

SAS[®] High-Performance Computing Management Console 1.6

User's Guide

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2012. *SAS® High-Performance Computing Management Console 1.6: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® High-Performance Computing Management Console 1.6: User's Guide

Copyright © 2012, SAS Institute Inc., Cary, NC, USA

All rights reserved. Produced in the United States of America.

For a hard-copy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

For a web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

The scanning, uploading, and distribution of this book via the Internet or any other means without the permission of the publisher is illegal and punishable by law. Please purchase only authorized electronic editions and do not participate in or encourage electronic piracy of copyrighted materials. Your support of others' rights is appreciated.

U.S. Government Restricted Rights Notice: Use, duplication, or disclosure of this software and related documentation by the U.S. government is subject to the Agreement with SAS Institute and the restrictions set forth in FAR 52.227-19, Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

December 2012

SAS provides a complete selection of books and electronic products to help customers use SAS® software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit **support.sas.com/bookstore** or call 1-800-727-3228.

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are registered trademarks or trademarks of their respective companies.

Contents

<i>What's New In SAS High-Performance Computing Management Console</i>	<i>v</i>
<i>Accessibility Features of the SAS High-Performance Computing Management Console</i>	<i>ix</i>
Chapter 1 • Overview of the Console	1
About the Console	1
Accessing the Console	2
Specifying Console Preferences	2
Configuring Console Network Settings	5
Configuring the Default Language	6
Configuring Console Logging	8
Chapter 2 • Managing Console Users and Groups	11
Managing Console Users and Groups	11
Force Off Console Sessions	19
Chapter 3 • Viewing Console Logs	21
View Logs for an Active Session	21
Search Console Logs	22
Understanding Log Details	23
Chapter 4 • Managing Users and Groups	27
About Managing Users and Groups	27
Managing Users and Groups	28
Managing User Sessions	37
Perform an SSH Lockout	39
Chapter 5 • Managing Database Resources	43
Managing Active Queries and Connections	43
Managing Temporary Tables	47

Chapter 6 • Managing Cube Data Files	49
Understanding the Cube Cleanup Feature	49
Specifying the Directory to Manage	50
Searching for Cube Data Files	51
Deleting Cube Data Files	53
Chapter 7 • Managing CPU and Memory Resources	55
Understanding the CGroup Resource Management Feature	55
About Resource Allocation	56
CGroups Resource Management Page	57
Creating a Control Group	58
Editing and Deleting Control Groups	60
Example: Resource Management	61
CGroup Resource Management Administration	62
Chapter 8 • Gridhosts File Management	63
What Is the Gridhosts File?	63
Managing the Gridhosts File	64
Chapter 9 • Simultaneous Shell Utilities	67
Simultaneous Shell Command	67
Simultaneous Copy Command	68
Index	69

What's New

What's New In SAS High-Performance Computing Management Console

Overview

The SAS High-Performance Computing Management Console is used for managing high-performance computing environments. The console assists with managing operating system user IDs and groups. The console can also be used to perform administration of SSH lockouts, manage temporary tables, manage cubes, and assist with managing CPU and memory resources.

The SAS High-Performance Computing Management Console is enhanced to provide the following changes and enhancements:

- [support for managing CPU and memory resources](#)
- [support for additional languages](#)
- [enhancements to user management](#)

Support for Managing CPU and Memory Resources

The console is enhanced to support managing resource consumption by groups and users. This feature uses the Linux kernel control groups feature. You can use the console to create control groups and then allocate a CPU share percentage and a memory limit.

Support for Additional Languages

The console is enhanced to support internationalization for the following languages:

- English
- French
- German
- Italian
- Japanese
- Korean
- Polish
- Spanish
- Simplified Chinese
- Traditional Chinese

Enhancements to User Management

The following enhancements for system user management have been made in this release:

- When you create a user, you can choose from several shells, in addition to `/bin/ksh`. You can also change the shell for a user by editing the account with the console.
- When you create a user, you can create a new group with the same name and use it as the primary group. Alternatively, you can create a new group with a different name, or select from existing groups.
- You can edit a user and change both primary group membership as well as secondary group membership all at once.
- The following are propagated to each machine in the cluster when you create or edit a user:
 - passwords
 - primary and secondary group membership
 - real name (this is an optional field when you create a user)

Accessibility

Accessibility Features of the SAS High-Performance Computing Management Console

Overview

SAS High-Performance Computing Management Console has not been tested for compliance with U.S. Section 508 standards and W3C web content accessibility guidelines. If you have specific questions about the accessibility of SAS products, send them to accessibility@sas.com or call SAS Technical Support.

Documentation Format

Please contact accessibility@sas.com if you need this document in an alternative digital format.

Using Assistive Technology via a Remote Server

If you have questions about running SAS High-Performance Computing Management Console and operating assistive technology via a remote server such as Citrix, contact your assistive technology vendor.

User Interface Layout

The application window contains two sections:

- The top of the window contains a toolbar with links for **Console Management** and **HPC Management** on the left and a **Log out** link on the right side.
- The center of the window contains the workspace. The workspace contains links for the modules that apply to either console management or HPC management.

For information about using the console management features, see [“Managing Console Users and Groups” on page 11](#). For information about using the HPC management features, see [“Managing Users and Groups” on page 27](#).

Overview of the Console

<i>About the Console</i>	1
<i>Accessing the Console</i>	2
<i>Specifying Console Preferences</i>	2
Specifying User Interface Preferences	2
Changing Your Console Password	4
<i>Configuring Console Network Settings</i>	5
<i>Configuring the Default Language</i>	6
<i>Configuring Console Logging</i>	8

About the Console

The SAS High-Performance Computing Management Console is a web application that is used by system administrators to manage high-performance computing environments that use SAS software. Maintaining high-performance computing (HPC) environments is challenging because of the large number of machines that are used in the distributed computing environment. In order to ensure peak performance and reliability, whatever action is taken on one machine must be performed on all the machines. Tasks that might be simple on a single machine, such as adding a user or changing a system configuration option, become more challenging in a distributed computing environment. Most changes must be propagated to the other machines and this can be labor

intensive and time consuming. The console is used to ease the administration of the environment.

The console includes its own embedded web server. When administrators log on to the console, where authentication occurs is determined by the user name provided. If the user name matches an operating system user account, then the operating system performs authentication. If the user name does not match an operating system user account, then the console performs authentication. In either case, the console does not perform any interaction with the SAS Metadata Server.

Accessing the Console

You can access the console with either of the following methods:

- Follow a web link or enter the URL to access the console directly. For example, you might enter `http://hostname.example.com:10020`.
- For SAS Visual Analytics deployments, you can access the console by selecting **Tools ► SAS HPC Management Console** from SAS Visual Analytics Administrator.

Log on with a console user account. For more information about console user accounts, see [“Managing Console Users and Groups” on page 11](#).

Specifying Console Preferences

Specifying User Interface Preferences

To specify your preferences:

- 1 Click **Console Management** from the toolbar.
- 2 Click **Console Configuration** and then click **User Interface**.

- 3 The most important setting is to make sure that the **HPC Environment being managed** menu is set appropriately. This menu choice determines the features that are presented in the console.

The following table identifies the options that are available:

Table 1.1 SAS High-Performance Computing Management Console User Interface Options

Option	Description
Page background Normal text Table background Table header Link text	Specify the color to use for each of these options as the RGB hexadecimal code for a color.
Display titles as text	Specifies whether any title is displayed as text in the console. The default value is No .
Display user ID and hostname for non-framed themes	Specify the location of the client login and host name when non-framed themes are used. Use the menu to select a choice. The default value is At bottom of browser .
Hostname to display in the console	This option enables you to specify the how the host name for the SAS High-Performance Computing Management Console is displayed to clients. Use the menu to select a choice. The default value is Real hostname .
Prefix user ID to page titles Prefix host name to page titles	Specify whether the user name or host name of the connected user appears in the page title bar of the web browser.
Date format	Specify how the date is displayed in the console. Use the menu to select a choice. The default value is dd/mon/yyyy .

Option	Description
File chooser size User chooser size Multiple users chooser size Date selector size Module chooser size Multiple modules chooser size	Specify the size (in pixels) to use for the dialog box for each of these options. The default value is to use the minimum number of pixels that are needed to display the contents of the dialog box.
HPC Environment being managed	Specify the environment type. Note: For SAS Visual Analytics deployments, select a menu item that includes SAS LASR Analytic Server.
Console server host name in <code>/etc/gridhosts</code>	Specify the host name for the machine that is used for running the SAS High-Performance Computing Management Console. Enter the host name exactly as it appears in the <code>/etc/gridhosts</code> file.

4 Click **Save**.

Changing Your Console Password

Console users can also change their own passwords. This feature is for console users that do not use an operating system user ID to log on to the console.

To change your password:

- 1 Click **Console Management** from the toolbar.
- 2 Click **Language and Password**.
- 3 The following table identifies the options that are available:

Table 1.2 *Change Language and Password Options*

Option	Description
Management Console UI language	To specify a language, select Personal choice and then select a language.
Management Console login password	To change your password, select Set to and then specify your new password in the text field and the Re-enter password field. The default is Leave unchanged .

4 Click **Make Changes**.

The change is immediate and the console does not need to be restarted for the change to take effect.

Configuring Console Network Settings

By default, the server process for the SAS High-Performance Computing Management Console listens on port 10020 and on all IP addresses for the machine on which it is installed. However, the port and address settings are configurable.

Note: In a SAS Visual Analytics deployment, do not change the port number.

To specify the port and address settings:

- 1 Click **Console Management** from the toolbar.
- 2 Click **Console Configuration** and then click **Console Server Network Configuration**.
- 3 The following table identifies the options that are available:

Table 1.3 SAS High-Performance Computing Management Console Ports and Addresses Options

Option	Description
Listen on specified IP addresses and port numbers	Select the Only Address option and specify an IP address to restrict the console to listen for connections on the specified IP address only. The default value is Any Address . Specify a port number in the Port Number field. The default value is 10020.
Accept IPv6 connections	Specifies whether the console accepts IPv6 connections. The default value is No .
Listen for broadcasts on UDP port	Specify whether to listen for UDP broadcasts from other console installations. Available options are Don't Listen or a specified port. The default value is to listen on port 10020. This option is not used at this time.
Console server host name	Specify how the console server determines its host name. Available options are Work out from browser or specifying a host name. The default value is to determine it from the web browser.
Reverse-resolve client IP address	Specify whether the IP addresses of the client machines are resolved to host names. The default value is Yes .

4 Click **Save**.

Configuring the Default Language

To specify the default language for the console:

1 Click **Console Management** from the toolbar.

- Click **Console Configuration** and then click **Language**.
- The following table identifies the options that are available:

Table 1.4 Management Console Language Options

Option	Description
Language	Use the menu to select a default language for the console.
Character set for HTML pages	Specify the character set that the console uses for sending HTML to the web browser. Some sample values are ASCII, ISO-8859-1, and ISO-2022-JA. The default is Determined by language .
Use the language that is specified by the client Web browser?	<p>If you set this value to Yes, then the console changes the language if one of the following locales is detected from the browser:</p> <div> <div>en</div> <div>English</div> <div>it</div> <div>Italian</div> <div>fr</div> <div>French</div> <div>de</div> <div>German</div> <div>es</div> <div>Spanish</div> <div>pl</div> <div>Polish</div> <div>ja_JP</div> <div>Japanese</div> <div>ko_KR</div> <div>Korean</div> <div>zh_TW</div> <div>Traditional Chinese</div> <div>zh_CN</div> <div>Simplified Chinese</div> </div>

4 Click **Change Language**.

The change is immediate and the console does not need to be restarted for the change to take effect.

Configuring Console Logging

To configure logging settings:

- 1 Click **Console Management** from the toolbar.
- 2 Click **Console Configuration** and then click **Logging**.
- 3 The following table identifies the options that are available:

Table 1.5 SAS High-Performance Computing Management Console Logging Options

Option	Description
Logging Enabled	Select Disable logging or Enable logging . The default value is Enable logging .
Log resolved hostnames?	Specifies whether to log host names in the logs. The default value is No .
Use combined log format (including referrer and user agent)?	Specifies whether a combined log format is used. The default value is No .
Periodically clear log files?	Specify the interval to use for clearing log files. The default value is not to clear log files.
Users to log	Specify the user accounts to record in log reports. The default is Log actions by all users .
Modules to log	Specify the modules to log. The default is Log actions in all modules .

Option	Description
Include Console logins and logouts in actions log?	Specify whether logon to and logoff actions from the console are logged. The default value is No .
Log changes made to files by each action?	Specify whether changes that are made to files are logged. The default value is No .
Record all modified files before actions, for rollbacks?	Specify whether modified file changes are recorded for possible rollback. The default value is No .
Permissions for log files	Specify the mode permissions for the log files. The default setting is 600 and owned by root.
Add log to syslog?	Specify whether actions are logged to syslog in addition to log files. The default value is No .

4 Click **Save**.

2

Managing Console Users and Groups

<i>Managing Console Users and Groups</i>	11
Understanding Console Users and Groups	11
Create a Console User	12
Create a Console Group	16
Best Practices for Managing Console Users and Groups	19
<i>Force Off Console Sessions</i>	19

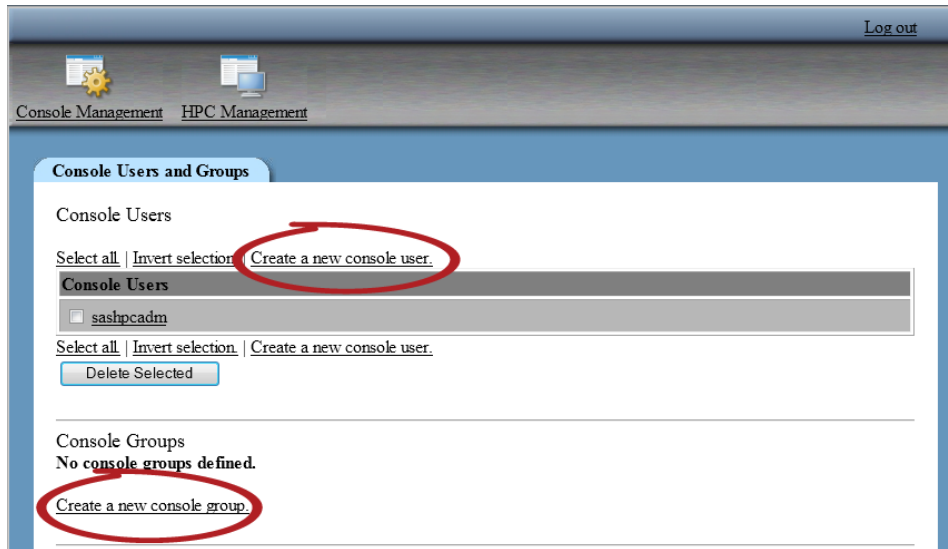
Managing Console Users and Groups

Understanding Console Users and Groups

An important concept to understand is that the SAS High-Performance Computing Management Console manages two types of users and groups. The first type is for console users and groups. This provides access control for the user accounts that can use the SAS High-Performance Computing Management Console. The second type of users and groups are the operating system users and groups. The console does assist with managing these users and groups. For information about managing operating system users and groups with the console, see [“Managing Users and Groups” on page 27](#).

The following display shows the interface for managing console users and groups. The links for creating new console users and new console groups are highlighted.

Display 2.1 Console Users and Groups Management



Create a Console User

Be aware that a console user is not the same as a user that is created in the **HPC Management** section of the console. Console users are created strictly for using the console. The Users and Groups module in the **HPC Management** section manages the operating system users and groups.

To create a console user:

- 1 Click **Console Management** from the toolbar.
- 2 Click **Console Users and Groups**.
- 3 Click **Create a new console user**.

The following display shows the Create console user screen.

Display 2.2 Create Console User

The screenshot shows the 'Create console user' web interface. At the top, there's a blue header bar with a 'Log out' link on the right. Below the header, there are navigation links: 'Console Management', 'HPC Management', and 'Module Index'. The main content area is titled 'Create console user'. It features a 'Create console user' tab and several expandable sections: 'Console user access rights' (expanded), 'User interface options', 'Security and limits options', and 'Available console modules'. The 'Console user access rights' section contains the following fields: 'Username' (sashpcadm), 'Password' (Set to .. with a dropdown menu and a checkbox for 'Force change at next login'), and 'Real name' (SAS HPC environment administrator). A 'Create' button is located at the bottom of the form. A 'Return to user list' link is at the bottom left of the page.

The following table describes the additional fields on the Create console user screen. They become available after clicking **Security and limits options** or **Available console modules**.

Table 2.1 Create Console User Field Descriptions

Field	Description
Username	Specify the user account.
Member of group	Use the menu to select a group.

Field	Description
Password	<p>Specify the password for the console user. The menu offers the following choices:</p> <p>Set to sets the password to the value entered into the text field.</p> <p>UNIX authentication passes the authentication request to the operating system. This is the most commonly used choice.</p> <p>No password accepted locks the user out of the system.</p>
Real name	(Optional) Specify an alternative description for the console user.

4 Click **Create**.

The follow table describes the additional options for creating a user.

Table 2.2 Create Console User Options

Field	Description
Language	Leave this option set to Default . English is the only supported language for this release.
Inactivity logout time	Specify an inactivity time-out. The default is no time-out.
Minimum password length	Specify the minimum number of characters for the user's password. The default is no minimum length.
IP access control	Select the option that meets the security requirements for the environment. If Deny from listed addresses is selected, then be aware that the console user can log on from any IP address that is not listed. The default value is Allow from all addresses .

Field	Description
Allowed days of the week	Specify the days of the week that the console user is permitted to log on. The default value is Every day .
Allowed times of the day	Specify the range of hours that the console user can log on. The default value is Any time .
Console Configuration	Select this check box to enable the console user to configure console settings such as the user interface, network settings, and console logging.
Console Users and Groups	Select this check box to enable the console user to configure console users and groups. A best practice is to limit access to this feature. For more information, see “Best Practices for Managing Console Users and Groups” on page 19 .
Console Logs	Select this check box to enable the console user to view and search the console logs. The console logs contain information about the actions performed with the console.
Active Queries and Connections	<p>Select this check box to enable the console user to view and terminate database queries and connections.</p> <p>Note: This option is available for environments that are configured to manage Greenplum Database deployments only.</p>
Cube Cleanup	<p>Select this check box to enable the console user to list and remove cube data files from the machines in the cluster.</p> <p>Note: This option is available for SAS High-Performance Risk environments only.</p>
Gridhosts File Management	Select this check box to enable the console user to modify the <code>/etc/gridhosts</code> file. This file is used to identify the machines in the HPC environment and is very important to the operation of SAS High-Performance Analytics software and the console itself.

Field	Description
Users and Groups	Select this check box to enable the console user to manage the operating system user accounts in the HPC environment.
CGroup Management	Select this check box to enable the console user to manage CPU and memory resources.
SSH Lockout	Select this check box to enable the console user to use the SSH lockout feature.
Temporary Table Maintenance	Select this check box to enable the console user to view and drop temporary database tables. Note: This option is available for environments that are configured to manage Greenplum Database deployments only.

Check boxes that are marked with an 'x' identify features that are enabled through group membership.

Create a Console Group

To create a console group:

- 1 Click **Console Management** from the toolbar.
- 2 Click **Console Users and Groups**.
- 3 Click **Create a new console group**.

The following display shows the Create Console Group screen.

Display 2.3 Create Console Group

Log out

Console Management HPC Management

Module Index

Create Console Group

↓ Console group access rights

Group name: consoleadmins

Group description: Administrators for the SAS High-Performance Computing Environ

↓ Available console modules

Select all | Invert selection

Console Management

☒ Console Configuration ☒ Console Logs

☒ Console Users and Groups ☒ Language and Password

HPC Management

☒ Cube Cleanup ☒ Gridhosts File Management

☒ SSH Lockout ☒ Users and Groups

Select all | Invert selection

Create

The following table describes the fields on the Create Console Group screen.

Table 2.3 Create Console Group Field Descriptions

Field	Description
Group name	Specifies the console group name.
Group description	Specifies a description for the console group.
Member of users and groups	When a console group that has members is edited, this field lists the members of the group.
Console Configuration	Select this check box to enable the members of the group to configure console settings such as the user interface, network settings, and console logging.

Field	Description
Console Users and Groups	Select this check box to enable members of the group to configure console users and groups. A best practice is to limit access to this feature. For more information, see “Best Practices for Managing Console Users and Groups” on page 19 .
Console Logs	Select this check box to enable members of the group to view and search the console logs. The console logs contain information about the actions performed with the console.
Active Queries and Connections	Select this check box to enable members of the group to view and terminate database queries and connections. Note: This option is available for environments that are configured to manage Greenplum Database deployments only.
Cube Cleanup	Select this check box to enable members of the group to list and remove cube data files from the machines in the cluster. Note: This option is available for SAS High-Performance Risk environments only.
Gridhosts File Management	Select this check box to enable members of the group to modify the <code>/etc/gridhosts</code> file. This file is used to identify the machines in the HPC environment and is very important to the operation of SAS High-Performance Analytics software and the console itself.
Users and Groups	Select this check box to enable members of the group to manage the operating system user accounts in the HPC environment.
SSH Lockout	Select this check box to enable members of the group to use the SSH lockout feature.

Field	Description
Temporary Table Maintenance	<p>Select this check box to enable members of the group to view and drop temporary database tables.</p> <p>Note: This option is available for environments that are configured to manage Greenplum Database deployments only.</p>

- 4 Click **Create**.

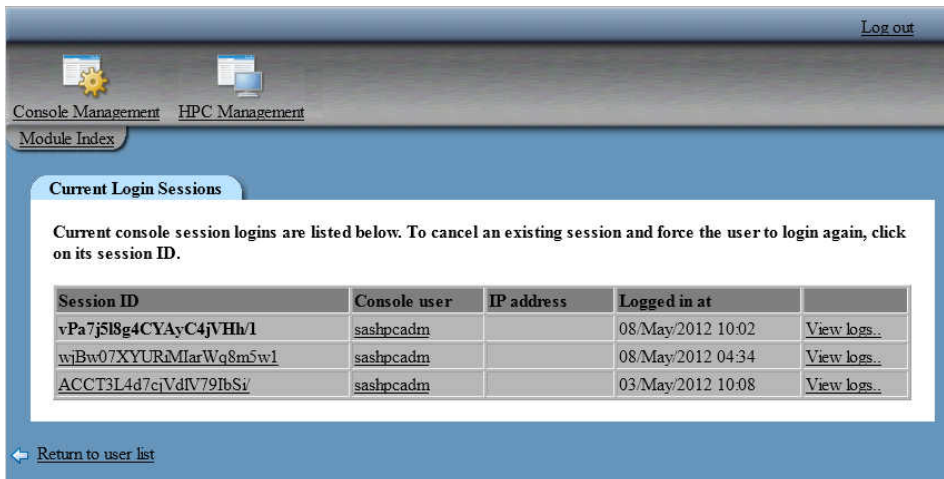
Best Practices for Managing Console Users and Groups

Console groups can be used to grant identical access rights to groups of users that share the same administration responsibilities. The **Console Users and Groups** module enables an administrator to manage the access permissions for console users. Limit access to this module because a user can change his access permissions and dilute the security of the SAS High-Performance Computing Management Console.

Force Off Console Sessions

The console enables administrators to force off other console users. To force off a console user:

- 1 Click **Console Management** from the toolbar.
- 2 Click **Console Users and Groups**.
- 3 Click **View Login Sessions** near the bottom of the page.
- 4 Select the session ID from the table to force off the console user. There is no confirmation prompt.

Display 2.4 SAS High-Performance Computing Management Console Current Login Sessions

The screenshot shows the SAS High-Performance Computing Management Console interface. At the top right is a "Log out" link. Below the header are navigation links: "Console Management" (with a gear icon), "HPC Management" (with a document icon), and "Module Index". The main content area is titled "Current Login Sessions". Below this title, a message states: "Current console session logins are listed below. To cancel an existing session and force the user to login again, click on its session ID." A table follows with four columns: "Session ID", "Console user", "IP address", and "Logged in at". The first column contains three session IDs, with the first one, **vPa7j5l8g4CYAyC4jVHh/1**, being bolded. The other two columns are empty for the first row. The second and third rows have values in the "Console user" and "Logged in at" columns. Each row has a "View logs..." link. At the bottom left is a "Return to user list" link with a left-pointing arrow.

Session ID	Console user	IP address	Logged in at	
vPa7j5l8g4CYAyC4jVHh/1	sashpcadm		08/May/2012 10:02	View logs...
wjBw07XYURaMlarWq8m5w1	sashpcadm		08/May/2012 04:34	View logs...
ACCT3L4d7cIVdIV79IbSi/	sashpcadm		03/May/2012 10:08	View logs...

Your session ID is shown in bold text.

Viewing Console Logs

<i>View Logs for an Active Session</i>	21
<i>Search Console Logs</i>	22
<i>Understanding Log Details</i>	23

View Logs for an Active Session

Follow the steps described in [“Force Off Console Sessions” on page 19](#). Instead of selecting the session ID, select **View logs** for the console user to monitor. The console uses the session ID to perform a search of the console logs. The following display shows an example of the search results:

Display 3.1 Console Logs by User and Session ID

Select the action to review and then review the details. For information about the details, see [“Understanding Log Details” on page 23](#).

Search Console Logs

To search console logs:

- 1 Click **Console Management** from the toolbar.
- 2 Click **Console Logs**.
- 3 Specify the search criteria in the **Search the console log for actions** fields and click **Search**.

Note: The **By non-console user** option does not work in a SAS Visual Analytics deployment.

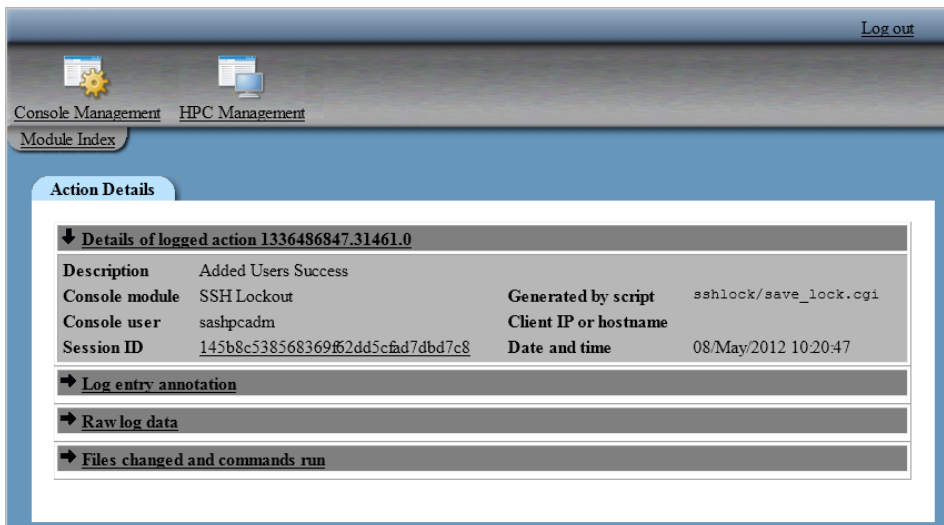
- 4 Select the action to review and then review the details. For information about the details, see [“Understanding Log Details” on page 23](#).

Understanding Log Details

The console logs record information such as the action that was performed, the module that was used, the date and time, and the parameters associated with the action. The information in these logs is valuable when reviewing the actions that were performed by a console user. In addition to viewing logs, the console enables an administrator to annotate the log record with comments.

The following display shows a console log example for starting an SSH lockout. The information that the console shows in most of the fields, and especially in the **Raw log data** section, differs according to which action is performed.

Display 3.2 Console Log Action Details



The **Details of logged action** section includes the following information:

Table 3.1 Console Log Details of Logged Action

Field	Description
Description	Specifies the action that was performed.

Field	Description
Console module	Specifies the console module that was used to perform the action.
Console user	Specifies the user name that performed the action.
Session ID	Specifies the web session that was used when the action was performed. You can select the link to view the console logs for all actions performed during the web session.
Generated by script	Specifies the console script that was used to perform the action.
Client IP or hostname	Specifies the IP address or host name that was used by the console user to perform the action.
Date and time	Specifies the date and time for when the action was performed. The date and time format can be changed by specifying a different format for displayed dates as described in “Specifying Console Preferences” on page 2 .

The **Log entry annotation** section enables an administrator to save a comment about the action. Click **Save** after adding a comment.

The information in the **Raw log data** section differs according to the action that was performed. The information is shown as a two-column table of parameter names and parameter values. Most of the parameters are related to the descriptions provided in the previous table. However, the appearance of some fields is different. The following list describes those fields:

sid

This is the value for the web session ID that was in use when the action was performed.

time

This value is a UNIX time representation for when the action was performed. When the data and time are shown in the **Details of logged action** section, the time is formatted for display from this value.

id

This value is the identifier for the log record.

The **Files changed and commands run** section shows the operating system commands that were executed to perform the action and the files that were used. This section is populated when the console is configured to log such information. By default, the console is not configured to log commands and files because the list of commands can be large and can contribute to a full system disk.

4

Managing Users and Groups

<i>About Managing Users and Groups</i>	27
<i>Managing Users and Groups</i>	28
Understanding Users and Groups	28
Configuring the Middle-Tier Shared Key	30
Create a User	31
About Editing and Deleting Users	34
SSH Key Management Features	35
Create a Group	35
About Editing and Deleting Groups	36
<i>Managing User Sessions</i>	37
Features in User Session Management	37
Review Recent SSH Logon Activity	37
Show Currently Logged On Users	38
<i>Perform an SSH Lockout</i>	39
Understanding the SSH Lockout Feature	39
Perform an SSH Lockout	41

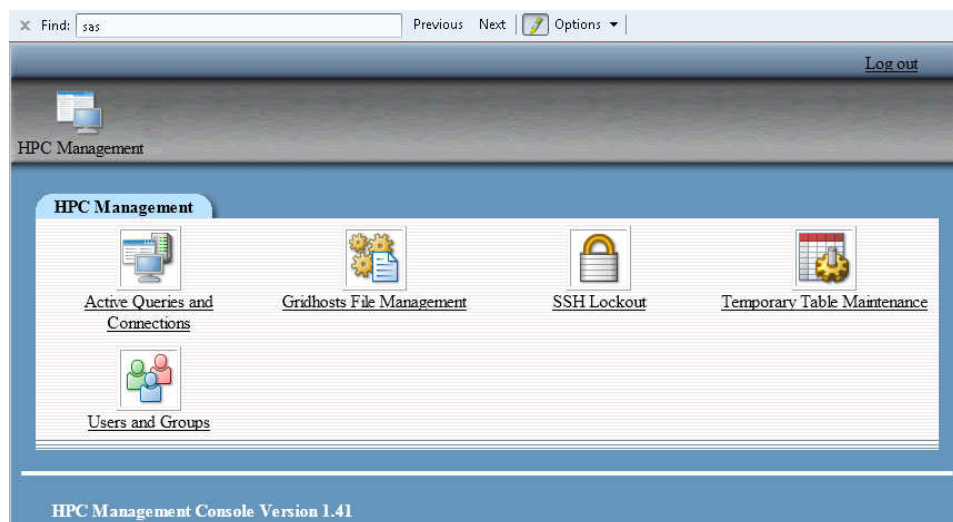
About Managing Users and Groups

The SAS High-Performance Computing Management Console enables two types of management for users and groups. The first type, for console users and groups, is described in [“Managing Console Users and Groups” on page 11](#). This section describes

the second type of user and group management, operating system users and groups. This management is performed with the High-Performance Computing (HPC) management section of the interface. It is also used to perform an SSH lockout. A lockout is used to limit access to machines in the cluster.

The following figure shows the interface for an administrator who has permission to perform HPC management only. The **Active Queries and Connections** and **Temporary Table Maintenance** modules apply to a SAS LASR Analytic Server with Greenplum environment only. The options that are used to perform console management are not available for this administrator.

Display 4.1 HPC Management Interface



Managing Users and Groups

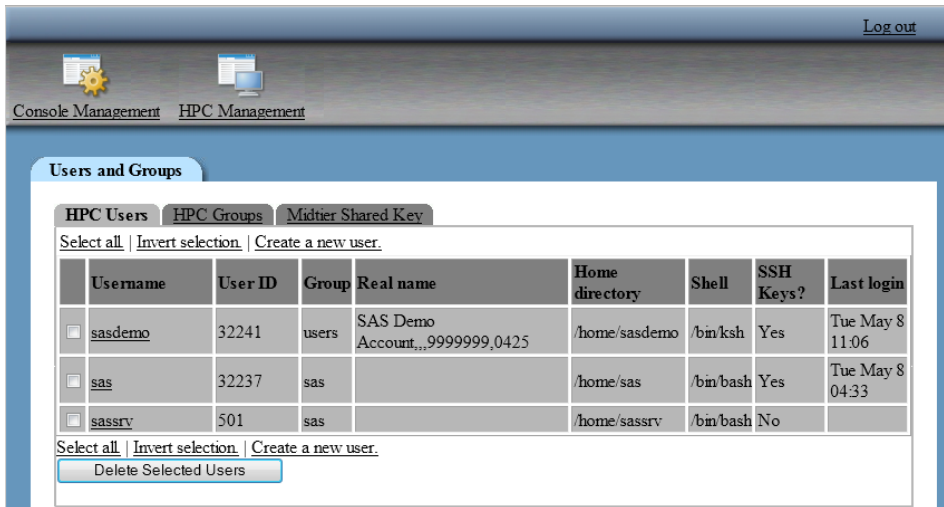
Understanding Users and Groups

The users and groups module is used to create, manage, and propagate operating system user accounts and groups throughout the machines in the cluster. It also enables an administrator to generate and distribute SSH keys for each user account.

The key generation can be performed as the account is propagated to the machines in the cluster, or it can be performed on existing accounts.

The following display shows the interface for managing operating system users and groups.

Display 4.2 Managing Users and Groups



The **HPC Users** tab shows the following information for each user account:

Table 4.1 User Accounts Field Descriptions

Field	Description
Username	Specifies the user account name. If the name is displayed in italic text, then the account has a null password. If the name is displayed in bold, then the account is disabled.
User ID	Specifies the UID for the user account.
Group	Specifies the primary group for the user account.
Real name	Specifies details about the user that make the user account more identifiable.

Field	Description
Home directory	Specifies the home directory for the user account.
Shell	Specifies the UNIX shell to use for the user account.
SSH Keys?	This field is set to Yes if the <code>\$HOME/.ssh/id_rsa</code> file exists. For environments that use Network File System (NFS) to manage home directories, this field is set to NFS . The console does not attempt to detect the existence of the <code>id_rsa</code> file.
Last login	Specifies the date of the last login.

Configuring the Middle-Tier Shared Key

Note: Configuring a middle-tier shared key for SAS Visual Analytics deployments is no longer necessary beginning with the 6.1 release.

You can use this feature to append an SSH public key to the `authorized_keys` files for user accounts. This task can be performed automatically when creating user accounts with the console. The public key can also be appended to the `authorized_keys` file for existing accounts.

To configure the middle-tier shared key:

- 1 Access the `.ssh/id_rsa.pub` file for the user account that is used to run JBoss. This file contains the SSH public key. Copy the contents of the file to your clipboard.
- 2 Click **HPC Management** from the toolbar.
- 3 Click **Users and Groups**.
- 4 Click **Midtier Shared Key** and specify values for the following fields:

Table 4.2 Middle-Tier Shared Key Field Descriptions

Field	Description
TKlasrkey location	Specify the fully qualified path to the tklasrkey file. The default location is <code>/opt/TKGrid/bin/tklasrkey</code> .
Shared Public Key	Paste the contents of the id_rsa.pub file from your clipboard.
Mid Tier Hostname	(Optional) If you specify the host name for the machine that is used to run JBoss, then the host name is included in the authorized_keys file. This provides an additional measure of security.

- 5 Click **Save**.

Create a User

To create a user:

- 1 Click **HPC Management** from the toolbar.
- 2 Click **Users and Groups**.
- 3 Click **Create a new user** and specify values for the following fields:

Table 4.3 Create User Field Descriptions

Field	Description
Username	Specify the user account name, such as <i>madupr</i> or <i>team1usr</i> .

Field	Description
User ID	<p>Select an option:</p> <p>Automatic the operating system selects an unused UID.</p> <p>Calculated (rarely used) the UID is created based on a Berkeley CRC and mkuid. The mkuid command assumes a standard naming convention for user names.</p> <p>Specified specify the UID to use. By default, this field shows the UID that is assigned if the Automatic option is used.</p> <p>The default option is Automatic.</p>
Real name	Specify details about the user that make the user account more identifiable.
Home directory	<p>Select an option:</p> <p>Automatic <code>/home/\$username</code> is the home directory value.</p> <p>Directory specify the fully qualified path to use as the home directory or click the Browse button to select a location.</p> <p>The default is Automatic.</p>
Shell	Select a shell from the menu. Alternatively, you can enter the path to the shell to use.
Password	<p>Select an option:</p> <p>No password required sets the password to null.</p> <p>Normal password enter the plain-text password in the field.</p> <p>Pre-encrypted password enter a password in encrypted form.</p>

Field	Description
Password changed	Specifies the last time the password was changed.
Expiry date	Specify the date on which the password should expire. You can enter the date or use the calendar. The default is no date.
Minimum days	Specifies the minimum number of days between password changes. The default is zero.
Maximum days	Specifies the maximum number of days between password changes. The default is 99999.
Warning days	Specifies the number of days that a password expiration warning is generated before a password expires. The default is zero.
Inactive days	Specifies the number of days that a user must be inactive before the user account is locked. The default is zero.
Force change at next login?	<p>Select Yes to force the user to change the password after the next logon. The default is No.</p> <p>Note: This option is not available for deployments that use SUSE Linux Enterprise Server. In this case, edit the user account after it is created to force the change.</p>
Primary group	You can select Existing group and then click ... to browse for an existing group. Alternatively, you can create a new group and enter the name with the New group option.
Propagate User	Select Yes to add the user to each machine in the environment. The default is No .

Field	Description
Generate and Propagate SSH Keys	Select Yes to generate SSH keys and propagate them when the user is created. The default is No .
Add Shared Midtier Key	Select Yes to include the information from the Midtier Shared Key tab in the authorized_keys file for the user. The default is No .
Create home directory?	Select Yes to create the user's home directory. Select No if the directory already exists. The default is Yes .
Copy template files to home directory?	Select Yes to copy standard environment files to the user's home directory at creation time. The default is Yes .

4 Click **Create**.

If **Propagate User** was set to **Yes**, then the progress of adding the user to the machines in the environment is shown.

If a mismatch is detected on a machine, such as a UID already in use, then the change fails and the mismatch is reported.

About Editing and Deleting Users

You can edit a user by selecting the user name on the **HPC Users** tab. The following list identifies the fields that can be changed:

- **Real name**
- **Shell**
- **Password**
- **Expiry date**
- **Minimum days**
- **Maximum days**

- **Warning days**
- **Inactive days**
- **Force change at next login?**
- **Primary group** and **Secondary groups**
- **Generate and Propagate SSH Keys**
- **Add Shared Midtier Key**

The field descriptions for these options are provided in [Table 4.3 on page 31](#).

You can also delete users individually or in groups by selecting the check box for the user on the **HPC Users** tab and then clicking **Delete Selected Users**.

SSH Key Management Features

As described in the section about creating users, the console can generate and propagate SSH keys. In addition, for SAS Visual Analytics deployments, the console can append a middle-tier shared key to the `authorized_keys` file for the user. These two features, SSH key generation and propagating the middle-tier shared key, can also be performed on existing user accounts.

To use these features, edit the user by selecting the user name on the **HPC Users** tab, select the radio button for the features that you want to use, and click **Save**.

Create a Group

To create a group:

- 1** Click **HPC Management** from the toolbar.
- 2** Click **Users and Groups**.
- 3** Select the **HPC Groups** tab and then click **Create a new group**. Specify values for the following fields:

Table 4.4 Create Group Field Descriptions

Field	Description
Group name	Specify the name for the group, such as <i>finance</i> or <i>team1</i> .
Group ID	<p>Select an option:</p> <p>Automatic the operating system selects an unused GID.</p> <p>Calculated (rarely used) the GID is created based on a Berkeley CRC and mkgid.</p> <p>Specified specify the GID to use. By default, this field shows the GID that is assigned if the Automatic option is used.</p> <p>The default option is Automatic.</p>
Members	Select users from the list and use the buttons to specify the members of the group.

4 Click **Create**.

If a mismatch is detected on a machine, such as a GID already in use, then the change fails and the mismatch is reported.

About Editing and Deleting Groups

You can edit a group by selecting the group name on the **HPC Groups** tab. Users can be added as members of the group. This action sets secondary group membership for the user account.

You cannot delete a group if it is used as the primary group for any user accounts.

Managing User Sessions

Features in User Session Management

The HPC management interface of the SAS High-Performance Computing Management Console enables administrators to monitor and manage user logon sessions. The following features are available to administrators:

- review recent SSH log on activity for one user or all users
- display a list of users currently logged on through SSH
- limit SSH logon actions by performing an SSH lockout

Review Recent SSH Logon Activity

To review recent SSH logon activity:

- 1 Click **HPC Management** from the toolbar.
- 2 Click **Users and Groups**.
- 3 Scroll to the bottom of the table of user accounts and then select either **All users** or **Only user**. If you select **Only user**, then enter the user account to review or click the button to select the user account from a list.
- 4 Click **Display Logins By**.

The recent logon activity is listed in the **Recorded Logins** table.

Display 4.3 Recent SSH Logon Activity



The screenshot shows a web-based management console. At the top right is a 'Log out' link. Below it are icons for 'Console Management' and 'HPC Management'. The 'Recorded Logins' section is highlighted with a blue header. Below this header, the text 'Recorded logins for sas' is displayed. A table follows with five columns: 'Login From', 'TTY', 'Login At', 'Logout At', and 'On For'. The table contains seven rows of logon data.

Login From	TTY	Login At	Logout At	On For
	pts/0	Tue May 8 04:33	Still logged in	
	pts/3	Wed May 2 10:23	10:23	00:00
	pts/3	Wed May 2 10:11	10:11	00:00
	pts/3	Wed May 2 09:53	09:53	00:00
	pts/3	Wed May 2 09:41	09:41	00:00
	pts/3	Wed May 2 09:40	09:40	00:00
	pts/0	Tue Apr 24 11:44	21:57	10:12

The **Login From** field is used to show the IP address or host name that was used for the logon session.

Show Currently Logged On Users

To show the users that are currently logged on with SSH:

- 1 Click **HPC Management** from the toolbar.
- 2 Click **Users and Groups**.
- 3 Scroll to the bottom of the table of user accounts and click **Show Logged In Users**.

The currently logged on users are listed **Logged In Users** table. You can select the user name to review the recent SSH logon activity for the user account.

Display 4.4 Logged In Users



Unix user	TTY	Logged in at	Logged in from
sas	pts/0	May 8 04:33	
sasdemo	pts/1	May 8 11:06	

[Return to users and groups list](#)

The **Logged in from** field is used to show the IP address or host name that was used for the logon session.

Perform an SSH Lockout

Understanding the SSH Lockout Feature

The SSH Lockout module enables an administrator to limit SSH logon access to the environment. This feature provides a mechanism for ensuring that processing resources are available for high-priority tasks as well as to assist with managing scheduled maintenance. Be aware that existing SSH logons are not affected by a lockout. A lockout prevents new SSH logons only. Access to software that does not rely on SSH is unaffected.

The feature provides user-level controls for denying SSH logons broadly while permitting SSH logons from specific user accounts. This level of control is available

based on how you set the lockout controls. The following figure shows the fields that are available on the **Lockout Whitelist** tab.

Display 4.5 SSH Lockout

The screenshot shows a web-based configuration interface for SSH lockout. At the top right is a "Log out" link. Below the navigation bar are two tabs: "Console Management" (active) and "HPC Management". The main content area is titled "Lockout Whitelist".

Lockout Status:
☐ Enabled ☒ Disabled

Admin Users:

sashpcadm
root

Regular Users:

sasdemo

Disabled Users:

frbeal

The following table provides a description for each of the fields on the **Lockout Whitelist** tab and describes how the settings interact with each other. The lockout status is not changed until the **Commit** button is clicked.

Table 4.5 Lockout Whitelist Field Descriptions

Field	Description
Lockout Status	<p>This field shows the current lockout status and is used to set the lockout status:</p> <p>Enabled only the user accounts listed in the Admin Users and Regular Users are permitted SSH logon access.</p> <p>Disabled all users, except the users listed in the Disabled Users are permitted SSH logon access.</p> <p>Note: User accounts in the Disabled Users are denied SSH logon access at all times.</p>
Admin Users	Specifies the user accounts that are permitted SSH logon access when the lockout is active.
Regular Users	Specifies the user accounts that are permitted SSH logon access when the lockout is active.
Disabled Users	Specifies the user accounts that are denied SSH logon access at all times. This list can be used to deny access to single users even when a lockout is not active.

Perform an SSH Lockout

To perform an SSH lockout:

- 1 Click **HPC Management** from the toolbar.
- 2 Click **SSH Lockout**.

- 3** Select **Enabled** on the **Lockout Status**.
- 4** Review the **Admin Users**, **Regular Users**, and **Disabled Users** lists.
For information about the lists, see [Table 4.5 on page 41](#).
- 5** Click **Commit**.

To disable an SSH lockout, set the Lockout Status to **Disabled** and then click **Commit**.

5

Managing Database Resources

<i>Managing Active Queries and Connections</i>	43
Understanding the Active Queries and Connection Feature	43
Configuring the Greenplum Database Connection	44
Restrict the Database User Account	45
Viewing Active Queries and Connections	46
<i>Managing Temporary Tables</i>	47
About Temporary Tables	47
Viewing Temporary Tables	47

Managing Active Queries and Connections

Understanding the Active Queries and Connection Feature

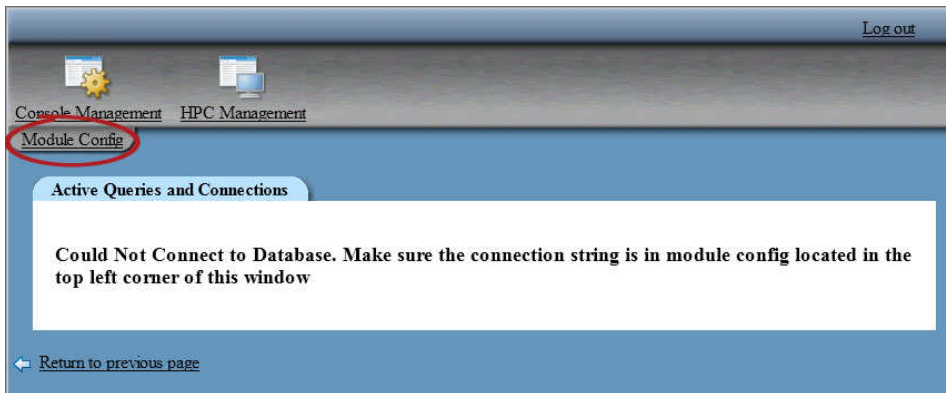
For deployments that use Greenplum Data Computing Appliance (DCA) as a co-located data provider, this feature enables an administrator to monitor long running queries and terminate database connections.

In order to monitor queries and terminate connections, a database user account is needed to connect to the database with administrative privileges. In order to list the active queries for all connections and to be able to terminate connections, the database user account should be granted the superuser role.

Configuring the Greenplum Database Connection

When you first access the **Active Queries and Connections** module, the console cannot connect to the Greenplum database because it is not configured with the connection parameters. The following display shows an example of the console before it is configured.

Display 5.1 Unconfigured Active Queries and Connections



To configure the module:

- 1 Click **HPC Management** from the toolbar.
- 2 Click **Active Queries and Connections**.
- 3 Select **Module Config**. (Shown in the preceding display.)
- 4 Specify the following parameters:

Table 5.1 *Active Queries and Connections Configuration*

Field	Description
Account Name for DB Connections	Enter the name of a database user account that has been granted the superuser role. For information about securing this user account, see “Restrict the Database User Account” .
Account Password	<p>This field is used for storing a password. This field is also used to avoid changing a stored password:</p> <p>Don’t change Select this option to avoid changing a stored password.</p> <p>Set to Select this option and specify the password for the database user account. The password is stored in <code>/opt/webmin/etc/actpsgqrys/config</code>.</p>
Database Name	Specifies the database name to use.
MDW Hostname	Specifies the host name for the Greenplum Database Master Host.

5 Click **Save**.

Restrict the Database User Account

The database user account that is used for the Greenplum Database connection must have superuser privileges in order to monitor queries and terminate connections. You can increase the security for the account by configuring Greenplum Database to accept connections for the account from the machine that hosts the console.

To restrict the database user account:

- 1 Edit the `$MASTER_DATA_DIRECTORY/pg_hba.conf` file.

2 Add a line to the file that is similar to the following example:

```
# permit the 'sashpcadm' role access to any database from
# IP address 192.168.1.1 only and use md5 encrypted password
host    all    sashpcadm    192.168.1.1/32    md5
```

3 Save and close the file.**4** Run `gpstop -u` so that Greenplum Database rereads the file and applies the change.

The same database user account can be used for the active queries and connections module and the temporary table maintenance module. For more information about editing the `pg_hba.conf` file, see *Greenplum Database Administrator Guide*.

Viewing Active Queries and Connections

Once the database connection information has been configured, the active queries and connections module shows the active query and connection information. Each row provides information about a database server process.

Display 5.2 Active Queries and Connections

The screenshot shows the 'Active Queries and Connections' module in the Greenplum Database web interface. The interface has a top navigation bar with 'Log out', 'Console Management', 'HPC Management', and 'Module Config'. The main content area displays a table with columns: PID, Username, Current State or Query, Last Query Started, and Client IP Address. Two rows are visible, both with state 'Idle'. Below the table are links for 'Select all' and 'Invert selection', and a 'Kill Selected' button.

PID	Username	Current State or Query	Last Query Started	Client IP Address
7478		Idle	2012-05-08 13:11:44.408076-04	
7467		Idle	2012-05-08 15:18:20.520483-04	

Select the check box for a server process and click **Kill Selected** to terminate the queries, connection, and server process. In the event that the process is not terminated, see Greenplum Database documentation for information about terminating processes.

Managing Temporary Tables

About Temporary Tables

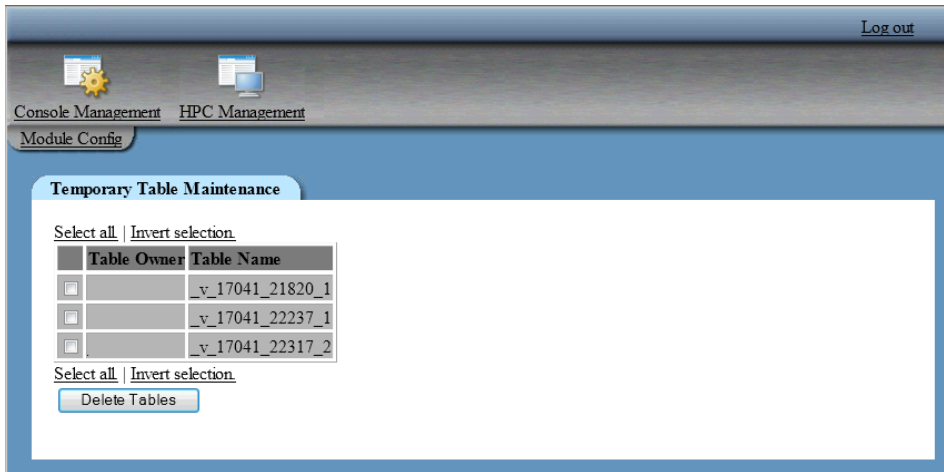
Temporary tables are a feature of Greenplum Database and several other third-party DBMS vendors. In some circumstances temporary tables are created by processes but are not dropped. The console provides the temporary table maintenance module for listing temporary tables and dropping them to restore system resources.

Before using the temporary table maintenance module, you must configure the Greenplum Database connection. Select **HPC Management ► Temporary Table Maintenance ► Module Config**. The configuration parameters are identical to those described in [“Configuring the Greenplum Database Connection” on page 44](#).

Viewing Temporary Tables

To view temporary tables:

- 1 Click **HPC Management** from the toolbar.
- 2 Click **Temporary Table Maintenance**.

Display 5.3 Temporary Table Maintenance

- 3 Select the check box for a table to drop, and click **Delete Tables** to drop the tables.

Note: It is possible to drop a table that is in use. Make sure that the table that you have selected to drop is not in use. Perform this action during scheduled maintenance or after using some other method to ensure that the table is not in use.

6

Managing Cube Data Files

<i>Understanding the Cube Cleanup Feature</i>	49
<i>Specifying the Directory to Manage</i>	50
<i>Searching for Cube Data Files</i>	51
<i>Deleting Cube Data Files</i>	53

Understanding the Cube Cleanup Feature

TIP The cube cleanup feature is available for deployments that use SAS High-Performance Risk only.

When the HPRISK procedure is used in a High-Performance Computing environment, the procedure creates data files for cubes on each machine in the cluster. The data files are created in the directory that is specified in the GRIDPATH= option.

The data files can be removed with the TASK=CLEAN option to the HPRISK procedure. However, there are some reasons that it might be necessary to use the console to manage cube data files. Some of these reasons are identified in the following list:

- The cube descriptor was deleted before using the HPRISK procedure with the TASK=CLEAN option.

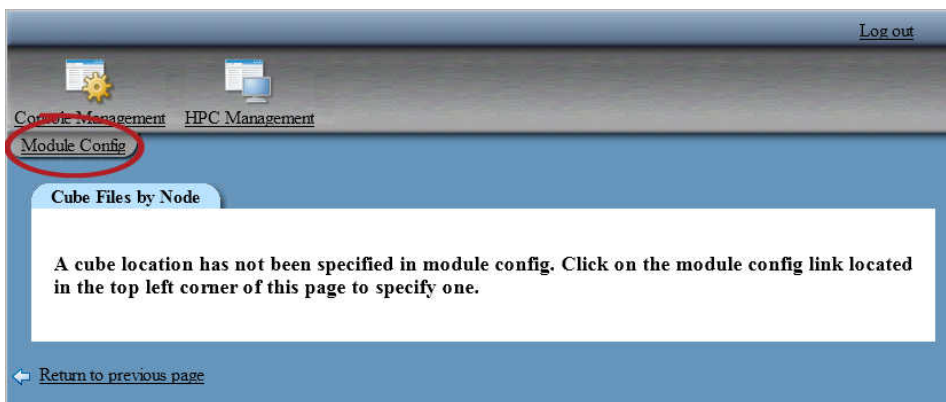
- The cube descriptor file is corrupted or, for some reason, is no longer readable by the HPRISK procedure.
- The cube descriptor file references host names for machines in the cluster that are no longer in service. This can be caused by an outage or renaming host names.

As indicated by the preceding list, there can be several reasons that data files are left on the machines in the cluster. The cube cleanup feature is used to list the data files so that you can delete them.

Specifying the Directory to Manage

When you first access the **Cube Cleanup** module, the console is not configured to monitor a directory for data files. The following display shows an example of the console before it is configured.

Display 6.1 Unconfigured Cube Cleanup



To configure the module:

- 1 Click **HPC Management** from the toolbar.
- 2 Click **Cube Cleanup**.
- 3 Select **Module Config**. (Shown in the preceding display.)

4 Specify the following parameters:

Field	Description
Location of Cubes	Specifies the path that is used in the GRIDPATH= option to the HPRISK procedure. If there are sub directories in the specified path, they are monitored as well.
Max Files per Page	Specifies the number of files to list in the console. If the number of files in the monitored directory exceeds this number, then the console provides a search window instead of listing the files.

5 Click **Save**.

Searching for Cube Data Files

To search for cube data files:

- 1 Click **HPC Management** from the toolbar.
- 2 Click **Cube Cleanup**.
- 3 If files are listed, click **Search**. If the monitored directory has more files than the **Max Files per Page** setting, then you are automatically presented with the search options.

Field	Description
Display Files By	<p>Use the menu to select a value:</p> <p>Owner Select this option to search for files by the user ID that owns the data file.</p> <p>Node Hostname Select this option to search for files by the host name.</p> <p>Cube File Name Select this option to search for files by the name of the data file.</p>
Which	<p>Use the menu to select either Contains or Equals.</p>
This Text	<p>Specifies the user ID, host name, or filename to use for the search. If Which is set to Equals, then the search criteria must be an exact match. If Which is set to Contains, then the specified text is used to match substrings.</p>
Or Files That Are	<p>The search for data files always includes date-based criteria. Use the menu to specify whether the search includes files that are Older or Newer than the specified date.</p>
Than This Date	<p>Specifies the date to use for the date-based search criteria.</p>

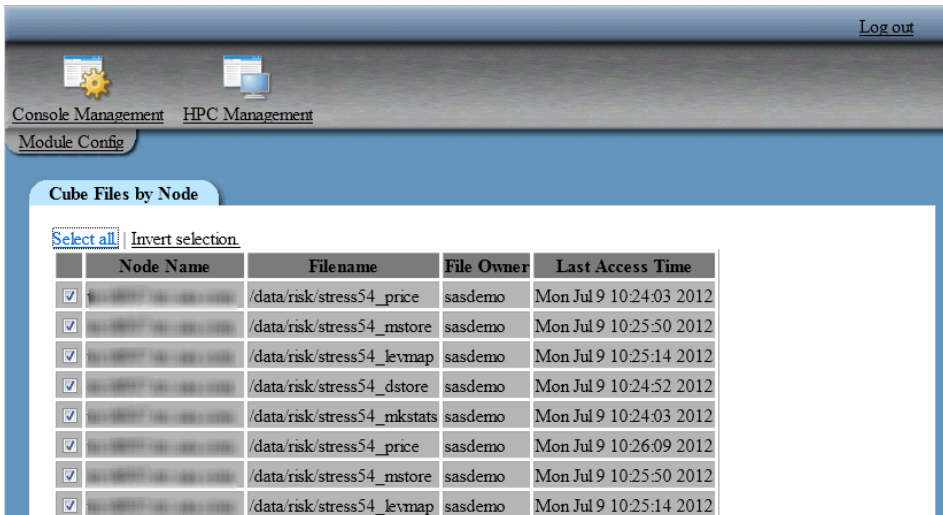
4 Click **Find**.

Deleting Cube Data Files

Once you have specified the directory to monitor for cube data files, you can delete files that are no longer needed. To delete the cube data files:

- 1 Click **HPC Management** from the toolbar.
- 2 Click **Cube Cleanup**.
- 3 Select the check box beside the cube data file to delete.

Display 6.2 Cube Files to Delete



The screenshot shows the HPC Management interface. At the top, there are icons for Console Management and HPC Management, and a 'Log out' link. Below these are tabs for 'Module Config' and 'Cube Files by Node'. The 'Cube Files by Node' tab is active, displaying a table of files to be deleted. The table has columns for 'Node Name', 'Filename', 'File Owner', and 'Last Access Time'. There are 8 rows of data, each with a checked checkbox in the first column. Above the table, there are links for 'Select all' and 'Invert selection'.

	Node Name	Filename	File Owner	Last Access Time
<input checked="" type="checkbox"/>	node1	/data/risk/stress54_price	sasdemo	Mon Jul 9 10:24:03 2012
<input checked="" type="checkbox"/>	node1	/data/risk/stress54_mstore	sasdemo	Mon Jul 9 10:25:50 2012
<input checked="" type="checkbox"/>	node1	/data/risk/stress54_levmmap	sasdemo	Mon Jul 9 10:25:14 2012
<input checked="" type="checkbox"/>	node1	/data/risk/stress54_dstore	sasdemo	Mon Jul 9 10:24:52 2012
<input checked="" type="checkbox"/>	node1	/data/risk/stress54_mkstats	sasdemo	Mon Jul 9 10:24:03 2012
<input checked="" type="checkbox"/>	node1	/data/risk/stress54_price	sasdemo	Mon Jul 9 10:26:09 2012
<input checked="" type="checkbox"/>	node1	/data/risk/stress54_mstore	sasdemo	Mon Jul 9 10:25:50 2012
<input checked="" type="checkbox"/>	node1	/data/risk/stress54_levmmap	sasdemo	Mon Jul 9 10:25:14 2012

- 4 Click **Delete**. Scroll to the bottom of the page if the button is not visible.

You can also use the **Select all** and **Invert selection** links to control which check boxes are selected. You can select the table heading, such as **Last Access Time**, to sort the data files. Click the same heading again to reverse the sort order.

7

Managing CPU and Memory Resources

- Understanding the CGroup Resource Management Feature* 55
- About Resource Allocation* 56
 - About CPU Allocation 56
 - About Memory Allocation 56
- CGroups Resource Management Page* 57
- Creating a Control Group* 58
- Editing and Deleting Control Groups* 60
- Example: Resource Management* 61
- CGroup Resource Management Administration* 62
 - CGroup System Services 62
 - Administrative Functions 62

Understanding the CGroup Resource Management Feature

You can use the CGroup Resource Management feature to manage CPU and memory consumption for users and groups. Resource management is performed through configuring kernel control groups (cgroups). The Linux kernel used for the SAS High-Performance Analytics Environment supports for cgroups. The Linux kernel supports

managing several subsystems to limit, prioritize, and account for resources. The SAS High-Performance Computing Management Console supports managing the CPU and memory subsystems because these are the most contentious resources in a deployment that uses the SAS High-Performance Analytics Environment.

In order to use the CGroup Resource Management feature, the operating system must include support for it through the Linux kernel version and the CGroups libraries. The Linux kernels used in Red Hat Enterprise Linux 6 and SuSE Linux Enterprise Server 11 contain support for CGroups. Older versions from both vendors do not include support for CGroups.

About Resource Allocation

About CPU Allocation

When you add a control group, you specify a share of the CPU, as a percentage, that the control group can use. When you view the **CGroup Resource Management** page, the allocated, available, and maximum percentages are listed in the **CGroup Resource Allocation** table.

Be aware that the allocation of CPU resources is for processes that are managed with control groups. Some CPU resources are consumed by system processes and processes that are started by users that are not part of control groups. In general, the percentages that are specified should be very close to actual CPU utilization when the system reaches full load.

About Memory Allocation

When you add a control group, you specify a maximum memory limit (in kilobytes) for the control group. When a process within the control group reaches that limit, any further memory is allocated from swap. Using swap as memory reduces performance significantly.

As with the CPU allocation, some RAM in the system is used by system processes and processes that are started by users that are not part of control groups.

CGroups Resource Management Page

The following display shows the CGroup Resource Management page.

Display 7.1 CGroup Resource Management

CGroup Resource Management

CGroup System Services [view all services](#)

cgred: 4 Started, 0 Stopped
 cgconfig: 4 Started, 0 Stopped

Configured Cgroups [add new cgroup](#)

Select all | [Invert selection](#)

	CGroup Name	CGroup Description	Subsystems Used
<input type="checkbox"/>	highpriority	High priority users	cpu,memory
<input type="checkbox"/>	lowpriority	Lower priority users	cpu,memory
<input type="checkbox"/>	midpriority	Middle priority user	cpu,memory

Select all | [Invert selection](#)
[Delete](#)

CGroup Resource Allocation

	Allocated	Available	Maximum
CPU Share Percent	100	0	100
Memory in KB	7500000	393420	7893420

Administrative Functions

[Compare Config Files](#) [Copy Config Files](#) [Restart All Services](#)

The operations that can be performed at the top of the page are described in “CGroup System Services”.

The **Configured CGroups** section is used to manage the control groups, including [“Creating a Control Group”](#).

The **CGroup Resource Allocation** table lists the resources that are allocated to control groups. For more information about the resources, see [“About Resource Allocation”](#).

The [“Administrative Functions”](#) operations are available from the bottom of the page.

Creating a Control Group

To create a control group:

- 1 Click **HPC Management** from the toolbar.
- 2 In the **Configured CGroups** section, click **add new cgroup**.

Log Off

Console Management HPC Management

Add New CGroup

Group Details

Group Name: salesgrp

Group Description: Analysts for sales

Subsystem and Resource Options

CPU Share Percentage: 35 (80 Percent Available)

Memory Limit Per Node (in KB): 5000 (5000000 KB Available)

Resource Managed User Options

Select Resource Managed Users

All Users

Managed Users

Task Admin User

Select Single Task Admin User

All Users

Task Admin User

Save

[Return to CGroup Resource Management](#)

3 Specify the following parameters:

Field	Description
Group Name	Specify a name for the control group such as salesgrp .
Group Description	Specify a description.
CPU Share Percentage	Specify the share of CPU utilization, as percentage, to allocate to the control group.

Field	Description
Memory Limit Per Node (In KB)	Specify the maximum memory limit for the control group.
Select Resource Managed Users	Specify the user IDs to manage in the control group.
Select Single Task Admin User	Specify the user ID that can add a process to the control group task list. For more information about the task admin user, see the information that follows this procedure.

4 Click **Save**.

The configuration files are updated, copied to the machines in the cluster, and the control group configuration start-up (cgconfig) service and the control group rules engine daemon (cgroupd) service are restarted.

The task admin user is more important for Linux environments that manage control groups manually. In a SAS High-Performance Analytics Environment, the choice of a task admin user is less critical. For example, in a SAS deployment, any process that is started by a resource managed user in a control group is subject to the resource allocations. Therefore, the ability for the task admin user to add unmanaged processes to a task file so that the process can become a managed task is not critical in a SAS deployment.

Editing and Deleting Control Groups

To edit a control group, select the group name from the **CGroup Name** column. Specify the same parameters that are used for creating a control group.

To delete a control group, select the check box beside the group name and click **Delete**.

When you change the share percentage, memory limit, or delete a group, the **CGroup Resource Allocation** table is updated to show the resources that remain available.

Example: Resource Management

The memory limit that is specified for a group is firm. Any further requests for memory are allocated from swap, and performance is likely to suffer. However, CPU allocation is performed slightly differently. CPU resources are granted freely until the system is under load and there is contention for CPU time. Consider the following table that specifies three sample control groups and CPU share percentages.

Table 7.1 Sample CPU Share Percentage Allocation

Control Group Name	CPU Share Percentage
High priority	60
Medium priority	30
Low priority	10

To understand how CPU utilization fluctuates until resource contention emerges, consider the following sequence of events:

- 1 A user in the low priority control group starts a process. Because there is no demand for CPU resources from any other user, the process receives 100% of the CPU time.
- 2 A user in the medium priority control group starts a process. The two users share CPU resources.
- 3 A user in the high priority control group starts a process. The contention for CPU resources causes the allocations to be enforced. The medium priority user receives 30% and the low priority user receives 10% until the process for the user in the high priority group ends or does not require CPU resources.

CGroup Resource Management Administration

CGroup System Services

The **CGroup System Services** section of the page shows the number of processes and status for the cgroupd and the cgconfig services.

You can click **view all services** to view a table of host names, services, and service statuses.

Administrative Functions

The most commonly used administrative functions are listed as buttons at the bottom of the page. The buttons are as follows:

Compare Config Files

Click this button to ensure that the configuration files for the control groups are the same on each machine in the cluster.

Copy Config Files

Click this button to copy the configuration files from the machine that is running the console to the other machines in the cluster.

Restart All Services

Click this button to restart the cgroupd and cgconfig services on each machine in the cluster.

8

Gridhosts File Management

<i>What Is the Gridhosts File?</i>	63
<i>Managing the Gridhosts File</i>	64
Creating the File	64
Editing Host Names in the Gridhosts File	65

What Is the Gridhosts File?

Identifying all the machines in the environment is critical to ensuring that the console can perform operations in parallel or serially on the machines in the cluster. The `/etc/gridhosts` file is used to identify the machines. It is a list of the machines, with one host name on each line. The `/etc/gridhosts` file is not the same as the `/etc/hosts` file.

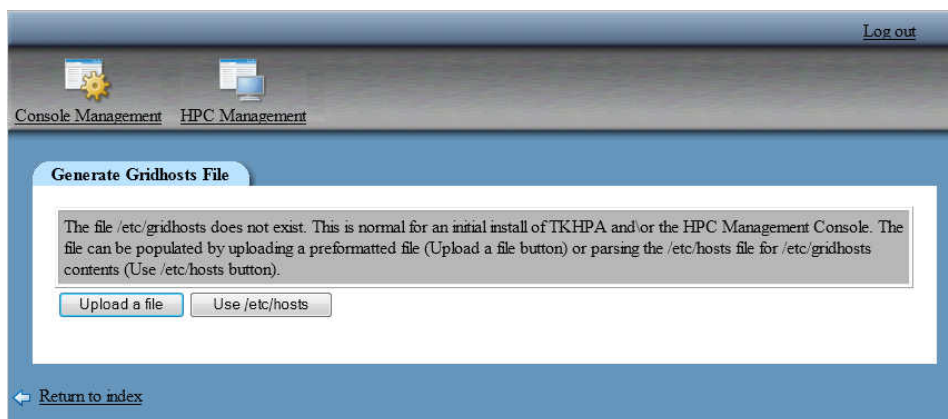
The `/etc/gridhosts` file is created immediately after installing SAS High-Performance Computing Management Console. If the `/etc/gridhosts` file has not been created already, then follow the steps described in [“Creating the File” on page 64](#). The console is also used to manage the host names that are listed in the file. The following sections provide more detail.

Managing the Gridhosts File

Creating the File

The contents of the file (host names) are typically created in a text editor or with shell scripting to avoid repetitive typing. If the file is not present on the machine that is hosting the console, then the following page is seen after accessing the **Gridhosts File Management** module of the console.

Display 8.1 Generate Gridhosts File



Note: The **Use /etc/hosts** button cannot be used in a SAS Visual Analytics deployment to create the `/etc/gridhosts` file.

If you can access the file, then you can upload it to the machine that is hosting the console. The file must contain one host name on each line. The host name must be the host name of the network interface to use for communication between the machines in the grid.

To upload a file:

- 1 Click **Upload a file**.

- 2 Click **Browse** and then navigate to the file that contains the host name of the machines in the cluster.

Click **Upload**.

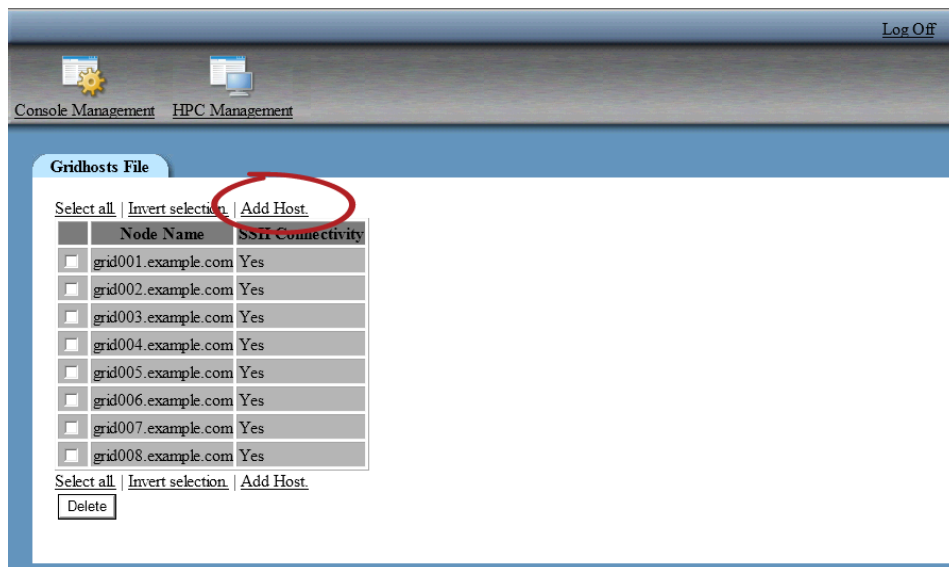
- 3 Review the host names that are listed in the table. If the list of host names is accurate, click **Select all** at the top of the table. Click **Submit**.

After the **Submit** button is clicked, the `/etc/gridhosts` file is created on the machine that is hosting the console.

Editing Host Names in the Gridhosts File

You can use the console to add and delete host names from the `/etc/gridhosts` file. The following display shows the link for adding host names and the button for deleting them.

Display 8.2 Gridhosts File



To add a host name:

- 1 Click **Add Host**.
- 2 Enter the host name in the **Hostname** field and click **Add**.

Note: If the machine that you specify is running and passwordless SSH is configured for the root user, then the **SSH Connectivity** field indicates **Yes**.

To delete a host name:

- 1** Select the check box beside the host name to remove.
- 2** Click **Delete**.

This action prevents the console from performing configuration operations on the machine. If the SAS High-Performance Computing Environment software or SAS High-Performance Deployment of Hadoop is configured, then those products need to be reconfigured to remove references to the machine too.

9

Simultaneous Shell Utilities

Simultaneous Shell Command 67

Simultaneous Copy Command 68

Simultaneous Shell Command

The simultaneous shell command, **simsh**, is installed with the console. The default location is `/opt/webmin/utilbin/simsh`. The **simsh** command relies on passwordless SSH and the `/etc/gridhosts` file. When you use the **simsh** command, passwordless SSH is used to log on to each machine identified in the `/etc/gridhosts` file and run the command line that follows the **simsh** command. The following example counts the number of files in `/tmp` on each machine in the cluster:

```
/opt/webmin/utilbin/simsh "ls /tmp | wc -l"
```

```
grid001.example.com: 6
grid002.example.com: 7
grid003.example.com: 12
...
grid016.example.com: 6
```

By default, the **simsh** command returns the host name, a colon, and the output of the command that ran on the machine. Some commands do not provide output, such as the **mkdir** command. In order to determine whether the **mkdir** command succeeds, you must check the return code. The following example uses the `-r` option to report the return code for the **mkdir** command:

```
/opt/webmin/utilbin/simsh -r "mkdir -p /tmp/sas/somedir"
```

```
grid001.example.com:0
grid002.example.com:0
grid003.example.com:0
...
grid016.example.com:0
```

In the preceding example, the return code follows the colon. You can also use the `-r` option with commands that do return output. The following example uses the `grep` command to determine whether `grp1` is listed in `/etc/group` on every machine in the cluster:

```
/opt/webmin/utilbin/simsh -r "grep '^grp1:' /etc/group"
```

```
grid001.example.com:0:grp1:x:32242:jemcki,madupr
grid002.example.com:0:grp1:x:32242:jemcki,madupr
grid003.example.com:1
...
grid016.example.com:0:grp1:x:32242:jemcki,madupr
```

In the preceding example, the return code for machine `grid003.example.com` is 1. This machine does not have `grp1` listed in `/etc/group`.

Simultaneous Copy Command

The simultaneous copy command, `simcp`, is installed with the console. The default location is `/opt/webmin/utilbin/simcp`. The `simcp` command relies on passwordless SSH and the `/etc/gridhosts` file. When you use the `simcp` command, passwordless SSH is used to copy a file or a directory securely to each machine identified in the `/etc/gridhosts` file. The following example copies the `/etc/hosts` file to each machine in the cluster:

```
/opt/webmin/utilbin/simcp /etc/hosts /etc
```

You can also use the `simcp` command to copy files recursively. For example, if the `/home/sas/inst` directory has multiple files, the following example copies the files to `/tmp`:

```
/opt/webmin/utilbin/simcp /home/sas/inst /tmp
```

Index

C

CGroup Resource Management 55
 cube data files
 SAS High-Performance Risk 49

G

Greenplum Database
 drop temporary tables 47
 superuser privilege 45
 terminate long running queries 46
 gridhosts file 63
 location and format 64

I

IP address and port number 5

L

logs
 configuring 8
 viewing console logs 21

M

memory and CPU utilization 55
 middle-tier shared key 30

N

network
 IP address and port number 5

R

resource management 55
 return code
 simsh 67

S

security
 limit logon access 39
 simcp 68
 simsh 67
 SSH keys 30
 generating and propagating 34
 middle-tier shared key 30
 SSH lockout 39, 42

U

UNIX groups

- creating [35](#)
- members [36](#)

UNIX user accounts

- creating [31](#)
- deleting [35](#)
- distribute middle-tier shared
key [30](#)

- secondary group membership
[36](#)

user management

- audit SSH logon activity [37](#)
- creating users [31](#)
- deleting users [35](#)
- force log off console user [19](#)
- group membership [36](#)
- logged on users [38](#)
- prevent SSH logons [39](#)