NA Administration Guide, October 2012*

Feedback: