



Sybase, Inc.
One Sybase Drive
Dublin, CA 94568
www.sybase.com

Sybase Avaki EII Command Reference

Release 7.0

August 24, 2006

This publication pertains to Sybase software and to any subsequent release until otherwise indicated in new editions or technical notes. Information in this document is subject to change without notice. The software described herein is furnished under a license agreement, and it may be used or copied only in accordance with the terms of that agreement.

To order additional documents, U.S. and Canadian customers should call Customer Fulfillment at (800) 685-8225, fax (617) 229-9845.

Customers in other countries with a U.S. license agreement may contact Customer Fulfillment via the above fax number. All other international customers should contact their Sybase subsidiary or local distributor. Upgrades are provided only at regularly scheduled software release dates. No part of this publication may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical, or otherwise, without the prior written permission of Sybase, Inc.

Sybase, SYBASE (logo), ADA Workbench, Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise, Adaptive Server Enterprise Monitor, Adaptive Server Enterprise Replication, Adaptive Server Everywhere, Advantage Database Server, Afaia, Answers Anywhere, Applied Meta, Applied Metacomputing, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-Translator, APT-Library, ASEP, Avaki, Avaki (Arrow Design), Avaki Data Grid, AvantGo, Backup Server, BayCam, Beyond Connected, Bit-Wise, BizTracker, Certified PowerBuilder Developer, Certified SYBASE Professional, Certified SYBASE Professional Logo, ClearConnect, Client-Library, Client Services, CodeBank, Column Design, ComponentPack, Connection Manager, Convoy/DM, Copernicus, CSP, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DataWindow .NET, DB-Library, dbQueue, Dejima, Dejima Direct, Developers Workbench, DirectConnect Anywhere, DirectConnect, Distribution Director, Dynamic Mobility Model, e-ADK, E-Anywhere, e-Biz Integrator, E-Whatever, EC Gateway, ECMAP, ECRTP, eFulfillment Accelerator, EII Plus, Electronic Case Management, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise Portal (logo), Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, eProcurement Accelerator, eremote, Everything Works Better When Everything Works Together, EWA, ExtendedAssist, Extended Systems, ExtendedView, Financial Fusion, Financial Fusion (and design), Financial Fusion Server, Formula One, Fusion Powered e-Finance, Fusion Powered Financial Destinations, Fusion Powered STP, Gateway Manager, GeoPoint, GlobalFIX, iAnywhere, iAnywhere Solutions, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InstaHelp, Intelligent Self-Care, InternetBuilder, iremote, iScript, Jaguar CTS, jConnect for JDBC, KnowledgeBase, Legion, Logical Memory Manager, IRLite, M2M Anywhere, Mach Desktop, Mail Anywhere Studio, Mainframe Connect, Maintenance Express, Manage Anywhere Studio, MAP, M-Business Anywhere, M-Business Channel, M-Business Network, M-Business Suite, MDI Access Server, MDI Database Gateway, media.splash, Message Anywhere Server, MetaWorks, MethodSet, mFolio, Mirror Activator, ML Query, MobiCATS, MobileQ, MySupport, Net-Gateway, Net-Library, New Era of Networks, Next Generation Learning, Next Generation Learning Studio, O DEVICE, OASIS, OASIS logo, ObjectConnect, ObjectCycle, OmniConnect, OmniQ, OmniSQL Access Module, OmniSQL Toolkit, OneBridge, Open Biz, Open Business Interchange, Open Client, Open ClientConnect, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, Partnerships that Work, PB-Gen, PC APT Execute, PC DB-Net, PC Net Library, Pharma Anywhere, PhysicalArchitect, Pocket PowerBuilder, PocketBuilder, Power++, Power Through Knowledge, power.stop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, Powering the New Economy, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Pylon, Pylon Anywhere, Pylon Application Server, Pylon Conduit, Pylon PIM Server, Pylon Pro, QAnywhere, Rapport, Relational Beans, RemoteWare, RepConnector, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RFID Anywhere, RW-DisplayLib, RW-Library, SAFE, SAFE/PRO, Sales Anywhere, Search Anywhere, SDF, Search Anywhere, Secure SQL Server, Secure SQL Toolkit, Security Guardian, ShareLink, ShareSpool, SKILS, smart.partners, smart.parts, smart.script, SOA Anywhere Trademark, SQL Advantage, SQL Anywhere, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Edit/TPU, SQL Everywhere, SQL Modeler, SQL Remote, SQL Server, SQL Server Manager, SQL SMART, SQL Toolkit, SQL Server/CFT, SQL Server/DBM, SQL Server SNMP SubAgent, SQL Station, SQLJ, Stage III Engineering, Startup.Com, STEP, SupportNow, S.W.I.F.T. Message Format Libraries, Sybase Central, Sybase Client/Server Interfaces, Sybase Development Framework, Sybase Financial Server, Sybase Gateways, Sybase Learning Connection, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase Synergy Program, Sybase Virtual Server Architecture, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SybFlex, SybMD, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, The Enterprise Client/Server Company, The Extensible Software Platform, The Future Is Wide Open, The Learning Connection, The Model For Client/Server Solutions, The Online Information Center, The Power of One, TotalFix, TradeForce, Transact-SQL, Translation Toolkit, Turning Imagination Into Reality, UltraLite, UltraLite.NET, UNIBOM, Unilib, Uninull, Unisep, Unistring, URK Runtime Kit for UniCode, Viafone, Viewer, VisualWriter, VQL, WarehouseArchitect, Warehouse Control Center, Warehouse Studio, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web Deployment Kit, Web.PB, Web.SQL, WebSights, WebViewer, WorkGroup SQL Server, XA-Library, XA-Server, XcelleNet, XP Server, XTNDAccess and XTNDConnect are trademarks of Sybase, Inc. or its subsidiaries. 07/06

Unicode and the Unicode Logo are registered trademarks of Unicode, Inc.

All other company and product names used herein may be trademarks or registered trademarks of their respective companies.

Use, duplication, or disclosure by the government is subject to the restrictions set forth in subparagraph (c)(1)(ii) of DFARS 52.227-7013 for the DOD and as set forth in FAR 52.227-19(a)-(d) for civilian agencies.

Set in *Arial*, *Courier New*, and *Times New Roman*. Stanley Morison, the creator of *Times New Roman*, said of it: "By the vice of Mammon and the misery of the machine, it is bigoted and narrow, mean and puritan."

Credits

This product includes software developed by the Apache Software Foundation (<http://www.apache.org>). This product includes Hypersonic SQL and ANTLR. This product includes code licenses from RSA Security, Inc. Some portions licensed from IBM are available at <http://oss.software.ibm.com/icu4j/>. Contains IBM® 64-bit Runtime Environment for AIX™, Java™ 2 Technology Edition Version 1.4 Modules © Copyright IBM Corporation 1999, 2000 All Rights Reserved. Contains the SAXON XSLT Processor from Michael Kay, which is available at <http://saxon.sourceforge.net>. This product includes software developed by the Proxool Project (<http://proxool.sourceforge.net>).

Sybase Avaki EII Command Reference

Written by Cheryl Magadiou and Beth Thoenen

Sybase, Inc., One Sybase Drive, Dublin, CA 94568.

Table of contents

Preface xi

Organization **xii**

Related documentation and online help **xiii**

Conventions **xv**

 Command syntax conventions **xv**

 Conventions for screen examples **xvi**

How to contact Avaki support at Sybase, Inc. **xvii**

Chapter 1

Using the CLI 1

Before using the CLI **1**

When to use the CLI **2**

Accessing the CLI **2**

Chapter 2

Starting and stopping Avaki servers 5

Automatically restarting servers and GDCs **5**

`dgas --register` **6**

`dgas --start` **7**

`dgas --stop` **8**

`dgas --unregister` **9**

`grid-server --register` **9**

`grid-server --start` **11**

`grid-server --stop` **11**

`grid-server --unregister` **12**

`proxy-server --register` **13**

`proxy-server --start` **14**

`proxy-server --stop` **14**

`proxy-server --unregister` **15**

share-server --register 16
share-server --start 17
share-server --stop 17
share-server --unregister 18

Chapter 3

Command reference 19

avaki attribute --delete 19
avaki attribute --list 21
avaki attribute --update 21
avaki backup 24
avaki cache --evict 25
avaki cache --evict --all 26
avaki cache --evict --deleted 26
avaki cache --get 27
avaki cache --invalidate 28
avaki cache --invalidate --all 29
avaki cache --invalidate-dataservice-results 30
avaki cache --invalidate-dbop-results 31
avaki cache --list 32
avaki cache --refresh 33
avaki cache --set 34
avaki cache --unset 35
avaki cat 35
avaki categories --add-to-category 36
avaki categories --create 37
avaki categories --delete 37
avaki categories --describe 38
avaki categories --list 38
avaki categories --remove-from-category 39
avaki categories --set-description 40
avaki categories --show categories 41
avaki cd 41
avaki chmod 42
avaki chown 45
avaki client 46
avaki cp 47

avaki dataservice --add-schedule 48
avaki dataservice --create 52
avaki dataservice --delete 54
avaki dataservice --delete-schedule 54
avaki dataservice --depends 55
avaki dataservice --execute 56
avaki dataservice --generate-sql-view 57
avaki dataservice --info 58
avaki dataservice --list-schedules 58
avaki dataservice --update 59
avaki dbconn --allow-dbop-creation 61
avaki dbconn --delete 62
avaki dbconn --disallow-dbop-creation 63
avaki dbconn --execute 64
avaki dbconn --info 66
avaki dbconn --jdbc 66
avaki dbconn --provision-tables 72
avaki dbconn --test 73
avaki dbconn --show-tables 73
avaki dbop --add-schedule 74
avaki dbop --delete 78
avaki dbop --delete-schedule 79
avaki dbop --depends 80
avaki dbop --execute 81
avaki dbop --generate-sql-view 82
avaki dbop --info 83
avaki dbop --jdbc 84
avaki dbop --jdbc --create-virtual-dbop 88
avaki dbop --list-schedules 91
avaki dgas --add-admission-policy 92
avaki dgas --add-group-mapping 94
avaki dgas --add-user-mapping 95
avaki dgas --cifs-share-info 95
avaki dgas --clear-cached-credentials 96
avaki dgas --create-cifs-share 97
avaki dgas --delete-admission-policy 98

avaki dgas --delete-cache **99**
avaki dgas --delete-cifs-share **99**
avaki dgas --delete-group-mapping **100**
avaki dgas --delete-user-mapping **100**
avaki dgas --disconnect-cifs-client **101**
avaki dgas --get-cache-size **102**
avaki dgas --get-cache-statistics **102**
avaki dgas --get-free-disk-space **103**
avaki dgas --get-properties **103**
avaki dgas --get-property **104**
avaki dgas --get-property-list **104**
avaki dgas --initialize **105**
avaki dgas --list-cifs-clients **105**
avaki dgas --list-cifs-shares **106**
avaki dgas --list-group-mappings **106**
avaki dgas --list-user-mappings **107**
avaki dgas --modify-cifs-share **107**
avaki dgas --read-log-properties **108**
avaki dgas --reset-cache-statistics **109**
avaki dgas --save-cache **109**
avaki dgas --self-map **110**
avaki dgas --self-unmap **111**
avaki dgas --set-admission-policy-domain **112**
avaki dgas --set-default-admission-policy **113**
avaki dgas --set-property **114**
avaki dgas --show-admission-policies **114**
avaki dgas --show-admission-policy-domain **115**
avaki dgas --sync-cache **116**
avaki dgas --unset-admission-policy-domain **117**
avaki dgas --unset-property **117**
avaki directory --add-schedule **118**
avaki directory --cache **122**
avaki directory --delete-schedule **126**
avaki directory --do-not-cache **126**
avaki directory --list-schedules **127**
avaki domain --create **127**

avaki domain --disconnect **128**
avaki domain --info **128**
avaki domain --interconnect **129**
avaki executionservice --info **129**
avaki executionservice --set **130**
avaki file --add-schedule **131**
avaki file --cache-on-demand **135**
avaki file --delete-schedule **136**
avaki file --do-not-cache **136**
avaki file --list-schedules **137**
avaki file --pin **138**
avaki group --add --user **141**
avaki group --create **143**
avaki group --delete **144**
avaki group --delete --user **145**
avaki group --info **147**
avaki group --list-user **148**
avaki help **149**
avaki id **149**
avaki ldap --add-schedule **150**
avaki ldap --delete **154**
avaki ldap --delete-schedule **155**
avaki ldap --import **155**
avaki ldap --info **157**
avaki ldap --integrate **158**
avaki ldap --list-schedules **160**
avaki ldap --searchbase **160**
avaki ldap --update **161**
avaki ln **163**
avaki locks --clear **164**
avaki locks --list **164**
avaki login **165**
avaki logout **166**
avaki ls **167**
avaki mkdir **167**
avaki monitor --add **168**

avaki monitor --clear **170**
avaki monitor --create **171**
avaki monitor --delete **172**
avaki monitor --list **172**
avaki monitor --result **173**
avaki monitor --start **174**
avaki monitor --stop **176**
avaki mv **177**
avaki nis --add-schedule **178**
avaki nis --delete **181**
avaki nis --delete-schedule **182**
avaki nis --import **183**
avaki nis --info **184**
avaki nis --integrate **184**
avaki nis --list-schedules **185**
avaki nis --update **185**
avaki passwd **186**
avaki permissions **187**
avaki plugin --generate **187**
avaki proxy --add **191**
avaki proxy --delete **192**
avaki proxy --list **193**
avaki pwd **193**
avaki replica --add **193**
avaki replica --config **194**
avaki replica --delete **194**
avaki replica --info **195**
avaki replica --synch **195**
avaki rm **196**
avaki schedule --delete **197**
avaki schedule --info **197**
avaki schedule --list **198**
avaki schedule --print-iterations **198**
avaki scheduleexclusion --create --custom **199**
avaki scheduleexclusion --create --daily **201**
avaki scheduleexclusion --create --monthly **203**

avaki scheduleexclusion --create --weekly **205**
avaki scheduleexclusion --create --yearly **207**
avaki scheduleexclusion --delete **208**
avaki scheduleexclusion --info **209**
avaki scheduleexclusion --list **209**
avaki search --create **210**
avaki search --delete **211**
avaki search (execute) **212**
avaki search --get-rehash-interval **214**
avaki search --info **215**
avaki search --rehash **215**
avaki search --set-rehash-interval **216**
avaki security --config **217**
avaki security --default-gid **218**
avaki security --default-group **219**
avaki security --default-uid **220**
avaki security --default-user **221**
avaki security --gid **222**
avaki security --info **222**
avaki security --uid **223**
avaki server --dgas --connect **224**
avaki server --dgas --destroy **225**
avaki server --dgas --stop **225**
avaki server --grid --connect **226**
avaki server --grid --destroy **227**
avaki server --grid --stop **228**
avaki server --proxy **228**
avaki server --share --connect **229**
avaki server --share --disconnect **230**
avaki server --share --stop **231**
avaki share --add-rehash-schedule **232**
avaki share --add-share-servers **235**
avaki share --create **236**
avaki share --delete-rehash-schedule **238**
avaki share --disconnect **238**
avaki share --get-local-path **239**

avaki share --get-status **239**
avaki share --list-rehash-schedules **240**
avaki share --list-share-servers **240**
avaki share --rehash **241**
avaki share --remove-share-servers **241**
avaki share --set-local-path **242**
avaki share --set-share-servers **243**
avaki share --set-status **244**
avaki share --update-share-servers **245**
avaki shell **246**
avaki sqlview --delete **246**
avaki sqlview --get-description **247**
avaki sqlview --set-description **248**
avaki status **249**
avaki upgrade **250**
avaki upgrade --info **250**
avaki user **251**
avaki user --create **252**
avaki user --db-mapping --add **253**
avaki user --db-mapping --delete **255**
avaki user --db-mapping --list **257**
avaki user --delete **258**
avaki user --info **258**
avaki user --list-group **259**
avaki view --add-schedule **260**
avaki view --create --database **264**
avaki view --create --data-service **267**
avaki view --create --file **270**
avaki view --delete **272**
avaki view --delete-schedule **272**
avaki view --depends **273**
avaki view --garbage-collect **273**
avaki view --generate **274**
avaki view --info **274**
avaki view --list-schedules **275**
avaki view --set-property **275**

avaki view --update **280**
avaki virtualdatabase --allow-dbop-creation **281**
avaki virtualdatabase --disallow-dbop-creation **282**
avaki virtualdatabase --execute **283**
avaki virtualdatabase --show-tables **284**
avaki virtualeschema --deploy **286**
avaki virtualeschema --undeploy **286**
avaki --whoami **287**

Appendix A *Attribute reference* **289**

Avaki grid server attributes **290**
File attributes **290**
Cache service attributes **291**

Appendix B *Setting up DGAS properties files* **293**

About the properties file **293**
Properties file format **294**

Appendix C *Configuring cron schedules* **297**

Appendix D *Data type mapping files* **301**

Format and contents of type mapping files **302**
Setting the source data type: data type inconsistencies **304**
Logging of type mapping **305**

Glossary **307**

Preface

This *Sybase Avaki EII Command Reference* describes commands in the Sybase Avaki EII command line interface (CLI). The CLI is an alternative to Avaki Studio and to Avaki's web user interface. In addition, the CLI provides ways to perform certain actions that cannot be performed in Avaki Studio or through the web user interface.

Avaki supports the following types of commands:

- CLI commands. These commands all begin with “avaki” or are accessible within an Avaki shell, and are available on grid servers and command clients; and
- Avaki server start/stop commands. These commands do not begin with “avaki” and are not accessible within an Avaki shell. Start/stop commands are available on any machine that has an appropriate Avaki software installation.

Note This book and the product's user interfaces refer to Sybase Avaki EII software as *Avaki* or *Avaki Data Grid*.

Organization

This book is organized as follows:

Chapter 1 Using the CLI	Shows how to access the Avaki CLI and when to use it.
Chapter 2 Starting and stopping Avaki servers	Describes the commands for starting and stopping Avaki grid, DGAS, and proxy, share servers.
Chapter 3 Command reference	Explains how to use Avaki CLI commands.
Appendix A Attribute reference	Describes attributes you can set with the avaki attribute commands.
Appendix B Setting up DGAS properties files	Provides a sample properties file for data grid access servers.
Appendix C Configuring cron schedules	Describes the values that may be used to create a cron specification to indicate the schedule by which various tasks may be performed.
Appendix D Data type mapping files	Describes the format and contents of files you can use to map data types for provisioned SQL views.
Glossary	Defines terms used in this guide.

Related documentation and online help

Manuals

These manuals make up the Avaki documentation set:

- *Sybase Avaki EII Overture*
- *Sybase Avaki EII Administration Guide* (includes installation instructions)
- *Data Integration with Sybase Avaki Studio*
- *Sybase Avaki EII Provisioning and Advanced Data Integration Guide*
- *Sybase Avaki EII API Guide*
- *Sybase Avaki EII Command Reference*

The manuals are included, in PDF format, on the CD with the Avaki software. They are stored in the docs subdirectory of the Avaki installation directory.

To access the manuals via Avaki's web user interface, log in to your Avaki domain and click the **Help** link at the top right corner of any page of the web UI.

Online help

In addition to the manuals, Avaki provides online help for commands.

To display a list of Avaki commands with brief descriptions, log in to Avaki and enter **avaki help**:

```
% avaki help
List of domain commands:
attribute
backup
cache
cat
categories
cd
chmod
chown
client
cp
dataservice
dbconn
dbop
dgas
```

```
directory
domain
executionservice
file
group
help
id
ldap
ln
locks
login
logout
ls
mkdir
monitor
mv
nis
passwd
permissions
plugin
proxy
pwd
replica
rm
scheduleexclusion
search
security
server
share
shell
sqlview
status
systemproperty
upgrade
user
view
virtualdatabase
virtualschema
whoami
```

To display a description of a particular command and the syntax, enter a command of the form **avaki help** *<command>*. For example:

```
% avaki help mv
usage: avaki mv <source-grid-path> <target-grid-path>
Description: Move or rename a grid directory or a file
in a grid directory. Similar to the Unix mv command.
```


Conventions

This section describes text conventions used in this guide to represent elements of commands and screen displays.

Command syntax conventions

This table describes conventions that this book uses in command syntax statements. The “Enter this” column tells you whether you need to enter the characters when you type a command. The examples in the “Examples” column are not necessarily complete commands.

Convention	Description	Enter this?	Examples
[]	Square brackets surround optional arguments.	no	<code>avaki login [<user-id>]</code>
{ }	Curly brackets surround groups of required arguments.	no	<code>avaki chmod {--allow --deny --unset}</code>
{ } []	Vertical bars separate alternative options within square or curly brackets. If the brackets are square, you need not enter any of the options; if the brackets are curly, you must choose one of the options.	no	<code>avaki backup {--snapshot --recover}</code> <code>avaki share --create [--background --bg]</code>
< >	Angle brackets surround placeholder arguments that you must replace with a value such as a path or file name. Square brackets outside the angle brackets indicate that the placeholder is optional.	no	<code>avaki help [<command-name>]</code> <code>avaki help share</code>
*	An asterisk follows an argument that can be entered zero or more times.	no	<code>avaki plugin --generate [--input=<stream-spec>*]</code> <code>avaki plugin --generate --input="name=input1,type=XML"</code>
+	A plus sign follows an argument that can be entered one or more times. Use spaces to separate the values.	no	<code>avaki cat <grid-path>+</code> <code>avaki cat /home/fred/file1 /home/fred/file2</code>

Convention	Description	Enter this?	Examples
-	Enter a hyphen or minus sign before a single-letter command option.	yes	avaki mkdir -p
=	Enter an equal sign before the value of an option.	yes	avaki login --auth-service= <auth-service-name>
(space)	A space separates multiple arguments.	yes	avaki cat file1 file2 file3

Conventions for screen examples

This table describes conventions this book uses in examples of user input and system output.

Convention	Description	Example
\$ or C:>	The command prompt	\$
<>	A placeholder; replace the text inside the brackets with an option or value	\$ avaki ls <grid-path>
screen font	Text that appears on the screen	sample text
bold screen font	User input—commands that you enter	\$ avaki ls

How to contact Avaki support at Sybase, Inc.

For general information about Sybase technical support, see the *Customer Service Reference Guide* at

<http://www.sybase.com/support/aboutsupport/guide/csrg>

Please contact us with any questions or difficulties you encounter.

By telephone

In North America, call toll free: 1-800-8SYBASE

Outside North America, follow the link below to see a list of Sybase offices and phone numbers around the world.

<http://www.sybase.com/contactus/support>

On the web

If you are a designated contact for a technical support plan, you can log and track cases on the web using the Case Express application. At <http://www.sybase.com>, mouse over the **Support and Services** tab and select **Case Management** from the drop-down list. Use the email address and password for your mysybase account to log in.

Using the CLI

This chapter discusses when the Avaki command line interface (CLI) is useful and how to access it. The chapter is organized as follows:

- [“Before using the CLI,”](#) below
- [“When to use the CLI”](#) on page 2
- [“Accessing the CLI”](#) on page 2

Before using the CLI

This *Sybase Avaki EII Command Reference* assumes that you know about the components of Avaki data grids. These include database tools such as data services, database operations, and SQL views; Avaki’s access control and caching schemes; and the servers that make up an Avaki domain—grid domain controllers, grid servers, share servers, etc. If you’re not familiar with Avaki data grids, read the *Sybase Avaki EII Overture* before you use the commands.

When to use the CLI

Most Avaki tasks are easiest to perform using Avaki Studio or the web user interface. However, you can often save time on frequently performed, repetitious, or complex tasks if you use the CLI for scripts.

Accessing the CLI

The Avaki CLI is installed automatically when grid servers (including grid domain controllers) are installed. You can install the Avaki CLI on any Avaki server by running the installer and choosing to install the command line client feature.

Note The start and stop commands for Avaki servers are not considered part of the CLI and are, of course, available on the servers. For information on these commands, see [Chapter 2, “Starting and stopping Avaki servers”](#).

Before you can use the CLI, Avaki software must be installed and started on at least one machine, your grid domain controller (GDC). If you plan to use the CLI on a grid server or a command-line client, an appropriate installation must be performed on the machine you plan to use. See the *Sybase Avaki EII Administration Guide* for installation instructions, or ask your Avaki domain administrator for assistance.

Some of the commands in this command reference pertain to administrative tasks that may be performed only by users who have administrative privileges.

In the Avaki CLI, relative paths are relative to the current working directory of the client session. To set the current working directory, use the DOS command **CD** or the Unix command **cd**.

Before using Avaki commands on any machine (including a machine running a GDC), follow these steps to connect to a GDC:

Step 1 Open a command window and enter a command of this form:

```
avaki client --connect <GDC-machine>
```

<GDC-machine> is the DNS name or IP address of a machine on which an Avaki GDC is installed and operating. For example:

```
$ avaki client --connect Flintstone
```

- Step 2** Log in to the domain. In this example, the user enters a password at the Password prompt, and at the next two prompts presses **Enter** to accept the defaults (DefaultAuthService and Grid service type):

```
$ avaki login Fred
Password:
Auth service [DefaultAuthService]:
Service type [(Grid), Nis, Ldap]:
Fred logged in successfully
$
```


Starting and stopping Avaki servers

This chapter describes the commands for starting and stopping Avaki grid, DGAS, proxy, and share servers. The commands are listed alphabetically.

The commands described in this chapter must be entered in the Avaki installation directory (on Windows, the default is C:\AvakiDataGridxx, where xx represents the release number).

Automatically restarting servers and GDCs

Avaki provides an automatic restart feature for Avaki servers and grid domain controllers. When auto-restart is enabled, the Avaki server or GDC restarts automatically when the computer restarts. If auto-restart is not enabled, you must bring the GDC or server back up manually each time its machine reboots. If you plan to use Avaki's auto-restart feature in Unix, you must be logged in to the machine as a user who has administrative privileges when you set up the GDC or server. If you plan to use Avaki's auto-restart feature in Windows, you must be logged in to the machine as a user who has permission to log on as a service. In Windows, when you start a server with auto-restart enabled, the server is registered as a service. You can find an entry for the service in the machine's Services list (Start > Control Panel > Administrative Tools > Services). In Unix, when you start a server with auto-restart enabled, the run command script for the server will be placed in the machine's /etc/rc/* directory.

dgas --register

Start a data grid access server with the automatic restart feature enabled.

Note In Windows, you can use Windows shortcuts to start Avaki servers with auto-restart enabled. To start a data grid access server with auto-restart enabled, open the Start menu, then select the **Avaki Data Grid** program group. Select **DGAS**, and then select the **Register Data Grid Access Server as Windows service** shortcut.

The server takes a minute or two to start.

Syntax

```
dgas --register [--user=<user>] [--name=<server-name>]
  [--db-path=<local-path>] [--cache-path=<local-path>]
  [--port=<connect-port>] [--rmi-port=<rmi-server-port>]
```

`--user=<user>` Optional. The name of the nonadministrative user account that you set up to run Avaki services. The account must have logon-as-service privileges.

For details about configuring user accounts, see the *Sybase Avaki EII Administration Guide*.

`--name=<server-name>` Optional. Specify a name for the data grid access server. If you don't specify a name for a new DGAS, the system uses the default name `Access1`.

`--db-path=<local-path>` Optional. Specify the location of the DGAS state database, which stores property settings and other information required to restart the DGAS.

Default: `<Avaki-install-directory>/DGAS/dgas_db`

`--cache-path=<local-path>` Optional. Specify the location of the local directory that holds all DGAS internal caches.

Default: `<Avaki-install-directory>/DGAS/cache`

- `--port=<connect-port>` Optional. Specify the connect port that the DGAS will use to communicate with other objects in the data grid.
Default value: 1399
- `--rmi-port=<rmi-server-port>` Optional. Specify the RMI server port that the DGAS will use.
Default value: 1599

dgas --start

Start a data grid access server (DGAS) without enabling the automatic restart feature. For information about enabling auto-restart, see [dgas --register on page 6](#).

The DGAS takes a minute or two to start.

Syntax

```
dgas --start [--name=<server-name>] [--db-path=<local-path>]
  [--cache-path=<local-path>] [--new] [--port=<connect-port>]
  [--rmi-port=<rmi-server-port>]
```

- `--name=<server-name>` Optional. Specify a name for the data grid access server. If you don't specify a name for a new DGAS, the system uses the default name.
- `--db-path=<local-path>` Optional. Specify the location of the DGAS state database, which stores property settings and other information required to restart the DGAS.
Default: <Avaki-install-directory>/DGAS/dgas_db
- `--cache-path=<local-path>` Optional. Specify the location of the local directory that holds all DGAS internal caches.
Default: <Avaki-install-directory>/DGAS/cache

<code>--new</code>	Optional. Wipe out the state database for the previous DGAS but retain any DGAS internal caches.
<code>--port=<connect-port></code>	Optional. Specify the connect port that the DGAS will use to communicate with other objects in the data grid. Default value: 1399
<code>--rmi-port=<rmi-server-port></code>	Optional. Specify the RMI server port that the DGAS will use. Default value: 1599

dgas --stop

Stop a data grid access server that is running on Unix. For information about disabling auto-restart, see [dgas --unregister](#) on page 9.

Note The procedure for stopping a data grid access server on Windows depends on whether auto-restart is enabled for the server, as follows:

- If auto-restart is not enabled on the server, enter **Ctrl-C** in the server's window to stop it.
- If auto-restart is enabled on the server, use the Windows Services list to stop and restart the server. Right-click on the server name (for example, Avaki-DGAS-1096608202) and select **stop** or **restart**.

Syntax

```
dgas --stop [--name=<server-name>]
```

<code>--name=<server-name></code>	Optional. Specify a name for the data grid access server. If you don't specify a name for a new DGAS, the system uses the default name.
---	---

dgas --unregister

Stop a data grid access server and disable the automatic restart feature.

Note In Windows, you can use Windows shortcuts to stop Avaki servers and disable auto-restart. To stop a data grid access server and disable auto-restart, open the Start menu, then select the **Avaki Data Grid** program group. Select **DGAS**, and then select the **Unregister Data Grid Access Server as Windows service** shortcut.

Syntax

```
dgas --unregister [--name=<server-name>]
```

`--name=<server-name>` Optional. Specify a name for the data grid access server. If you don't specify a name for a new DGAS, the system uses the default name.

grid-server --register

Start a grid server or grid domain controller with the automatic restart feature enabled.

Note In Windows, you can use Windows shortcuts to start Avaki servers with auto-restart enabled. To start a grid server or grid domain controller with auto-restart enabled, open the Start menu, then select the **Avaki Data Grid** program group. Select **Grid Server**, and then select the **Register Grid Server as Windows service** shortcut.

The grid server takes a minute or two to start. Startup is complete when you see a message similar to one of these:

On Unix:

```
bash-2.05# ./grid-server --register
User to run as [wilma]:
The Avaki-GridService1395600136 service is starting...
Starting Avaki-GridService1395600136 daemon:
```

The Avaki-GridService1395600136 service started. For more details, see
/home/local/wilma/AvakiDataGrid70/jboss/server/grid-server/
log/server.log

On Windows:

```
C:\AvakiDataGrid70> grid-server --register
User to run as [BEDROCK\wilma] :
Password to run as BEDROCK\wilma:
The Avaki-GridService-1096607241 service is starting.
The Avaki-GridService-1096607241 service was started
successfully.
```

Syntax

```
grid-server --register [--user=<user> [--password=<password>]]
```

--user=<user> Optional. The name of the nonadministrative user account that you set up to run Avaki services. The account must have logon-as-service privileges.

For details about configuring user accounts, see the *Sybase Avaki EII Administration Guide*.

--password=<password> Optional. The password for the user account specified in <user>.

grid-server --start

Start a grid server or grid domain controller without enabling the automatic restart feature. For information about enabling auto-restart, see [grid-server --register on page 9](#).

The grid server takes a minute or two to start. Startup is complete when you see a message similar to this:

```
2003-06-30 13:50:26,951 INFO
[org.jboss.system.server.Server, main] JBoss
(MX MicroKernel) [3.0.6 (CVSTag=JBoss_3_0_6
Date=200301260037)] Started in 1m:3s:212ms
```

Syntax

```
grid-server --start
```

– This command has no options.

grid-server --stop

Stop a grid server or grid domain controller that is running on Unix. For information about disabling auto-restart, see [grid-server --unregister on page 12](#).

Note The procedure for stopping a grid server or grid domain controller on Windows depends on whether auto-restart is enabled for the server, as follows:

- If auto-restart is not enabled on the server, enter **Ctrl-C** in the server's window to stop it.
- If auto-restart is enabled on the server, use the Windows Services list to stop and restart the server. Right-click on the server name (for example, Avaki-GridService-1096608202) and select **stop** or **restart**.

Syntax

```
grid-server --stop
```

– This command has no options.

grid-server --unregister

Stop a grid server or grid domain controller and disable the automatic restart feature.

Note In Windows, you can use Windows shortcuts to stop Avaki servers and disable auto-restart. To stop a grid server or grid domain controller and disable auto-restart, open the Start menu, then select the **Avaki Data Grid** program group. Select **Grid Server**, and then select the **Unregister Grid Server as Windows service** shortcut.

Syntax

```
grid-server --unregister
```

– This command has no options.

proxy-server --register

Start a proxy server with the automatic restart feature enabled.

Note In Windows, you can use Windows shortcuts to start Avaki servers with auto-restart enabled. To start a proxy server with auto-restart enabled, open the Start menu, then select the **Avaki Data Grid** program group. Select **Proxy Server**, and then select the **Register ProxyServer as Windows service** shortcut.

The proxy server takes a minute or two to start.

Syntax

```
proxy-server --register [--user=<user> [--password=<password>]]
```

`--user=<user>` Optional. The name of the nonadministrative user account that you set up to run Avaki services. The account must have logon-as-service privileges.

For details about configuring user accounts, see the *Sybase Avaki EII Administration Guide*.

`--password=<password>` Optional. The password for the user account specified in `<user>`.

proxy-server --start

Start a proxy server without enabling the automatic restart feature. For information about enabling auto-restart, see [proxy-server --register on page 13](#).

The proxy server takes a minute or two to start.

Syntax

```
proxy-server --start
```

– This command has no options.

proxy-server --stop

Stop a proxy server that is running on Unix. For information about disabling auto-restart, see [proxy-server --unregister on page 15](#).

Note The procedure for stopping a proxy server on Windows depends on whether auto-restart is enabled for the server, as follows:

- If auto-restart is not enabled on the server, enter **Ctrl-C** in the server's window to stop it.
- If auto-restart is enabled on the server, use the Windows Services list to stop and restart the server. Right-click on the server name (for example, Avaki-ProxyService-1096608202) and select **stop** or **restart**.

Syntax

```
proxy-server --stop
```

– This command has no options.

proxy-server --unregister

Stop a proxy server and disable the automatic restart feature.

Note In Windows, you can use Windows shortcuts to stop Avaki servers and disable auto-restart. To stop a proxy server and disable auto-restart, open the Start menu, then select the **Avaki Data Grid** program group. Select **Proxy Server**, and then select the **Unregister ProxyServer as Windows service** shortcut.

Syntax

```
proxy-server --unregister
```

– This command has no options.

share-server --register

Start a share server with the automatic restart feature enabled.

Note In Windows, you can use Windows shortcuts to start Avaki servers with auto-restart enabled. To start a share server with auto-restart enabled, open the Start menu, then select the **Avaki Data Grid** program group. Select **Share Server**, and then select the **Register ShareServer as Windows service** shortcut.

The share server takes a minute or two to start.

Syntax

```
share-server --register [--user=<user> [--password=<password>]]
```

`--user=<user>` Optional. The name of the nonadministrative user account that you set up to run Avaki services. The account must have logon-as-service privileges.

For details about configuring user accounts, see the *Sybase Avaki EII Administration Guide*.

`--password=<password>` Optional. The password for the user account specified in `<user>`.

share-server --start

Start a share server without enabling the automatic restart feature. For information about enabling auto-restart, see [share-server --register on page 16](#).

The share server takes a minute or two to start.

Syntax

```
share-server --start
```

– This command has no options.

share-server --stop

Stop a share server that is running on Unix. For information about disabling auto-restart, see [share-server --unregister on page 18](#).

Note The procedure for stopping a share server on Windows depends on whether auto-restart is enabled for the server, as follows:

- If auto-restart is not enabled on the server, enter **Ctrl-C** in the server's window to stop it.
- If auto-restart is enabled on the server, use the Windows Services list to stop and restart the server. Right-click on the server name (for example, Avaki-ShareService-1096608202) and select **stop** or **restart**.

Syntax

```
share-server --stop
```

– This command has no options.

share-server --unregister

Stop a share server and disable the automatic restart feature.

Note In Windows, you can use Windows shortcuts to stop Avaki servers and disable auto-restart. To stop a share server and disable auto-restart, open the Start menu, then select the **Avaki Data Grid** program group. Select **Share Server**, and then select the **Unregister ShareServer as Windows service** shortcut.

Syntax

```
share-server --unregister
```

– This command has no options.

Command reference

This chapter describes Avaki CLI commands. The commands are listed alphabetically.

avaki attribute --delete

Delete an attribute for a grid file, directory, Avaki share, cache, server, or service.

Syntax

```
attribute {--delete | -d} [-r] --name=<attribute-name>  
  --type={string | integer | float | date | time | timestamp}  
  <target-grid-path>
```

-r Recursive mode (optional). If you are deleting an attribute for a directory, specify the **-r** option if you want to delete the attribute recursively from all the directory's contents. If you do not specify the **-r** option, the attribute will be deleted only from the directory itself.

--name=
<attribute-name> The name of the attribute to delete.

```
--type={string |
integer | float |
date | time |
timestamp}
```

The attribute's type, which is one of the following:

- | | |
|------------|--|
| string | Text, such as search strings or a project name. |
| integer | Any whole quantity, such as the maximum size of the cache on a machine. |
| float | A numeric value that can be fractional or very large. |
| date | Year, month, and date.

Format: yyyy-mm-dd
Example: 2003-09-21 |
| time | Hour, minute, and second that an event occurs.

Format: hh:mm:ss
Example: 23:59:59. |
| time-stamp | A precise time and date, including year, month, date, hour, minute, second, and (optionally) fraction of a second.

Format: yyyy-mm-dd hh:mm:ss[.ff]
There is a space between the day and the hour.

Example: 2003-09-21 04:05:06.57567 |

```
<target-grid-
path>
```

The data catalog path to the file, directory, or other object for which you are deleting an attribute.

avaki attribute --list

Display the name, type, and value for all of an object's attributes.

Syntax

```
avaki attribute {--list | -l} [-r] <target-grid-path>
```

<code>-r</code>	Recursive mode (optional). If you are obtaining a list of attributes for a directory, specify the <code>-r</code> option if you want information about the attributes for all of the directory's contents. If you do not specify the <code>-r</code> option, the output will include information about only the directory itself.
<code><target-grid-path></code>	The relative or absolute data catalog directory for the file, directory, or other object for which you are adding or modifying an attribute.

avaki attribute --update

Add or modify an attribute for a grid file, directory, Avaki share, cache, server, or service. See [Appendix A, "Attribute reference"](#), for a list of settable attributes.

Syntax

```
avaki attribute {--update | -u} [-r] --name=<attribute-name>
  --type={string | integer | float | date | time | timestamp}
  --value=<attribute-value> <target-grid-path>
```

<code>-r</code>	Recursive mode (optional). If you are configuring an attribute for a directory, specify the <code>-r</code> option if you want the attribute's settings to apply recursively to all the directory's contents. If you do not specify the <code>-r</code> option, the settings apply only to the directory itself.
<code>--name=<attribute-name></code>	The name of the attribute. The name may be up to 255 characters long, and it must not include spaces or special characters (such as a backslash, an ampersand, or an exclamation point).

<code>--type={string integer float date time timestamp}</code>	The attribute's type, which is one of the following:
string	Text, such as search strings or a project name. The string may be up to 255 characters long.
integer	Any whole quantity, such as the maximum size of the cache on a machine. The valid range for integers is -9,223,372,036,854,775,808 to +9,223,372,036,854,775,807.
float	A numeric value that can be fractional or very large. This might be the load on a computer or a percentage, for example. The valid range for float is $\pm 4.94065645841246544e-324$ to $\pm 1.79769313486231570e+308$.
date	Year, month, and date, such as a project date, file creation date, or a subject's birth date. Format: yyyy-mm-dd Example: 2003-09-21
time	Hour, minute, and second that an event occurs. This can be a celestial event or flight schedule information, for example. The supported range is 00:00:00 to 23:59:59. Format: hh:mm:ss Example: 04:05:06
time-stamp	A precise time and date, including year, month, date, hour, minute, second, and (optionally) fraction of a second. This attribute type can be used to mark the start and end of an experiment or a chemical reaction, for example. Format: yyyy-mm-dd hh:mm:ss[.ff] There is a space between the day and the hour. Example: 2003-09-21 04:05:06.57567

<code>--value= <attribute-value></code>	The value you're assigning to the attribute specified in the <code>--name</code> option.
	The formats for the date, time, and timestamp attribute types are as follows:
<code>date</code>	Year, month, and date; format: yyyy-mm-dd Example: 2003-09-21
<code>time</code>	Hour, minute, and second; format: hh:mm:ss Example: 04:05:06
<code>time- stamp</code>	A precise time and date, including year, month, date, hour, minute, second, and (optionally) fraction of a second. Format: yyyy-mm-dd hh:mm:ss[.ff] There is a space between the day and the hour. Example: 2003-09-21 04:05:06.57567
<code><target-grid- path></code>	The data catalog directory of the file, directory, database operation, cache service, or other object for which you are adding or modifying an attribute.

Example

This command creates a user-defined string attribute called `SlateRockDocs` for the share named `DocProjects` and sets the value to `ReleaseNotes`.

```
$ avaki attribute -u --name=SlateRockDocs --type=string
--value=ReleaseNotes /Shares/DocProjects
```

avaki backup

Back up or recover a grid server's persistent state information, including databases and internal shares. (The databases store state information about the grid domain, such as grid authentication services and access control lists.) We recommend running the backup command only when the grid is not under load.

Syntax

```
avaki backup {--snapshot | --recover}  
  [--grid-server=<grid-server>] <filename>.zip
```

<code>--snapshot</code>	Create a snapshot of the database.
<code>--recover</code>	Reset database state to a previously taken snapshot.
<code>--grid-server=<grid-server></code>	Optional. Specify the name of the Avaki server to back up. For example, Bedrock. If you omit this option, the GDC will be backed up.
<code><filename>.zip</code>	The name of the local file that you are saving a snapshot to or recovering a backup from. You can specify an absolute path or a relative path.

avaki cache --evict

Purge one or more database operations, data services, files, or directories from the cache service and unschedule them if they're pinned.

To use this command, you must have write permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default).

If you want to invalidate an item in the cache rather than permanently evicting it, see the invalidation commands, beginning with [avaki cache --invalidate](#) (page 28).

Syntax

```
avaki cache --evict <cache-service-path>
  [--dbop=<qualified-dbop-name>* |
  --data-service=<qualified-ds-name>*] [-r] <grid-path>*
```

<cache-service-path> The data catalog path of a cache service.

--dbop=
<qualified-dbop-
name>* Optional. The names of one or more database operations, in this format:
 [<domain-name>.]<database-connector-name>.
 <database-operation-name>

--data-service=
<qualified-
ds-name>* Optional. The names of one or more data services, in this format:
 [<domain-name>.]<data-service-name>

-r Optional. If an object specified in <grid-path> is a directory, recursively evict the directory and all its contents from the cache.

<grid-path>* Optional. The data catalog paths of one or more files or directories.

avaki cache --evict --all

Purge all database operations, data services, files, or directories from the cache service and unschedule them if they're pinned.

To use this command, you must have write permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default).

If you want to invalidate an item in the cache rather than permanently evicting it, see the invalidation commands, beginning with [avaki cache --invalidate](#) (page 28).

Syntax

```
avaki cache --evict --all <cache-service-path>
```

<cache-service-path> The grid path of a cache service.

avaki cache --evict --deleted

Purge from the cache any items whose source objects (database operations, data services, files, or directories) have been deleted from the domain. The cache service attempts to refresh a cache entry whenever the item's source objects have been deleted, but users may want to use [avaki cache --evict --deleted](#) to expedite the process in order to save disk space.

To use this command, you must have write permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default).

If you want to invalidate an item in the cache rather than permanently evicting it, see the invalidation commands, beginning with [avaki cache --invalidate](#) (page 28).

Syntax

```
avaki cache --evict -deleted [--dbops | --data-services |
  --filesystem-objects] <cache-service-path>
```

<code>--dbops</code>	Optional. Purge from the cache all database operations that have been deleted from the grid.
<code>--data-service</code>	Optional. Purge from the cache all data services that have been deleted from the grid.
<code>--filesystem-objects</code>	Optional. Purge from the cache all files or directories that have been deleted from the grid.
<code><cache-service-path></code>	The grid path of a cache service.

avaki cache --get

Display the name of the cache service that is associated with the specified DGAS, grid server, or local user. The association is valid across avaki logins to the same CLI.

To use this command, you must have read permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default) and on the DGAS, grid server, or local user.

Syntax

```
avaki cache --get {<DGAS> | <grid-server> | --cli}
```

<code><DGAS></code>	The data catalog path of a data grid access server.
<code><grid-server></code>	The data catalog path of a grid server.
<code>--cli</code>	Indicates the current local user's instance of the CLI.

avaki cache --invalidate

Flag one or more database operations, data services, files, or directories in a cache as invalid. When a cached copy of an object is invalidated, the cached database operation metadata is invalidated and the current cached data will no longer be used when a user requests the object. The object will be refreshed in the cache the next time it is accessed or upon its next scheduled refresh.

To use this command, you must have write permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default).

Syntax

```
avaki cache --invalidate <cache-service-path>
  --dbop=<qualified-dbop-name>*
  --data-service=<qualified-ds-name>* [-r] <grid-path>*
```

<code><cache-service-path></code>	The data catalog path of a cache service.
<code>--dbop=qualified-dbop-name*</code>	Optional. The names of one or more database operations, in this format: [<domain-name>.]<database-connector-name>.<database-operation-name>
<code>--data-service=qualified-ds-name*</code>	Optional. The names of one or more data services, in this format: [<domain-name>.]<data-service-name>
<code>-r</code>	Optional. If an object specified in <grid-path> is a directory, recursively invalidate the directory and all its contents.
<code><grid-path>*</code>	Optional. The data catalog paths to one or more files or directories.

avaki cache --invalidate --all

Flag all database operations, data services, files, or directories in a cache as invalid. When a cached copy of an object is invalidated, the cached copy data will no longer be used when a user requests the object; the object will be refreshed in the cache the next time it is accessed or upon its next scheduled refresh.

To use this command, you must have write permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default).

Syntax

```
avaki cache --invalidate --all <cache-service-path>
```

<cache-service-
path> The grid path of a cache service.

avaki cache --invalidate-dataservice-results

Invalidate the results from a cached data service that correspond to the specified input parameters, but do not invalidate cached data service metadata. When a cached copy of an object is invalidated, the current cached data will no longer be used when a user requests the object; the object will be refreshed in the cache the next time it is accessed or upon its next scheduled refresh.

To use this command, you must have write permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default).

Syntax

```
avaki cache --invalidate-dataservice-results  
  <cache-service-path> <qualified-ds-name>  
  [<param-name>:<param-value>] +
```

<cache-service-path> The data catalog path of a cache service.

qualified-ds-name Optional. The name of a data service, in this format:
 [<domain-name>.]<data-service-name>

[<param-name>:<param-value>] + Optional. Specify one or more input or output parameters and associated values required by the data service.

avaki cache --invalidate-dbop-results

Invalidate the results from a cached database operation that correspond to the specified input parameters, but do not invalidate cached database operation metadata. When a cached copy of an object is invalidated, the cached data will no longer be used when a user requests the object; the object will be refreshed in the cache the next time it is accessed or upon its next scheduled refresh.

To use this command, you must have write permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default).

Syntax

```
avaki cache --invalidate-dbop-results <cache-service-path>
  <qualified-dbop-name> [<parameter-value>]+
```

<cache-service-path> The grid path of a cache service.

qualified-dbop-name The names of a database operation, in this format:
 [<domain-name>.]<database-connector-name>.
 <database-operation-name>

<parameter-value>+ Optional. One or more parameter values required by the target database operation. The parameter values will be processed in the order specified in the database operation's query. Any parameter value that includes a space must be enclosed in quotation marks. Any string at the end of the command that is enclosed in quotation marks will be interpreted as a single value. (That is, enclose each parameter value in its own pair of quotation marks.)

avaki cache --list

Display the tagging information for the given cached objects (files, directories, database operations, and data services) contained in the specified cache. If no objects are specified for the given cache service, the tagging information for the entire collection of cached objects is returned.

To use this command, you must have read permission on the cache service.

Syntax

```
avaki cache --list <cache-service-path>
  [--dbop=<qualified-dbop-name>] *
  [--data-service=<qualified-ds-name>] * [-r] [<grid-path>*]
```

<code><cache-service-path></code>	The data catalog path of a cache service in the current Avaki domain.
<code>--dbop=<qualified-dbop-name>*</code>	Optional. Specify the name of a cached database operation, in this format: [<domain-name>.]<database-connector-name>.<database-operation-name>
<code>--data-service=<qualified-ds-name>*</code>	Optional. Specify the name of a cached data service, in this format: [<domain-name>.]<data-service-name>
<code>-r</code>	Optional. If an object specified in <code><grid-path></code> is a directory, recursively list the tagging information for the directory and any cached objects within its hierarchy.
<code><grid-path>*</code>	Optional. Specify the data catalog path of a cached grid file or directory.

avaki cache --refresh

Force the cache service to inspect (and update, if necessary) the specified cached entry.

To use this command, you must have write permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default) and on the target files or directories.

Syntax

```
avaki cache --refresh <cache-service-path>
  [--dbop=<qualified-dbop-name>] *
  [--data-service=<qualified-ds-name>] * [-r] [<grid-path>*]
```

<cache-service-path> The data catalog path of a cache service.

--dbop=
<qualified-
dbop-name>* Optional. The names of one or more database operations, in this format:
 [<domain-name>.] <database-connector-name>.
 <database-operation-name>

--data-service=
qualified-
ds-name* Optional. The names of one or more data services, in this format:
 [<domain-name>.] <data-service-name>

-r Optional. If an object specified in <grid-path> is a directory, recursively refresh the directory and all its contents in the cache.

<grid-path>* The data catalog paths of one or more files or directories.

avaki cache --set

Configure the specified client to use a cache service.

To use this command, you must have read permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default) and on the DGAS, grid server, or local user.

Syntax

```
avaki cache --set {<DGAS> | <grid-server> | --cli}  
    <cache-service-path>
```

<code><DGAS></code>	The data catalog path of a data grid access server.
<code><grid-server></code>	The data catalog path of a grid server.
<code>--cli</code>	Indicates the current local user's instance of the CLI.
<code><cache-service-path></code>	The data catalog path of a cache service.

avaki cache --unset

Uncouple any configured cache service from the specified client. If you uncouple a grid server from a cache service, the grid server uses its default setting, in which it is not associated with a cache service.

To use this command, you must have read permission on the cache service (members of the Administrators group and the DataProviders group have this permission by default) and on the DGAS, grid server, or local user.

Syntax

```
avaki cache --unset {<DGAS> | <grid-server> | --cli}
```

<DGAS>	The grid path of a data grid access server.
<grid-server>	The grid path of a grid server.
--cli	Indicates the current local user's instance of the CLI.

avaki cat

Display the specified Avaki file. Similar to the Unix **cat** command.

Syntax

```
avaki cat <grid-path>+
```

<grid-path>+	The path to the file or files to be displayed. The file(s) must be in a grid directory. You can enter an absolute path or a relative one.
--------------	---

avaki categories --add-to-category

Add objects to a category.

Syntax

```
avaki categories --add-to-category <category-path>
  [--dbop=<qualified-dbop-name>] *
  [--data-service=<qualified-ds-name>] * [<grid-path>] *
```

<category-path> The name of the category to which objects will be added, in this format:

```
/ParentCategoryName
```

To add objects to a subcategory, place a slash between the parent category name and the subcategory. For example:

```
/ParentCategoryName/SubcategoryName
```

<qualified-dbop-name>* The name of zero or more database operations, in this format:

```
[<domain-name>.] <database-connector-name> .
<database-operation-name>
```

<qualified-ds-name>* The name of zero or more data services to add to the category, in this format:

```
[<domain-name>.] <data-service-name>
```

<grid-path>* Zero or more data catalog paths of objects to add to the category.

avaki categories --create

Create a category.

Syntax

```
avaki categories --create [--description=<description>] [-D |
  --domain=<domain-name>] <category-path>
```

--description= <description-text> Optional. Enter details about this category.

[-D | --domain= <domain-name>] Optional. The name of the domain in which to create the category.

<category-path> The path of the category to create, in this format:

```
/ParentCategoryName
```

To create a subcategory, place a slash between the parent category name and the subcategory. For example:

```
/ParentCategoryName/SubcategoryName
```

avaki categories --delete

Deletes the category <category-path>.

Syntax

```
avaki categories [-d | --delete] [-D | --domain=<domain-name>]
  <category-path>
```

[-D | --domain= <domain-name>] Optional. The name of the domain from which the category will be deleted.

<category-path> The path of the category to delete.

avaki categories --describe

Show the description for a category.

Syntax

```
avaki categories --describe [-D | --domain=<domain-name>]
  <category-path>
```

[-D | --domain=<domain-name>] Optional. The name of the category's domain.

<category-path> The path of the category to describe.

avaki categories --list

List the categories that exist in a particular domain.

Syntax

```
avaki categories --list [-r | --recursive] [-D |
  --domain=<domain-name>] [<category-path>]
```

[-r | --recursive] Recursive mode (optional). Use the -r option with a blank category path to display all categories and subcategories for the domain.

