

Storage Services Connector 2.0 for SAP™

Administration Guide

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 U.S.A.
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>
Part number 215-04148_A0
November, 2008

Copyright and trademark information

Copyright information

Copyright © 1994–2008 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted Network Appliance material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETWORK APPLIANCE “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL NETWORK APPLIANCE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, Data ONTAP, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, SyncMirror, Topio, VFM, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. Cryptainer, Cryptoshred, Datafort, and Decru are registered trademarks, and Lifetime Key Management and OpenKey are trademarks, of Decru, a Network Appliance, Inc. company, in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexShare, FPolicy, HyperSAN, InfoFabric, LockVault, NOW, NOW NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simplicore, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, SohoFiler, SpinMirror, SpinRestore, SpinShot, SpinStor,

StoreVault, vFiler, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the U.S.A.

IBM, the IBM logo, AIX, and System Storage are trademarks and/or registered trademarks of International Business Machines Corporation.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the U.S.A. and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	vii
Chapter 1	Introduction to Storage Services Connector for SAP	1
	Storage Services Connector	2
	Storage Services Connector server	3
	Storage Services Connector agent	4
	Storage Services Connector subsystem applications	5
	Where to go for further information	7
Chapter 2	Installing Storage Services Connector	9
	Version compatibility requirements	10
	Licensing	11
	SSC Server and Agent setup	12
	Configuring SnapMirror and SnapVault access	19
	Configuring the ndmpd.preferred_interface option	20
Chapter 3	Configuring the subsystem applications	21
	Protection Manager overview and concepts	22
	Protection policy related tasks	25
	Resource pool related tasks	27
	Data back up and verification.	29
	Operations Manager concepts and tasks	31
	Managing storage system credentials.	34
Chapter 4	Storage Management Operations	35
	Archive log management	36
	SAP system backup management	38
	SAP system storage provisioning.	41

Chapter 5	Command Line Options.	43
	Credential Manager options	45
	System Command Line Options	47
	Storage Command Line Options	55
Chapter 6	Troubleshooting	63
	Glossary	65
	Index	67

Preface

About this guide	This guide describes how to use the NetApp® Storage Services Connector for the web services infrastructure. The guide also provides information about installing Storage Services Connector, subsystem applications.
Audience	This guide is written for system administrators and others interested in managing storage systems using Storage Services Connector. The user should have a working knowledge of NetApp storage systems, UNIX® servers, and SAP software. This guide assumes the following: ◆ You are familiar with the Data ONTAP operating system software. ◆ You are familiar with Web service integration using WSDL, SOAP, and HTTP(s).
Terminology	Storage systems that run Data ONTAP are also referred to as <i>NetApp storage system</i>, or <i>systems</i>. The terms used in Storage Services Connector guide reflect some of these common usages. When the term <i>appliance</i> is used in Storage Services Connector, the information applies to all supported storage systems, NearStore® systems, <i>Operations Manager</i> is the Web-based user interface of DataFabric® Manager, from which you can monitor and manage multiple storage systems, or storage system clusters. Operations Manager is used for day-to-day monitoring, alerting, and reporting on storage systems. <i>Protection Manager</i> helps you manage backup and mirror relationships easily and more efficiently than by using traditional methods of deploying and managing data protection.
Command conventions	You can enter storage system commands on the system console or from any client that can obtain access to the storage system using a Telnet session. In examples that illustrate commands executed on a UNIX workstation, the command syntax and output might differ, depending on your version of UNIX.

Formatting conventions

The following table lists different character formats used in this guide to set off special information.

Formatting convention	Type of information
<i>Italic type</i>	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fc test adaptername</code> command, you enter the characters “fc test” followed by the actual name of the adapter. ◆ Book titles in cross-references.
Monospaced font	◆ Command and daemon names. ◆ Information displayed on the system console or other computer monitors. ◆ The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

Keyboard conventions

This guide uses capitalization and some abbreviations to refer to the keys on the keyboard. The keys on your keyboard might not be labeled exactly as they are in this guide.

What is in this guide...	What it means...
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
<i>Enter</i>	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards.
<i>type</i>	Used to mean pressing one or more keys on the keyboard.
<i>enter</i>	Used to mean pressing one or more keys and then pressing the Enter key.

Special messages

This guide contains special messages that are described as follows:

Note

A note contains important information that helps you install or operate the system efficiently.

Attention

An attention contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

WARNING

A warning contains instructions that you must follow to avoid personal injury.

About this chapter

This chapter describes how to use Storage Services Connector server (SSC), Storage Services Connector agent, DataFabric Manager, and Protection Manager in the SAP environment. This chapter also describes where to find the Storage Services Connector for download and how to find Help and further related information.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Storage Services Connector”](#) on page 2
- ◆ [“Storage Services Connector server”](#) on page 3
- ◆ [“Storage Services Connector agent”](#) on page 4
- ◆ [“Storage Services Connector subsystem applications”](#) on page 5
- ◆ [“Where to go for further information”](#) on page 7

Storage Services Connector

What is new in Storage Services Connector 2.0

Following are the enhancements in Storage Services Connector 2.0:

- ◆ Support for multiple Protection Managers
- ◆ CLI interface
- ◆ Storage Services Connector deployment without Protection Manager
- ◆ Single install for Storage Services Connector server and agent
- ◆ Optional roll back of restore operation
- ◆ Restore to an alternate location
- ◆ AutoSupport
- ◆ Support for back up of offline database

What Storage Services Connector is

Storage Services Connector (SSC) is used by Web service infrastructure when some type of storage manipulation operation such as storage provisioning, backup, or restore is required. Storage Services Connector provides a Web service based interface between SAP application management framework and storage systems.

Storage Services Connector solution consists of the following two components:

- ◆ Storage Services Connector server
- ◆ Storage Services Connector agent

The main component is the SSC server which acts as an interface between the Web service tools and the storage system. SSC server receives the requests for storage operations from Web service tools. SSC then passes these requests to the various subsystem applications for storage manipulation operations.

Downloading Storage Services Connector

The Linux version of Storage Services Connector is available for downloading at the NOW™ NetApp on the Web™ site: <http://now.netapp.com>

Storage Services Connector server

What Storage Services Connector server is

SSC server is the main component of Storage Services Connector. An external application interacts with the SSC server using SOAP requests over HTTPS. This includes requests to access the storage device and perform backup, restore, and recover operations.

What Storage Services Connector server does

SSC server processes requests in an asynchronous manner. The response to the original request provides an operation ID for the operation that is being executed with SSC server. The user can use this operation ID to get the status of that operation until it is complete.

SSC server queues the requests based on different criteria. For example, Only a specific request type or a specific number of requests are processed at a time. The remaining requests are kept in a queue until the earlier ones are processed. This asynchronous processing allows requests to be managed in a manner that prevents the storage system from the impact of multiple and simultaneous storage requests.

SSC server interacts with Operations Manager and Protection Manager to complete the various storage manipulation request.

SSC server and Operations Manager: SSC server uses Operations Manager to make storage system requests and to manage the storage system credentials.

SSC server and Protection Manager: SSC server uses Protection Manager to monitor the status of backups and their protection status across the storage systems. It is also used to troubleshoot the errors that can occur during the back up process.

Storage Services Connector agent

What Storage Services Connector agent is

SSC agent is the other major component in the Storage Services Connector solution. The agent is installed on the SAP system host systems. SSC agent runs as the root user to reduce configuration needed on a per-host basis. The agent acts as an interface between the database and SSC server.

What Storage Services Connector agent does

SSC agent is used by SSC server to perform operations during the backup and recovery process. SSC agent accepts incoming SOAP requests that are only for SSC server use. SSC agent provides the following services to SSC server:

- ◆ Start the database
- ◆ Stop the database
- ◆ Quiesce the database
- ◆ Unquiesce the database
- ◆ Recover database to a point in time
- ◆ Recover database with all available log files

Storage Services Connector subsystem applications

The Storage Services Connector solution uses Operations Manager and Protection Manager applications as its subsystems. These applications are used, for example, to set up storage system credentials, assign capabilities to users, create resource pools, and so on. The following sections give a brief overview about these subsystem applications. For detailed information about the usage of these applications in the Storage Services Connector solution, see Chapter 2, [“Configuring the subsystem applications”](#) on page 21.

What DataFabric Manager is

The DataFabric Manager server provides infrastructure services such as discovery, monitoring, role-based access control (RBAC), auditing, and logging for products in the NetApp Storage and Data Suites. DataFabric Manager software runs on a separate server. It does not run on the storage systems.

DataFabric Manager has a command-line interface (CLI) for scripting commands that might otherwise be performed using its Web-based UI, Operations Manager.

DataFabric Manager server also provides infrastructure services for various applications, such as Protection Manager.

Note

DataFabric Manager is the subsystem application while Operations Manager is the Web-based user interface of DataFabric™ Manager.

Administrator accounts in DataFabric Manager

DataFabric Manager has two different administrator accounts:

- ◆ Administrator
- ◆ Everyone

The Administrator account grants full access for the administrator who installed DataFabric Manager software. An Everyone administrator account allows users to have read-only access without logging in. You can modify the permissions of the Everyone administrator to suit your needs.

What Protection Manager is

Protection Manager helps you unify and automate backup and disaster-recovery operations with a policy-based management tool. It provides end-to-end, workflow-based design and seamless integration of SnapVault®, SnapMirror™, and Open Systems SnapVault to enable you to manage large-scale deployments easily.

Data management objects in Protection Manager

Protection Manager uses objects called data sets, policies, and resource pools to help you organize your data. It also provides rules using which you can protect your data, and the storage resources in which you want to contain your backups and mirrors.

Protection Manager uses these objects to automate routine data protection tasks and help you manage your backup and mirror relationships more efficiently. For more information about data sets, policies and resource pools, see Chapter 3, [“Configuring the subsystem applications,”](#) on page 21.

Where to go for further information

Additional information

The following table provides information about the documentation available on the NOW™ site. This information can be useful when you are working in the NetApp system environment.

If you want...	See...
The most current information for the following NetApp products: ◆ DataFabric Manager ◆ Protection Manager	The following documents: ◆ DataFabric Manager with Operations Manager 3.6 <i>Administration Guide</i> ◆ Protection Manager <i>Administration Guide</i>

About this chapter

This chapter describes the process of installing Storage Services Connector server and agent and its version compatibility requirements. It also describes the various configuration options related to SSC server and SSC agent.

Topics in this chapter

This chapter contains the following topics:

- ◆ “[Version compatibility requirements](#)” on page 10
- ◆ “[Licensing](#)” on page 11
- ◆ “[SSC Server and Agent setup](#)” on page 12

Version compatibility requirements

Version requirements

Storage Services Connector supports the Web service environment platform. The products and the versions that are supported are as follows:

Operating System: Following is the operating system supported for the SSC server:

- ◆ SUSE® (64 bit) Linux® 9 on EM64T. This restriction is specific to Operations Manager platform support.
The operating system for the host agent system can be either SUSE (64 bit) Linux 9 or SUSE (64 bit) Linux 10.

Database: The databases that are supported with this version of Storage Services Connector are as follows:

- ◆ IBM® DB2 Version 9
- ◆ MySQL MaxDB Version 7.7

Data ONTAP: Data ONTAP 7.2.2 and later is supported with Storage Services Connector.

Storage Protocol: NFS v3 is the storage protocol that is supported with this version of Storage Services Connector.

DataFabric Manager and Protection Manager: Storage Services Connector requires DataFabric Manager 3.7 and Protection Manager. Protection Manager is a part of NetApp Management Console. You can download NetApp Management Console from Operations Manager.

Note

For more information on the supported applications, see the Support Matrix on the NOW site.

Licensing

The user needs to license the product to download the software. However, no license key required to set up and use the product.

Storage Services Connector requires Data Protection Manager and Operations Manager for set up. You need to purchase and install these products separately before installing SSC:

- ◆ Operations Manager core
- ◆ Protection Manager (a feature of Operations Manager)

You can access your license key from <http://now.netapp.com/licenses>.

SSC Server and Agent setup

Installing SSC server and the subsystems

SSC server is installed on a host system running SUSE (64 bit) Linux 10 on EM64T.

To set up the SSC server environment, complete the following steps.

Step	Action
1	Install SSC server, agent, and CLI using the Red Hat® Package Manager installer. Use the following command to install the SSC server, agent, and the CLI: <pre>rpm -I <package-name></pre> Note————— You should have standard J2SE version 5 installed on the system. The Red Hat Package Manager installer for SSC server does not include Java™ Virtual Machine
2	[Optional] If you want to use SSC with a crash database on MaxDB, Install MaxDB.
3	[Optional] Create a MaxDB instance for SSC server crash recovery. Remember the following when you create a MaxDB instance: ◆ The crash database should be in NO ARCHIVELOG mode. ◆ An SQL user has to be created. The SSC server is by default configured to use the user name csm, and the password netapp. If you want to use a different user name and password, then you need to configure the SSC server to use the new user name and password. ◆ Install the schema from the script <code>/opt/NetApp/ssc/bin/maxdb.sql</code> into the SQL user created.

Step	Action
4	Install the DataFabric Manager server 3.7 on the same host as the SSC server. For more information about how to install DataFabric Manager, see the DataFabric Manager <i>Installation and Upgrade guide</i>. Install Protection Manager. To download Protection Manager, in Operations Manager (Web UI of DataFabric Manager), click Setup > Download Management Console. Protection Manager is a part of Management Console. For more information about how to install Protection Manager, see the Protection Manager <i>Administration guide</i>.

Note

HTTPS is enabled by default. If you want to disable HTTPS, set `http.ssl.enable` to `false` in `services.config` file.

Enabling server or agent

To use SSC server, SSC agent or both, set the following parameters in the `services.config` file:

```
# Service initialization properties
# Service Initialization Mode
# web.interfaces=[agent|server|both]
# web.interfaces=both
```

Setting HTTP properties

Use the following parameters in the `services.config` file to set the HTTP properties:

```
# HTTP Properties
#http.port=8146
#http.ssl.enable=true
#http.ssl.keystore.file=etc/keystore
#http.ssl.keystore.password=NetAppCSMAgent
```

Configuring HTTPS on SSC server

The SSC server comes with a generic HTTPS keystore to allow basic HTTPS operation without any necessary configuration.

The SSC also ships with a generic truststore for communicating with the tenant agents. If the HTTPS keystore is changed on the tenant agents, then the truststore must also be updated.

There are several properties in the `services.config` file that you can configure. These properties are as follows:

- ◆ `http.ssl.enable`. The default value is `true`. You can set this value to `false` to turn off HTTPS.
- ◆ `http.ssl.keystore.file`. The default value is `etc/keystore`. You can change the location of the keystore file.
- ◆ `http.ssl.keystore.password` sets the pass-phrase of the keystore.
- ◆ `http.ssl.key.password` sets the certificate key password.

Configuring HTTPS on SSC agent

The SSC agent comes with a generic HTTPS keystore to allow basic HTTPS operation without any necessary configuration.

The SSC also ships with a generic truststore to allow it to communicate with the tenant agents. If the keystore is modified on the tenant agents, then the SSC server must also be updated with a new truststore.

There are several properties in the `agent.config` file that are necessary to configure SSL. These properties are:

- ◆ `http.ssl.enable`. The default value is `true`. You can set this value to `false` to turn off HTTPS.
- ◆ `http.ssl.keystore.file`. The default value is `etc/keystore`. You can change the location of the keystore file.
- ◆ `http.ssl.keystore.password` sets the pass-phrase of the keystore.
- ◆ `http.ssl.key.password` - sets the certificate key password.

[Optional] SSC server host configuration

The process that runs SSC server requires that the following conditions be met:

- ◆ The MaxDB `dbmcli` utility should be available on `$PATH`. This is required to determine the location of the JDBC drivers for MaxDB used in connecting to the SSC server crash-database.
- ◆ The `JAVA_HOME` environment variable should be defined to reference the root of a Java 5 runtime installation.

To configure SSC server, you need not log in as the root user. It is recommended that a local, non-login user be used.

Note

If SSC server is run as an user other than the root user, then the command `chown -R <user>:<group> <installation directory>` should be used to allow the correct user permissions.

SSL settings for storage system

Ensure that the following settings are correct with respect to SSL and storage system:

<code>httpd.admin.ssl.enable</code>	<code>on</code>
<code>httpd.admin.enable</code>	<code>on</code>
<code>ssl.enable</code>	<code>on</code>

Changing SSC crash-database to MaxDB

In SSC 2.0, the crash-database used for SSC server is an embedded database. No additional steps required to use the embedded database. However, if you want to change the crash-database to MaxDB, set the following parameter in the `csm.config` file:

```
jdbc.use.embedded=false
```

[Optional] Installing and configuring crash-database

The SSC server persistence database also referred to as crash-database is a small database. SSC uses crash-database to store metadata about backups and status of its work flows.

Note

Do not confuse crash-database with the SAP database. The SAP database can also be a MaxDB database.

The SSC might not be used to back up its own persistence database. In SSC 2.0, the default SSC server persistence is Derby. You can change the default database to MaxDB. To change the default to MaxDB, see [“Controlling lifetime of contexts”](#) on page 17

The SSC server stops processing requests when the crash-database goes down. Restarting the crash-database is detected and SSC server resumes the processing of current and new requests.

To install and configure the crash-database, complete the following steps:

1. Install MaxDB. To install MaxDB, unpack the .tar file and enter the start sdbinst command.

2. Create the directories for crash database.

```
mkdir -p /sapdb/csmcrash/sdbcsmcra
mkdir -p /sapdb/csmcrash/sapdata1
mkdir -p /sapdb/csmcrash/saplog
```

3. Change ownership by entering the following command:

```
chown -R sdb:sdba /sapdb
```

Create a crash database instance by entering the following command:

```
dbmcli db_create csmcrash dbm,dbm
dbmcli -d csmcrash -u dbm,dbm <<EOF
param_startsession
param_init OLTP
param_put RUNDIRECTORY /sapdb/csmcrash/sdbcsmcra
param_put SESSION_TIMEOUT 0
param_put CACHE_SIZE 2500
param_put _SERVERDB_FOR_SAP YES
param_put MAXDATAVOLUMES 64
param_put MAXLOGVOLUMES 8
param_put MAXUSERTASKS 100
param_put MAXLOCKS 500000
param_checkall
param_commitsession
param_addvolume 1 DATA /sapdb/csmcrash/sapdata1/DISKD0001 F 5000
param_addvolume 1 LOG /sapdb/csmcrash/saplog/DISKL0001 F 1000
db_admin
db_activate dba,dba
load_systab
auto_update_statistics ON
auto_extend ON
medium_put CSMCRASH_ARCHIVE_LOG
/sapdb/csmcrash/saplog/csmcrash_archive.log FILE AUTO
medium_put CSMCRASH_BACKUP /tmp/CSMCRASH_COMPLETE FILE DATA
db_execute SET LOG AUTO OVERWRITE OFF
db_connect
backup_start CSMCRASH_BACKUP DATA
autolog_on CSMCRASH_ARCHIVE_LOG
quit
EOF
```

5. Create a user for Storage Service Connector with user name csm and password csm.

```
sqlcli -d csmcrash -u dba,dba CREATE USER CSM PASSWORD CSM RESOURCE
NOT EXCLUSIVE ENABLE CONNECT
```

sqlcli -d csmcrash -u dba,dba is the command to start a SQL session to the CSM

crash-database with user dba and password dba.

*CREATE USER CSM PASSWORD CSM RESOURCE NOT EXCLUSIVE
ENABLE CONNECT* conveys the following information:

1. Create a database user CSM with password CSM
2. RESOURCE is the database user class
3. NOT EXCLUSIVE means that the user has permission to have multiple database connections
4. ENABLE CONNECT means that the user can open database connections

6. Load the schema

```
sqlcli -d csmcrash -u csm,csm -i  
{p4ciientroot}/opt/NetApp/ssc/bin/maxdb.sql
```

The following properties configure the crash database for SSC server in the csm.config file. The csm.config file is in the etc directory of the installation directory. The following are the default properties:

```
csm.persistence.jdbc.user=csm  
csm.persistence.jdbc.password=csm  
csm.persistence.jdbc.url=jdbc:sapdb://127.0.0.1/csmcrash  
csm.persistence.jdbc.maxActive=3  
csm.persistence.jdbc.initialSize=1
```

Controlling lifetime of contexts

You can also set the following options in the csm.config file to control the lifetime of the contexts. The context parameters defines the time duration for which the status of the completed operations is maintained by SSC server or SSC agent. During this time duration if the client requests for the status of the completed operation, it is provided with the status. Beyond the specified time duration, a message is returned informing that the operation status was not found.

The default values on the SSC server are as follows:

```
flow.persistence.contextLifetime=720 # 12 hours  
flow.persistence.scavengeInterval=30 # 30 minutes  
flow.persistence.requeueOnInit=true # Requeue running contexts at  
restart
```

Set the following option on the agent:

```
flow.persistence.contextLifetime=30 # 30 minutes
```

SSC server access to Protection Manager

To configure access of the SSC server to Protection Manager, use the following command line option:

```
ssc credential create
```

```
-type DFM
-target <dfm_server>
-credential
-username <username>
-password <password>
```

AutoSupport in SSC SSC 2.0 supports AutoSupport. AutoSupport gathers the following basic information:

- ◆ Operating system version
- ◆ DataFabric Manager version
- ◆ Data ONTAP version, and
- ◆ Operating system architecture, whether it is 32 bit or a 64 bit architecture

AutoSupport is enabled by default. To disable AutoSupport change the following variable in the `csm.config` file:

```
csm.asup.disable=TRUE
```

You can add multiple hosts for AutoSupport. To enable AutoSupport on multiple hosts, set the following parameter in the `csm.config` file:

```
csm.asup.host=<asup_hostname>[, ...]
```

The date and time for regular AutoSupport message is set by the following parameter:

```
csm.asup.date=Mon, 2:00
```

The date format is set by the following parameter:

```
csm.asup.date.format=E, HH:mm
```

You can set the time interval between successive AutoSupport messages by the following parameter:

```
csm.asup.interval=604800000
```

The default value for the above parameter is one week. The value is displayed in milliseconds.

Note

It is not recommended to switch off AutoSupport.

Configuring SnapMirror and SnapVault access

Protection Manager does not set up the storage system level access permissions for SnapMirror and SnapVault to be able to communicate between storage systems. These permissions have to be configured on the storage system as part of the deployment. Access options can be restricted based on the connections that are allowed. This depends on how the web service administrator configures it.

If specific storage system access is set up, then only the defined relationships work. The storage system access permissions can be set up using the following storage system commands:

```
options snapmirror.access *
```

```
options snapvault.access *
```

Alternatively, you can use the following commands also:

```
options snapmirror.access host=<comma separated list of hosts  
allowed to access>
```

```
options snapvault.access host=<comma separated list of hosts  
allowed to access>
```

Configuring the `ndmpd.preferred_interface` option

What this option is about

This option allows you set the preferred interface for back ups. It also lets you bind the backup relationship to an interface and still allows you to have all the interfaces available for DataFabric Manager.

Using the option

You can specify the storage system network interface to be used when establishing an NDMP data connection to another storage system.

By default, an NDMP data connection uses the same network interface as the NDMP control connection established by the NDMP backup application. However, when a data connection between an NDMP-enabled device is required over an alternate network, it is necessary to specify the storage system's interface through which the alternate network will be accessed.

For example, a UNIX or NT resident NDMP back up application and multiple NDMP-enabled storage system can be interconnected via a corporate network. The same NDMP-enabled devices can also be interconnected via an isolated private network. To minimize the load on the corporate network, the `ndmpd.preferred_interface` option can be used to direct all NDMP data connections over the isolated private network.

To specify the preferred network interface to be used for NDMP data connections, issue the following command on the storage system:

```
options ndmpd.preferred_interface <interface>
```

To disable a preferred network interface specification and force the NDMP default interface to be used for data connections, issue the following command:

```
options ndmpd.preferred_interface disable.
```

The default value for the `ndmpd.preferred_interface` option is `disable`.

The `ndmpd.preferred_interface` option retains its previous value when a storage system restarts.

For back ups, set the value on the primary system. For restores, set the preceding option on the secondary system.

About this chapter

This chapter provides information about the subsystem applications, Protection Manager, and DataFabric Manager. Protection Manager and Operations Manager provide many more features than what are used in the Storage Services Connector solution. The description of those is beyond the scope of this guide.

Topics in this chapter

This chapter contains the following topics:

- ◆ [“Protection Manager overview and concepts”](#) on page 22
- ◆ [“Protection policies”](#) on page 23
- ◆ [“Resource pool related tasks”](#) on page 27
- ◆ [“Operations Manager concepts and tasks”](#) on page 31
- ◆ [“Managing storage system credentials”](#) on page 34

Protection Manager overview and concepts

Purpose of Protection Manager

In Storage Services Connector solution, Protection Manager is used to manage back ups made using SSC server. Protection Manager is also used to manage the SnapVault relationships between primary and secondary storage systems. Protection Manager monitors the conformance of the storage system relationships and provides diagnostic results when these relationships do not conform according to the policies.

Protection Manager data management tools

Protection Manager provides three powerful tools to manage data. These are *data sets*, *resource pools*, and *protection policies*. The following section provides details about these storage tools.

Data sets

A data set is a collection of data managed as a single unit, and all the replicas of that data. The data is identified by the volume, qtree, or directory in which it is located.

You manage a data set as a single unit, therefore its members should have common management requirements. In Protection Manager, members of a data set should share the same data protection requirements.

Data sets are created by SSC server based on the back up request received. The data set created represents a SAP system.

Resource Pool

A resource pool is a collection of currently unused physical storage from which other storage containers can be provisioned. If you assign a storage system to a resource pool, all aggregates and spare disks on that storage system become available for provisioning.

In Protection Manager, you assign a resource pool to the backup and destinations of a policy which is assigned to a data set. Protection Manager can then provision volumes automatically. To prevent conflicts, physical storage assigned to one resource pool cannot be assigned to a second resource pool.

Protection policies

A policy is a set of rules that specify the intended management of data set members. You can apply the same policy to multiple data sets, leveraging your configuration of the policy across the data sets. If you update a policy, the update is propagated across all the data sets to which the policy is applied.

In Protection Manager, protection policies define how protection relationships should be structured and the property settings for each component of the structure. For example, a protection policy might specify that the primary data is backed up to a secondary location and that the secondary copies are mirrored to a tertiary location.

Policy used by SSC server

SSC server uses only the *back up* protection policy provided by the Protection Manager. You can copy and then customize this policy to meet your data protection requirements. In the back up policy a data set gets both local primary storage back up and remote back up from primary storage to secondary storage.

The Web service Administrator configures the back up policy using Protection Manager. This policy is provided to the SSC server during a back up request for the datasets created within the back up SAP system.

What is back up data retention

You can schedule four categories of backup (Hourly, Daily, Weekly, Monthly) for retention purposes. *Retention duration* is the minimum time that you want the backup copies to be preserved and protected from deletion in the primary or secondary storage.

Note

You can only specify retention durations for backup copies. To guarantee preservation of the most recent successful backups, Protection Manager always retains copies to the last two successful back ups even if those copies are older than their specified minimum retention duration.

How long you decide to retain a particular backup depends on the original purpose of that back up.

- ◆ An example of backups not required beyond a 24-hour duration:
Local backups at the primary storage location. These backups are executed every hour for quick restoration of any recent data accidentally corrupted or deleted.
- ◆ An example of backups that need to be retained for a week, a month, or longer:

A daily remote backup of accidentally deleted or corrupted data to a secondary storage location. It is carried out after working hours to ensure preservation. When required, this backup is used to restore a system to its state on a particular date.

Retention settings are made when policies are set up. When a back up is requested using the SSC server, only the retention type, for example, daily, weekly, and so on is specified.

For a detailed description of the back up retention options, see the Protection Manager *Administration Guide*.

Lag warnings

Protection Manager generates lag warnings and errors based on the threshold definitions in the *policy* when a back up or transfer has not happened in the specified threshold time. These warnings help you detect whether the Production Reporting and Control Center tool has requested back ups based on the expected schedules or not. If you do not want to generate lag warnings frequently, set the thresholds to a high enough value.

Choosing between SnapVault and qtree SnapMirror relationship

Protection Manager analyzes the protection policy, available resources, and licenses on the storage systems to determine whether it can create a qtree SnapMirror relationship. If the qtree SnapMirror relationship is not possible, then it uses SnapVault relationship. You can change this behavior using the following CLI command:

```
dfm option set pmQSMBackupPreferred=No
```

If the above mentioned option is set, then Protection Manager first tries to use SnapVault relationship. If the SnapVault relationship does not work, it uses qtree SnapMirror relationship (if licensed).

You need to keep the following points in mind when you set the above mentioned option:

- ◆ If the schedule specifies that the back up of data is to be taken up more frequently than once per hour, then Protection Manager must create a qtree SnapMirror relationship regardless of the preceding option.
- ◆ If the source qtree contains a space in its name, Protection Manager tries to create only SnapVault relationships regardless of preceding option. If the source data is an open system directory, the relationship must be Open Systems SnapVault.

Protection policy related tasks

Copying a protection policy

To use the backup policy with SSC server, you need to make a copy of a protection policy and then customize it so that you can apply it to data sets. The following table describes how to copy a protection policy.

Before starting, ensure that you are assigned the appropriate role to edit protection policies. Complete the following steps.

Step	Action
1	From the left navigation pane, click Set Up > Protected Data > Policies. The Policies tab displays a list of backup policy types to configure.
2	In the Policies list, select the Back up policy. Click Copy. This action creates a "Copy of Back up" policy.
3	To rename the copy of the back up policy created in step 2, select Copy of Back up and click Edit. A Copy of Back up Properties window appears. You can rename the policy by using this window.

Editing the back up protection policy

You can edit the back up protection policy by selecting it on the Policies tab and clicking its nodes and connections.

Before you edit the back up protection policy, you should observe the following:

- ◆ Ensure that you are assigned the appropriate role to edit protection policies.
- ◆ Confirm that the change you want to make to this protection policy is acceptable for all data sets using this policy.

Step	Action
1	From the left navigation pane, click Set Up > Protected Data > Policies. The Policies tab displays a list box of protection policies which includes the back up policy you created.

Step	Action
2	In the Policies list box, select the back up protection policy you created. Check the Dependencies list at the bottom to view all the data sets and resources that are assigned this policy. Confirm that you want to apply the edits that you make to this policy to all the data sets in the Dependencies list.
3	With the back up policy still selected, note the selectable node and connection components displayed in the protection policy diagram.
4	Click each node or connection in the policy diagram to display and edit the protection properties for that component. After modifying properties for one node or connection, click Apply before displaying and modifying the properties for another component for this policy.

For Storage Services Connector solution, you need to set the local back up and transfer schedules to *None*.

For more information about policies, see Protection Manager *Administration Guide*.

How does having more than one policy help

You can configure more than one policy using Protection Manager. By configuring different policies, you can have different retention periods for back ups. This is particularly useful when you have different types of SAP systems.

Resource pool related tasks

Creating a resource pool

This procedure describes how to create a new resource pool by using the Resource Pool wizard.

Ensure that you are assigned the appropriate role to create a resource pool, such as the GlobalFullControl role. If you encounter access problems, contact the administrator who maintains the DataFabric Manager roles and privileges.

Ensure the host storage systems you are adding to the resource pool have the proper configuration and licensing for their intended use.

Have the information available that you need to complete this task:

- ◆ Name you want assigned to the resource pool
- ◆ Description of the resource pool
- ◆ Owner of the resource pool
- ◆ E-mail address for contacting the owner
- ◆ Time zone used for the resource pool

Step	Action
1	From the left navigation pane, click Set Up > Resource Pools. The Resource Pools window opens.
2	Click the Add button to open the Resource Pool wizard and then complete the wizard.
3	Verify the creation and content of the resource pool by viewing the results that are displayed in the Set Up > Resource Pools window.

Which storage to include in the resource pool

To decide which storage to include in the resource pool, ensure the following:

- ◆ Whether the resource pool will contain individual aggregates, an entire storage system (all the aggregates in a storage system), or a combination
- ◆ Whether you want to use traditional volumes or flexible volumes in the resource pool
- ◆ How much available storage space you need for the resource pool, based on the resource pool's intended use

The decision can also depend on whether you want to segregate different SAP systems to a specific back up storage system. Alternatively, you can choose to have one resource pool with all of the secondary storage.

Monitoring the resource pool

You can monitor the resource pool that you created by using the dashboard provided in Protection Manager. The dashboard displays the resource pool name, and the total space allocated and utilized. The total size allocated is indicated in gigabytes or terabytes. The space utilization is indicated in percentage of the total size. You can access the dashboard by clicking Monitor > Dashboard.

Data back up and verification

Data back up using SSC server

SSC server backs up your local data to the primary storage system, and from the primary storage system to a secondary storage system. The information to back up data, such as the name of the policy and resource pool, is provided by web service tools to SSC server. The sequence of actions that place for the back up of data is as follows:

1. SSC server creates the dataset, if it does not already exist.
2. SSC server adds the storage members to the dataset created in step 1.
3. SSC server assigns the policy to the dataset. The policy assigned is the *back up* policy or a policy which is a copy of the back up policy.
4. SSC server assigns the resource pool to the destination node of the policy.
5. SSC server then does the following:
 - a. Requests the database agent to quiesce the SAP system database
 - b. Requests Protection Manager to take the back up and protect it
 - c. Requests the database agent to unquiesce the database (This happens when all the Snapshot™ copies are taken by Protection Manager)

The operations listed in steps 1 and 2 occur only during the first back up of the SAP system. This is because during subsequent back ups the dataset already exists and the SSC server does not need to create the dataset again.

The operations listed in steps 3 and 4 are not repeated for each back up unless they differ from the policy or resource pool assignments currently in the dataset. When you want to migrate your data back up to a different storage system, you need to update the resource pool to start using the new storage system.

Troubleshooting back up problems with Protection Manager

Protection Manager is used to troubleshoot errors like those that can occur due to policy non-conformance. Following are few possible type of problems that can occur:

- ◆ Invalid storage system configuration. This can prevent SnapVault to be set up between a primary and a secondary storage system.
- ◆ Errors can happen during the protection operation. These can occur because of storage system level failure during the SnapVault or SnapShot process, network failure, the resource pool not having enough resources, such as Aggregate free space or enough volumes.

To diagnose the cause of the problem, you can use the diagnose wizard in Protection Manager. To use this wizard for storage system, navigate to Control > Hosts > Storage System > Diagnose. Follow the instructions in the wizard to diagnose the problem.

Operations Manager concepts and tasks

What role based access control is

DataFabric Manager uses role-based access control (RBAC) for user login and role permissions. RBAC allows administrators to manage groups of users by defining roles.

Note _____
SSC server uses Operations Manager to set up and manage user credentials to access the storage system.

Configuring users

DataFabric Manager uses administrator accounts to manage access control and maintain security. When you install DataFabric Manager software, default administrator accounts are created: the Administrator and Everyone accounts. Both accounts are global and have predefined roles assigned to them.

Administrator account: The Administrator has super-user privileges and can perform any operation in the DataFabric Manager database. The Administrator account is given the same name as the name of the administrator who installed the software. Therefore, if you install DataFabric Manager on a UNIX workstation, the administrator account is called “root”.

Everyone: After installing DataFabric Manager, you must log in as the administrator and set up the Everyone account to grant view permission on this account.

Adding administrative users

The following table describes the procedure to create and edit administrator accounts using Operations Manager

Step	Action
1	In the Control Center, from the Setup menu, select Administration users. The Administrators page is displayed.
2	Log in to the Administrator account.

Step	Action
3	Type the name for the administrator or domain name for the group of administrators.
4	Optional: Enter the e-mail address for the administrator or administrator group.
5	Optional: Enter the pager address as an e-mail address or pager number for the administrator or administrator group.
6	Click Add.

To set up administrator accounts by using the CLI, use the `dfm user` command.

Default SSC server account

SSC server is configured with a default user account called `csm`. This default user account has to be configured on the host and the access is granted for it to be a user in DataFabric Manager.

The default user `csm`, and password `netapp` should be assigned the following three roles.

GlobalDataProtection: You can perform all the operations of Globalbackup, GlobalRead, and GlobalDataSet.

GlobalRestore : You can perform restore operations from back ups on secondary volumes.

GlobalWrite: You can view or write to the DataFabric Manager database.

Note


SAP administrator can change the user name and password. If the user name and password are changed, SSC server should be configured to use them.

Capabilities

When creating roles, you must assign *capabilities*, the combination of operations and resources, to the role. You can view capabilities or edit them by modifying the operations that are associated with the resource.

Setting up access capability

The DFM.Core.Control capability is required to enable an Operations Manager access to the Proxy API mechanism of Operations Manager. You access this capability from Setup > Roles. The Add Roles page displays. Click the Add Capabilities link in the Add Roles page. The following example (image) shows the Add Capabilities window.



Note

You need to assign the role created to the SSC server user.

For more information on roles and capabilities, see the DataFabric Manager with Operations Manager *Administration Guide*.

Managing storage system credentials

Managing the credentials using Operations Manager

You can set the credentials for your storage system using the Operations Manager user interface. To set the credentials complete the following steps.

Step	Action
1	Go to the Appliance Details page for the storage system and choose Edit Settings from the Appliance Tools menu (The Appliance Tools menu is on the lower left side of the page). The Edit Appliance Settings page displays.
2	In the Login field, enter a user name that DataFabric Manager uses to authenticate the storage system.
3	In the Password field, enter a password that DataFabric Manager uses to authenticate the storage system.
4	Click Update.

Note

After the password is changed for the SSC DataFabric Manager user and a subsequent SSC credential update with the new password, DataFabric Manager needs around 20 minutes to allow SSC access with the updated password.

Managing the credentials using Protection Manager

You can manage the credentials of the storage systems by also using the Protection Manager. To set the credentials using Protection Manager, click Set Up > Hosts. In the Set Up Hosts page, select the storage system for which you want to set or edit the credentials and then click Edit.

About this chapter

This chapter describes the storage management operations that can be performed using Storage Services Connector and gives the detailed assumptions to execute those operations. These functions are provided as a SOAP Web service from the SSC server.

For more information on the exact format of the interface methods used, see the WSDL documentation. You can access the web services description from https://<ssc_host>:<ssc_port>/csm/services/CSM?wsdl. The default port is 8146.

Topics in this chapter

This chapter contains the following topics:

- ◆ “[Archive log management](#)” on page 36
- ◆ “[SAP system backup management](#)” on page 38
- ◆ “[SAP system storage provisioning](#)” on page 41

Archive log management

Storage Services Connector supports management of database archive logs, which supports the SAP system restore and recovery process. For archive log management, the SAP administrator needs to complete the following steps:

1. Create an archive log volume on the primary and secondary storage systems.
2. Create a qtree for each SAP system in the archive log volume.
3. Mount the qtrees on the SAP system.
4. Configure the database to write its archive and overflow logs to the mounted qtrees.

The SAP system supporting the DB2 or MaxDB database utilize the NFS mount of the allocated qtree as the destination location for its archive log process.

When archive logs are moved to a secondary storage, the files written to the qtree in the primary volume are transitioned an allocated secondary volume or qtree. After the transition of the archive log is confirmed, it can be removed from the the primary storage system.

The secondary qtree created is also available through NFS to SAP system. All the logs which are moved to the secondary storage are available through this single mount.

Note

The archive log file processing does not process hidden files or directories.

Backup of the archive log available on primary volume

Archive logs are written only once by the database system to the primary archive log storage and are not modified after that by the database system. This condition is used in determining when an archive log file has been completely written to a primary storage and then subsequently copied to a secondary storage so that it can be removed from the primary storage system.

When the client requests a backup of a primary archive log volume to a secondary archive log volume, SSC server copies the files from the primary qtree to the secondary qtree. If the files are older than the specified threshold in the archive operation request, SSC server removes them from the primary storage system.

**Purge archive log
secondary volume**

SSC server, after receiving a request from web service, removes archive log files from the secondary storage. The files which that crossed the time limit for being stored on the secondary storage system are deleted.

SAP system backup management

Storage Services Connector supports backup and restore of SAP system and recovery processing of the primary databases supporting a SAP system hosted system.

Pre-requisites for SAP system backup management

- ◆ The database instance exists within the host configurations supported and is properly functioning.
- ◆ Storage Services Connector components are properly installed, configured, and initialized.
- ◆ The SAP system database agent components, are properly installed, configured, and initialized.
- ◆ The SAP system database agent component is accessible as a SOAP service from the SSC server host.

If at any time the database or underlying storage is unable to fulfill a request from Storage Services Connector, the operation fails and it attempts to return the system to the initial state.

Defining back up retention policy

Backup retention policy defines the minimum duration of time for which you want your backup copies to preserved. The back up retention policy is set by using Protection Manager. To know more about how to create secondary storage using Protection Manager, see *Protection Manager Administration Guide*.

Defining secondary storage resources

To take SAP system back up, you need to configure secondary storage. The secondary storage resources are defined using Protection Manager. To know more about how to create secondary storage using Protection Manager, see *Protection Manager Administration Guide*.

Taking SAP system's back up

For taking a SAP system back up, you should make sure that the following conditions are met:

- ◆ SAP system can be either online or offline.

- ◆ The necessary SAP system database agent is running on the SAP system host and accessible from the SSC server host as a SOAP service.

Web service by SOAP request provides SSC server the details of the SAP system which is to be backed up. After the SSC server receives the request, it requests the Protection Manager to complete the back up operation. During this operation, Storage Services Connector quiesces and unquiesces the database. Storage Services Connector returns the status of the back up to the client. The status can be in-progress, complete, warning, or error.

The back up operation can fail when any one of the following conditions occurs:

- ◆ Protection Manager is unavailable
- ◆ SAP host system is not reachable
- ◆ Primary storage back up failed
- ◆ Secondary storage back up failed

Inquiring about SAP system back ups

You can inquire about the existing SAP system back ups. SSC server provides back up history retained by Protection Manager. The back up data consists of back up dates, protection status, policy conformance status, and so on.

Restoring SAP system from back up

You can restore a SAP system from a back up using SSC server. To restore a SAP system, make sure that the SAP system-hosted application system is offline. This is required because the underlying file system contents are altered during this operation and this could cause unpredictable results to an operating application.

Web service provides information about the back up from which the SAP system is to be restored. When SSC server receives the requests, it verifies whether the database instance is shut down. After that, SSC server requests Protection Manager to restore the file components of the back up to the primary storage.

SSC may not be able to restore if a Snapshot is removed from the primary storage location by using an external application. In such a case it may be possible to manually Snapvault or SnapMirror restore the corresponding members of the dataset from the secondary to the primary storage.

Optional roll back of restore operation

If a restore operation does not complete successfully, SSC 2.0 lets you to roll back the system to its original state. You can use this option by setting the following parameter in the `csm.config` file:

```
csm.restore.rollback=true
```

By default this option is set to `true`. If you do not want to roll back the system state during an incomplete restore, you can set the above mentioned option to `false`.

Restoring to an alternative destination

SSC 2.0 lets you restore your data to some other volume as defined by you. The alternative destination is provided by using the parameter `-destination` with `ssc system restore` command.

Recovering restored SAP system database

You can recover a restored SAP system database. To recover a restored SAP system, the following conditions must be satisfied:

- ◆ SAP system-hosted application system should be offline before you execute this use case. This is required because the underlying file system contents are altered by the use case and can cause unpredictable results for an operating application.
- ◆ Back up of database storage components for the SAP system is restored.
- ◆ Archive log storage for this SAP system is locally available and accessible on the SAP system host.

Web service provides SSC server information to identify the host which runs the SAP system, primary database, and access credentials. Web service also provides information for database to be recovered to the primary database. Based on the request, SSC server recovers the database.

Offline database backup

Using SSC, you can force a offline back up of database. After the back up is complete, the database is brought back to the online mode. You should use the `ssc system backup` command with `-mode OFFLINE_FORCE` option. To achieve this set the `csm.backup.onlineAfterForceOffline=true` in the configuration file.

Deleting SAP system backup

Protection Manager is used to delete a SAP system backup. The backup deletion is based upon the parameters defined in the backup policy used by SSC server. If the backup falls outside the retention period, the backups are deleted.

SAP system storage provisioning

You must keep the following points in mind for SAP system provisioning operations:

- ◆ Storage Services Connector is properly installed, configured, and initialized
- ◆ Storage systems that are manipulated must be network-accessible from the local server.
- ◆ If at any time the underlying storage is unable to fulfill a request from Storage Services Connector, the operation will fail and the system returns to the original state.

Using Storage Services Connector, you can create, delete, clone split, and copy NFS volume on the storage system. You can also create, delete, and restore a snapshot of a SAP system. For more information about the interface method used, see the web services description at the following link:

https://<ssc_host>:<ssc_port>/csm/services/CSM?wsdl

SAP system removal from Protection Manager

Removing a SAP system volume from its primary storage system does not result in the removal of the SAP system's dataset in Protection Manager. In order to remove a SAP system, the dataset must be deleted manually by using Protection Manager. Any backup that were made must also be removed manually from the backup storage systems.

About this chapter This chapter describes the various command line interface (CLI) options.

Topics in this chapter This chapter contains the explains the following CLI options:

- ◆ “[Credential Manager options](#)” on page 45
 - ❖ “[create](#)” on page 45
 - ❖ “[update](#)” on page 45
 - ❖ “[delete](#)” on page 46
 - ❖ “[list](#)” on page 46
- ◆ “[System Command Line Options](#)” on page 47
 - ❖ “[backup](#)” on page 47
 - ❖ “[backupSystem](#)” on page 49
 - ❖ “[backupList](#)” on page 50
 - ❖ “[recover](#)” on page 50
 - ❖ “[archiveLogBackup](#)” on page 51
 - ❖ “[restore](#)” on page 52
 - ❖ “[restoreSystem](#)” on page 53
- ◆ “[Storage Command Line Options](#)” on page 55
 - ❖ “[createSnapshot](#)” on page 55
 - ❖ “[split](#)” on page 56
 - ❖ “[createPartitionedStorage](#)” on page 56
 - ❖ “[delete](#)” on page 57
 - ❖ “[setPriorityForStorage](#)” on page 57
 - ❖ “[setQuotaForStorage](#)” on page 58
 - ❖ “[resize](#)” on page 58
 - ❖ “[copy](#)” on page 59
 - ❖ “[clone](#)” on page 59
 - ❖ “[volumeForExportList](#)” on page 60
 - ❖ “[listSnapshot](#)” on page 60
 - ❖ “[listVolume](#)” on page 60
 - ❖ “[restoreSnapshot](#)” on page 60
 - ❖ “[volumeList](#)” on page 61
 - ❖ “[createStorage](#)” on page 61

❖ “[modifyExports](#)” on page 61

To find out the version of SSC, use the `ssc version` command.

Credential Manager options

The `credential` command is related with managing and updating credential for hosts and storage system using SSC. It requires one of the following sub-commands to complete the operation:

- ◆ `create`
- ◆ `update`
- ◆ `list`
- ◆ `delete`

create

The `create` sub-command adds a credential within the credential persistence layer. For types `STORAGE` and `DFM` the target is the host name.

For type `MAXDB` the target is identified in the following way:

<DB_ID>@<host>:<port>: Here *host* is the hostname on which the database `DB_ID` is located and *port* is the port of the database.

```
ssc credential create
  [-clientRequestId <clientRequestId>]
  -type STORAGE|DFM|MAXDB
  -target <target>
  -credential
    -username <username>
    -password <password>
```

update

The `update` sub-command updates the credential within the credential persistence layer. For types `STORAGE` and `DFM` the target is the host name.

For type `MAXDB` the target is identified in the following way:

<DB_ID>@<host>:<port>: Here *host* is the hostname on which the database `DB_ID` is located and *port* is the port of the database.

```
ssc credential update
  [-clientRequestId <clientRequestId>]
  -type STORAGE|DFM|MAXDB
  -target <target>
  -credential
```

```
-username <username>  
-password <password>
```

delete

The `delete` sub-command deletes credential from the credential persistence layer. For types `STORAGE` and `DFM` the target is the host name.

For type `MAXDB` the target is identified in the following way:

<DB_ID>@<host>:<port>: Here *host* is the hostname on which the database `DB_ID` is located and *port* is the port of the database.

```
ssc credential delete  
  [-clientRequestId <clientRequestId>]  
  -type STORAGE|DFM|MAXDB  
  -target <target>
```

list

The `list` sub-command lists the persistent credentials of SSC.

```
ssc credential list  
  [-clientRequestId <clientRequestId>]  
  -type STORAGE|DFM|MAXDB
```

System Command Line Options

The `system` command is related with the backup, restore, and recover operations. This command requires one of the following sub-commands:

- ◆ `backup`
- ◆ `backupSystem`
- ◆ `backupList`
- ◆ `recover`
- ◆ `archiveLogBackup`
- ◆ `restore`
- ◆ `restoreSystem`

backup

The `backup` sub-command creates a backup of the identified SAP system.

```
ssc system backup
  [-clientRequestId <clientRequestId>]
  -sapSystem
    -applicationSystemId <applicationSystemId>
    -systemId <systemId>
    -systemName <systemName>
  -database
    -id <id>
    [-credentials
      -password <db_password>
      -username <db_username>]
    -server <server>
    [-oscredentials
      -password <os_password>
      -username <os_username>]
    -serverPort <serverPort>
    -databasePort <databasePort>
    -type <type>
  -storage
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]
  -backupProtectionProperties
    -protectionType <protectionType>
    -resourcePool <resourcePool>
    -retentionType <retentionType>
    -policy <policy>
```

```

-label <label>
[-mode OFFLINE_FORCE]
[-command
  [-pre <command>]
  [-post <command>]
  [-host <server>
  [-port <port>]]

```

Following is the description of the various command line options:

clientRequestId: Request ID for tracking the client request status.

applicationSystemId: SAP system ID.

systemName: The name of the SAP system.

database:

- ❖ **id:** The database ID.
- ❖ **credential:** The username and password for the database.
- ❖ **server:** Server hostname on which the SSC agent or server listens to requests.

Note

The parameter `oscredentials` is optional. This parameter is not used in case of a MaxDB backup.

- ❖ **serverPort:** Port on which the SSC agent receives the requests from the SSC server.
- ❖ **databasePort:** Port through which the database communicates.
- ❖ **type:** The database type.

storage:

- ❖ **storageSystem:** Hostname of the storage system.
- ❖ **volume:** Specifies the volume name.

backupProtectionProperties:

- ❖ **protectionType:** The type of protection of data.
- ❖ **resourcePool:** The resource pool used.
- ❖ **retentionType:** Data retention type used for data, for example, hourly, daily, weekly etc.
- ❖ **policy:** The protection policy used.

label: The label of the backup.

backupSystem

The `backupSystem` sub-command creates Snapshots with the name `<snapshot>` on the volumes as defined by mandatory `-storage` arguments. The database parameters determine the database which will be set to back up mode before taking the Snapshot and set back after.

The optional `-secondary` parameter determines the secondary storage location to which the qtrees in the primary volumes will be backed up. The primary and secondary retention determine which Snapshots matching the defined retention class are purged or not.

Retention periods are only applied in each retention class, e.g. retention period applied to one retention class does not affect lifetime of Snapshots in other retention class.

```
ssc system backupSystem
  [-clientRequestId <clientRequestId>]
  -database
    -id <id>
    -credentials
      -password <db_password>
      -username <db_username>
    -server <server>
  -oscredentials
    -password <os_password>
    -username <os_username>
  -serverPort <serverPort>
  -databasePort <databasePort>
  -type <type>
  -storage
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]
  [-secondary
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]]
  [[-primaryRetain <primaryRetain>]
  [-secondaryRetain <secondaryRetain>]
  -retentionUnit <retentionUnit>]
  [-retentionClass <retentionClass>]]
  -snapshot <snapshot>
```

Following is the description of the few parameters. For the description of the rest of the parameters, see section “[backup](#)” on page 47

primaryRetain: The number of retention units for primary.

secondaryRetain: The number of retention units for secondary.

retentionUnit: Gives the retention unit.

backupList

The `backupList` sub-command requests a list of available backups which are maintained by Protection Manager.

```
ssc system backupList
  [-clientRequestId <clientRequestId>]
  -storageSystem <storageSystem>
  -sapSystem
    -applicationSystemId <applicationSystemId>
    -systemId <systemId>
    -systemName <systemName>
```

recover

The `recover` sub-command recovers the primary database of a SAP system to one of the following states:

- ◆ Point in time of the backup
- ◆ Recovery of archive logs to a given point in time
- ◆ Complete recovery of all archive logs

Recovery is necessary after restoration of storage elements which are utilized by the SAP system database. Recovery should only be performed after a restore of database files is complete.

Recovery also requires that the archive logs relevant to the database have been made locally available to the host system before this service is requested.

The `recover` operation requires that the SAP system be offline. Before `recover` operation is requested, ensure that the archive logs relevant to the specific SAP system's database are locally available to the SAP host system.

```
ssc system recover
  [-clientRequestId <clientRequestId>]
  -storageSystem <storageSystem>
  -sapSystem
    -applicationSystemId <applicationSystemId>
    -systemId <systemId>
    -systemName <systemName>
  -database
    -id <id>
    -credentials
      -password <db_password>
      -username <db_username>
```

```

-server <server>
-oscredentials
-password <os_password>
-username <os_username>
-serverPort <serverPort>
-databasePort <databasePort>
-type <type>
-logfiles <logfiles> [<logfiles> ...]
[-recoverPoint <recoverPoint>|
-recoverAllLogs]

```

Following is the description of the few parameters related to the recover operation. For the description of the rest of the parameters, see section “[backup](#)” on page 47

logfiles: Name of the database log file.

recoverPoint: This is an optional parameter. If provided, its value indicates the database time at which the recovery should be complete. It accepts the date in the form of `yyyy-MM-dd HH:mm:ss z`; here `z` is the general time zone.

The parameter `stroageSystem` indicates to the service which Protection Manager owns the Protection Manager dataset as defined by `sapSystemId`.

archiveLogBackup

Archive log backup performs a back up of files that match the regular expression specified `inclusiveFileNameFilter` from a primary to a secondary storage location. Files or directories that do not reside within a qtree or that begin with "." are ignored by the backup. The backup retains the original qtree and directory structure of the files from the primary storage location when copying them to the secondary storage location.

If a file that was previously backed up is modified, then the file will again be backed up and the previous version on the secondary storage location will be saved in a subdirectory with the name of the current `operationId`.

lastLog setting set to SKIP: When the `lastLog` setting is set to `SKIP`, the files that have the latest modification time in each qtree are not backed up to the secondary storage location.

lastLog setting set to COPY: When the `lastLog` setting is set to `COPY`, the files that have the latest modification time in each qtree are backed up. Also, the files that have the latest modification time that have already been backed up before are backed up again.

Archive log backup optionally purges old files from the primary or secondary storage location based on the values of the primaryToRetain and secondaryToRetain settings. If primaryToRetain is not specified, then no files on the primary storage location are purged. If secondaryToRetain is not specified then no files on the secondary storage location are purged.

If retentionUnits is not specified, then the values specified by primaryToRetain or secondaryToRetain are assumed to be in DAYS. No files that exist on the primary storage location are purged if they do not exist on the secondary storage location. No files that exist on the secondary storage location are purged if they still exist on the primary storage location.

```
ssc system archiveLogBackup
  [-clientRequestId <clientRequestId>]
  -primaryLogs
    -storageSystem <primary_storageSystem>
    -volume <primary_volumeName>
  -secondaryLogs
    -storageSystem <secondary_storageSystem>
    -volume <secondary_volumeName>
  -inclusiveFileNameFilter <inclusiveFileNameFilter>
  -lastLogMethod <lastLogMethod>
  -primaryRetain <primaryRetain>
  -secondaryRetain <secondaryRetain>
  -retentionUnit <retentionUnit>
```

Following is the description of the parameters used in the archiveLogBackup operation:

inclusiveFileNameFilter: Regular expression to filter file names, to be considered in the corresponding operation.

lastLogMethod: Indicates how to deal with the latest log files on the primary or secondary storage.

restore

Restore operation restores (from a backup) all or portions of the data from storage for a SAP system. Restore operation only restores the file system, it does not interact with the database. If you want level recovery after file system recovery, use the recover service method.

Restore will cause the storage file system to be reset. You need to ensure that the SAP system which utilized this storage is offline before this service is used. If the SAP system is not offline, then the behavior may be unpredictable because the underlying file system is restored and can interfere with a clean restoration.

The optional parameter `backupNode` may be used to force the restore from a specific node within the policy. This can be used, for example, to perform a restore to alternate location across storage systems from the backup node even if there is a valid Snapshot on the primary node. By default the back up is carried out from the primary node if there is a valid backup or Snapshot.

Using SSC you can restore the root qtree also. For restoring, the root partitions should be specified with `-partition "/"`.

Note

This option is only available if Protection Manager is configured.

```
ssc system restore
  [-clientRequestId <clientRequestId>]
  -sapSystem
    -applicationSystemId <applicationSystemId>
    -systemId <systemId>
    -systemName <systemName>
  -storage
    -partition <partition> [<partition> ...]
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]
  [-destination
    -partition <partition> [<partition> ...]
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]]
  -label <label>
  [-node <backupNode>]
```

restoreSystem

This command restores a volume from a Snapshot from primary volume. If the Snapshot is not available on the primary volume, it restores it from a secondary volume. This command does not require Protection Manager to be configured.

```
ssc system restoreSystem
  [-clientRequestId <clientRequestId>]
  -storage
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]
  -secondary
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]
  -snapshot <snapshot>
```

For a description of the parameters for `restoreSystem` command, see “[backup](#)” on page 47.

Storage Command Line Options

The `storage` command is related with creating and deleting Snapshot, creating and resizing storage, splitting, cloning and so on. It works with the following sub-commands:

- ◆ `createSnapshot`
- ◆ `deleteSnapshot`
- ◆ `split`
- ◆ `createPartitionedStorage`
- ◆ `delete`
- ◆ `setPriorityForStorage`
- ◆ `setQuotaForStorage`
- ◆ `resize`
- ◆ `copy`
- ◆ `clone`
- ◆ `volumeForExportList`
- ◆ `listSnapshot`
- ◆ `listVolume`
- ◆ `restoreSnapshot`
- ◆ `volumeList`
- ◆ `createStorage`
- ◆ `modifyExports`

createSnapshot

This operation creates one or more Snapshot copies. This operation will create either all or none of the requested Snapshot storage. For this operation to succeed, the specified storage volume must already exist. An error occurs if any of the named storage Snapshots already exist.

```
ssc storage createSnapshot
  [-clientRequestId <clientRequestId>]
  -storage
    -snapshot <snapshot> [<snapshot> ...]
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]e
```

Following is the description of the parameters used in the `createSnapshot` command:

snapshot: Specifies the Snapshot name.

storagesystem: Specifies the host name of the storage system.

volume: Specifies the volume name.

deleteSnapshot

This sub-command deletes a Snapshot.

```
ssc storage deleteSnapshot
  [-clientRequestId <clientRequestId>]
  -storage
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]
    -snapshot <snapshot> [<snapshot> ...]
```

split

This sub-command splits an existing clone volume, or ensures that the volume does not rely on underlying shared storage with any other volume. The namespace volume must exist as a storage volume. This operation succeeds when the specified storage volume no longer shares underlying storage. Cancelling this operation during execution discontinues any further storage-level split activity; but it ensures that the split activity already performed remains complete.

```
ssc storage split
  [-clientRequestId <clientRequestId>]
  -storage
    -storageSystem <storageSystem>
    -volume <volumeName>
```

For a description of the parameters of `split` command, see “[createSnapshot](#)” on page 55.

createPartitionedStorage

This sub-command creates one or more quota-level partitions within a volume of storage. This operation creates either all or none of the requested storage. For this operation to succeed, the volume of storage must already exist. An error occurs if any of the named storage partitions already exist. There is no partial success for this operation

```
ssc storage createPartitionedStorage
  [-clientRequestId <clientRequestId>]
```

```

-storage
  -storageSystem <storageSystem> [<storageSystem> ...]
  -volume <volumeName> [<volumeName> ...]
  -partition <partition> [<partition> ...]

```

Following is a description of the parameter used in the `createPartitionedStorage` command. For a description of remaining parameters, see, “[createSnapshot](#)” on page 55.

delete

This operation deletes one or more volumes of storage. For this operation to succeed volumes must exist within the storage system. Once the volume is deleted, it cannot be recovered. The final operation result warns you about any volume that shows an error during the deletion process. However, the operation will be successful.

```

ssc storage delete
  [-clientRequestId <clientRequestId>]
  -storage
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]

```

For a description of the parameters of delete command, see “[createSnapshot](#)” on page 55.

setPriorityForStorage

This sub-command assigns `PriorityProperties` to an element of partitioned storage.

```

ssc storage setPriorityForStorage
  [-clientRequestId <clientRequestId>]
  -priorityProperties
  -storage
    -storageSystem <storageSystem>
    -volume <volumeName>
  -prioritySystem <prioritySystem>
  -priorityDefault
  -priorityLevel <priorityLevel>

```

setQuotaForStorage **e**

This sub-command assigns quota properties to an element of a partitioned storage.

```
ssc storage setQuotaForStorage
  [-clientRequestId <clientRequestId>]
  -quotaProperties
    [-diskLimitHard
      -amount <hardamount>
      -unit <hardunit>]
    [-diskLimitSoft
      [-amount <softamount>
        -unit <softunit>]
      [-diskLimitThreshold
        [-amount <thresholdamount>]
        [-unit <thresholdunit>]
        [-diskLimitSoftUnlimited]
        [-diskLimitThresholdUnlimited]
        [-diskLimitHardUnlimited]
        [-filesLimitHard <filesLimitHard>]
        [-filesLimitSoft <filesLimitSoft>]
        [-filesLimitSoftUnlimited]
        [-filesLimitHardUnlimited]]]]
  -storage
    [-partition <partition>]
    -storageSystem <storageSystem>
    -volume <volumeName>
  -target <target>
  -quotaType <quotaType>
```

resize

This sub-command re-sizes the bytes of storage allocated or file count allowed within a volume of storage. To re-size the bytes of storage, either an absolute or the difference in size may be specified. Similarly to change the total file count, either an absolute or difference in file count may be specified.

```
ssc storage resize
  [-clientRequestId <clientRequestId>]
  -storage
    -storageSystem <storageSystem>
    -volume <volumeName>
  -absoluteSize
    -amount <amount>
    -unit <unit>
  -deltaSize
    -amount <amount>
    -unit <unit>
```

```
-absoluteFileCount <absoluteFileCount>
-deltaFileCount <deltaFileCount>
```

Following is the description of the parameters used in the resize command:

absolute size:

- ❖ **amount:** Number of units, usually expressed in relation with unit
- ❖ **unit:** Unit of data
- ❖ **deltasize:** Relative change in size.
- ❖ **absoluteFileCount:** Total number of files allowed in the volume.
- ❖ **deltaFileCount:** Relative change from current file count allowed that should be applied.

copy

This sub-command creates a new volume which is a copy of a specified source volume. This copied volume exists in the name-space of a different storage system. For this operation to succeed, the source storage volume must already exist but the destination volume with the same name should not. You cannot access the copied volume until the copy operation completes. There is no partial success for this operation.

```
ssc storage copy
  [-clientRequestId <clientRequestId>]
  -copyPair
    -destinationAggregate <destinationAggregate>
    -source
      -storageSystem <src_storageSystem>
      -volume <src_volumeName>
    -destination
      -storageSystem <dest_storageSystem>
      -volume <dest_volumeName>
```

Following is the description of the parameters used in the copy command:

copyPair:

- ❖ **destinationAggregate:** Destination aggregate.

clone

This operation creates one or more new volumes of storage which are a copy of a specified source volume. For this operation to succeed, the source storage volumes must already exist but the destination volume with the same name should not. In the cloning operation, both the source and destination volumes must be within the same storage system.

```

ssc storage clone
  [-clientRequestId <clientRequestId>]
  -clones
    -storageSystem <storageSystem> [<storageSystem> ...]
    -destinationVolume <destinationVolume> [<destinationVolume>
...]
  [-sourceSnapshot <sourceSnapshot> [<sourceSnapshot> ...]]
  -sourceVolume <sourceVolume> [<sourceVolume> ...]

```

volumeForExportList

This sub-command lists volume for a given export. This command does not validate the existence for the given export. Thus this sub-command may be used before cloning to determine the volume name for the target volume, based on the desired export; even if the volume and the export may not yet exist.

```

ssc storage volumeForExportList
  [-clientRequestId <clientRequestId>]
  -storageExport
    -storageSystem <storageSystem>
    -exportName <exportName>

```

listSnapshot

This sub-command lists the Snapshots.

```

ssc storage listSnapshot
  [-clientRequestId <clientRequestId>]
  -storage
    -storageSystem <storageSystem>
    -volume <volumeName>

```

listVolume

This sub-command lists volume(s) on a storage system. The sub-command also provides additional information such as used space, available space, reserve, clone parent, and clone children.

```

ssc storage listVolume
  [-clientRequestId <clientRequestId>]
  -storageSystem <storageSystem>
  [-volume <volume>]

```

restoreSnapshot

This sub-command restores a snapshot.

```
ssc storage restoreSnapshot
  [-clientRequestId <clientRequestId>]
  -storage
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]
    -snapshot <snapshot> [<snapshot> ...]
```

volumeList

This sub-command lists the volume(s) on a storage system.

```
ssc storage volumeList
  [-clientRequestId <clientRequestId>]
  -storageSystem <storageSystem>
  [-volume <volume>]
```

createStorage

This sub-command creates one or more new volumes on storage system. Storage is created based on the data provided in storageSpecifications. This operation creates either all or none of the requested storage. An error occurs if any of the named storage volumes already exist. There is no partial success for this operation.

```
ssc storage createStorage
  [-clientRequestId <clientRequestId>]
  -storage
    -aggregate <aggregate> [<aggregate> ...]
    -languageCode <languageCode> [<languageCode> ...]
    -reserve <reserve> [<reserve> ...]
    -size
      -amount <amount> [<amount> ...]
      -unit <unit> [<unit> ...]
    -storageSystem <storageSystem> [<storageSystem> ...]
    -volume <volumeName> [<volumeName> ...]
```

modifyExports

This sub-command configures the NFS export settings of a storage volume.

```
ssc storage modifyExports
  [-clientRequestId <clientRequestId>]
  -exportProperties
    -storage
      -partition <partition>
      -storageSystem <storageSystem>
```

```
    -volume <volumeName>
-anonymous <anonymous>
-noSuid
-readOnlyACL <readOnlyACL> [<readOnlyACL> ...]
-readOnlyType <readOnlyType> -readWriteACL <readWriteACL>
[<readWriteACL> ...]
-readWriteType <readWriteType>
-rootACL <rootACL> [<rootACL> ...]
-rootType <rootType>
-security <security>
```

About this chapter

This chapter describes work-around for some problems with SSC 2.0

When CrashDB is down, "ssc system backup" gives an error message

When `jdbc.use.embedded=false` in `csm.config` file and `ssc` system backup operation is carried out with CrashDB down, SSC gives the following error message:

```
[ERROR] CLI-00052: Could not invoke remote operation.
```

In such a case, you should check the CrashDb database state. If it is in offline state, then make it online using the following command:

```
/opt/sapdb/programs/bin/dbmcli -d csmcrash -u dbm,dbm db_online
```

Backup with protectionType LOGICAL fails with error message "There is no dp policy"

If you get this error, ensure that for LOGICAL backup also you should have a policy and a resource pool defined. This is required because SSC needs a resource pool for creation of datasets

The log file split feature

Splitting log file: Storage Services Connector generates a server log file for various storage operations. This log file is generated every day and contains the log of all the storage operations that were carried out that day.

The log file split feature helps you create different log files for individual storage operations. For example, if the log file has entries for three different operations then the log file split feature will create three log files, one for each operation. Creating separate log files for individual operations helps you to troubleshoot an error easily.

The following is the syntax of the log split command:

```
bin/logsplit server log file output directory
```

server log file is the path to the server log file. The default value is "-". If the value is "-", then the command will read from the standard input.

output directory is the directory where the operation-specific log files are written. This directory must exist before executing the command. The default directory for the command output is the current directory.

The following is an example of the above command:

```
bin/logsplit log/csm_server.log operationLogs
```

The name of the created split log file follows the following format:

operation name_operation id[_<client request ID>].log

A new server log file is created everyday at 12 AM. If an operation starts before and continues after 12 AM, the operation is split into two different log files.

Combining split log files: You can also combine log files that are split using the log file split feature. The following example shows how to combine multiple operation-specific log files into a single log file named operationLogs.

```
cat log/csm_server.log.2007-01-01 log/csm_server.log.2007-01-02  
log/csm_server.log | bin/logsplit "-" operationLogs
```

It is important that the log files are specified chronologically, otherwise the resulting log files will contain entries that are out of order..

Glossary

SID	Abbreviation for System ID, an Identifier for a SAP system and its database. Similar to ORACLE_SID.
BPP	Business Process Platform.
SSC	Storage Services Connector server, a component of the BPP architecture.
Provider	A provider is a role within the BPP. It represents a Hosting Provider of SAP application services.
SAP system	A SAP system is a client of a Hosting Provider and refers to the SAP application services and resources configured to provide those services.
Master Template	A Provider defines Master Templates which are used (via cloning) to initiate a new SAP system environment.

Index

A

- administrators
 - accounts 5
 - default accounts 5
- archive logs 36
- audience, intended for this book vii

C

- contexts 17
- conventions
 - command vii
 - formatting viii
 - keyboard viii
- crash database 15
- crash-database 12
- credentials 34

D

- data set 22
- DataFabric Manager 31

G

- GlobalDataProtection 32
- GlobalRestore 32
- GlobalWrite 32

I

- Install SSC 12
- Installing Host System Agent 18

L

- Lag warnings 24
- Licensing 11

- log file split 63

O

- Operations Manager 5

P

- policy 23
- Protection Manager 5

R

- RBAC 31
- resource pool 22
- Retention duration 23

S

- special messages ix
- SSC agent 4
- SSC server 3
- Storage Services Connector 2

T

- Tenant backup management 38
- Tenant storage provisioning 41
- terminology vii

V

- Version compatibility requirements 10

W

- WSDL 35