[-D | --domain=<domain-name>] Optional. The name of the domain for which a list of categories will be obtained.

<category-path> Optional. List the subcategories of <category-path>.

avaki categories --remove-from-category

Remove objects from a category.

Syntax

```
avaki categories --remove-from-category [-D |
  --domain=<domain-name>] <category-path>
  [--dbop=<qualified-dbop-name>] *
  [--data-service=<qualified-ds-name>] * [<grid-path>*]
```

`[-D | --domain=<domain-name>]` Optional. The name of the domain from which the category will be removed.

`<category-path>` The path of the category from which the object will be removed.

`<qualified-dbop-name>*` The name of zero or more database operations, in this format:
 [`<domain-name>.`] `<database-connector-name>.`
`<database-operation-name>`

`<qualified-ds-name>*` The name of zero or more data services, in this format:
 [`<domain-name>.`] `<data-service-name>`

`<grid-path>*` Zero or more grid directories of objects that you are removing from a category.

avaki categories --set-description

Set the description for a category.

Syntax

```
avaki categories --set-description [-D |  
  --domain=<domain-name>] <category-path> <description-text>
```

`[-D | --domain=<domain-name>]` Optional. The name of the domain for which the category description will be set.

`<category-path>` The path of the category for which the description will be set.

`--description=<description-text>` Enter details about this category; enclose the entire text string in quotation marks. Length limit: 255 characters

avaki categories --show categories

Shows the categories that an object is a member of.

Syntax

```
avaki categories --show-categories
  [--dbop=<qualified-dbop-name> |
  --data-service=<qualified-ds-name> | <grid-path>]
```

<qualified-
dbop-name>

The name of a database operation, in this format:

[<domain-name>.]<database-connector-name>.
<database-operation-name>

<qualified-
ds-name>

The name of a data service, in this format:

[<domain-name>.]<data-service-name>

<grid-path>

The data catalog directory of the object for which you want to obtain a list of categories.

avaki cd

Change the current working directory in a grid directory. Similar to the Unix and DOS **cd** commands.

Syntax

```
avaki cd <grid-path>
```

<grid-path>

The data catalog directory to which you are changing.

avaki chmod

Modify read, write, execute, and delete permissions for users and groups in the access control list for an Avaki object.

Syntax

```
avaki chmod {<qualified-user-name> |
  --user=<qualified-user-name> | <qualified-group-name> |
  --group=<qualified-group-name> | everyone |
  --filename=<interconnect-ID-file>}
  {--delete | {--allow | --deny | --unset}=<permission-name>}+
  {--target=<target> | <target>} [-r | --recursive] [--lock]
```

```
<qualified-user-name> |
--user=<qualified-user-
name> | <group-grid-path>
| --group=<group-grid-
path> | <qualified-group-
name> | --group=
<qualified-group-name> |
everyone | --filename=
<interconnect-ID-file>
```

Specify the qualified name of the users or groups for which permissions will be changed.

The format for the qualified user name is as follows:

```
<user>@<authservice>.<authservice-
type>.<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of `<qualified-user-name>` are as follows:

<code><user></code>	The name of the user to whose permissions you are modifying.
<code><authservice></code>	The name of the user's authentication service.
<code><authservice-type></code>	The type of the user's authentication service (Grid, Nis, or Ldap).
<code><domain></code>	The name of the user's Avaki domain.

The format for the qualified group name is as follows:

```
<group>@<authservice>.<authservice-
type>.<domain>
```

For example:

```
Development@DefaultAthService.Grid.Bed-
rock
```

The elements of `<qualified-group-name>` are as follows:

<code><group></code>	The name of the group whose permissions you are modifying.
<code><authservice></code>	The name of the group's authentication service.
<code><authservice-type></code>	The type of the group's authentication service (Grid, Nis, or Ldap).
<code><domain></code>	The name of the group's Avaki domain.

Specify everyone if you want to change the permissions for all known and unknown users and groups.

```
--filename=
<interconnect-ID-file>
```

Specify the path to an interconnect ID file for the user or group whose permissions you want to add or change. This option lets you edit permissions for a user or group in another grid domain. The file is assumed to be in the local file system.

```
--delete | {--allow |
--deny | --unset}
```

For each read, write, execute, or delete permission, specify the delete, allow, deny, or unset option, or leave the current value as is.

Option	Description
delete	The user or group can delete the object.

allow	The user or group can perform the specified action on the object.
deny	The user cannot perform the specified action on the object. Note: You cannot deny permissions for a group.
unset	The user's or group's permissions for this action on this object have not been specified. Note that if you set all of an object's permissions to unset and the user is not the ACL's owner, the user will be removed from the ACL.

<permission-name>

Specify read, write, execute, or delete permission for each user or group in the ACL for an object.

Permission	Description
r	The user or group can read the object.
w	The user or group can write to the object.
x	The user or group can execute the object.
d	The user or group can delete the object.

--target=<target> |
<target>

The path to the object whose access control list permissions you are modifying.

<code>[-r --recursive]</code>	Recursive mode (optional). If you are configuring permissions for a directory, specify the <code>-r</code> option if you want the permission settings to apply recursively to all the directory's contents. If you do not specify the <code>-r</code> option, the settings apply only to the directory itself.
<code>--lock</code>	If you specify the <code>--lock</code> option, all objects in a directory are locked while the permissions on the directory are modified. Locking the objects reduces the time required to perform a <code>chmod</code> operation on a large share. Use the <code>--lock</code> option only if the <code>chmod</code> operation has exclusive access to the directory on which the permissions will be modified. If the <code>chmod</code> operation does not have exclusive access, there is a strong potential for a deadlock to occur. To provide exclusive access, stop all rehash, search, and other operations that involve the directory.

avaki chown

Change the owner of an object in the data catalog.

Syntax

```
avaki chown [-r | recursive] <user-grid-path>
            <target-grid-path>
```

<code>-r recursive</code>	Recursive mode (optional). If you are changing ownership permissions for a directory, the system changes the permissions of all the directory's contents, including all subdirectories and their contents.
<code><user-grid-path></code>	The path of the user or group whose permissions you are changing in an object's access control list.
<code><target-grid-path></code>	The path to the file or directory whose ownership you are modifying.

avaki client

Connect an Avaki command-line client to an Avaki domain, disconnect a command client from a domain, or get information about a client connected to a domain.

Syntax

```
avaki client {--connect <machine>[:<port>] | --disconnect |  
  --info}
```

<code>--connect</code> <code><machine>[:<port>]</code>	Connect a command client to the grid domain controller or grid server running on <code><machine></code> . For <code><machine></code> , specify a DNS name or IP address. For <code><port></code> , specify the Avaki server's connect port number.
<code>--disconnect</code>	Disconnect a command client from an Avaki domain.
<code>--info</code>	Display the name of the Avaki domain to which a client is connected.

avaki cp

Copy files and directories (Avaki or local) from the specified source file or path to the specified destination. You can use **avaki cp** to copy from a grid directory to the local file system, from one grid directory to another, from the local file system to a shared directory, or from the local file system to the local file system. You may not copy from a shared directory to a nonshared directory. This command is similar to the Unix **cp** command.

Caution When you copy a file or directory into an Avaki share, out of an Avaki share, or within an Avaki share, the change is reflected in the source file system—not in the grid only. You cannot copy an entire shared directory.

Syntax

```
avaki cp [-r] [--localsrc] [--localdst] <copy-source>+
        <copy-destination>
```

<code>-r</code>	Recursive mode (optional). If the source or destination is a directory, the system copies the contents, including all subdirectories and their contents.
<code> [--localsrc]</code>	Optional. Indicates that the source file or path is in the local file system. If you omit this option, the system assumes that the source file or path is in a grid directory.
<code> [--localdst]</code>	Indicates that the destination file or path is in the local file system. If you omit this option, the system assumes that the destination file or path is in a grid directory.
<code><copy-source>+</code>	One or more files, directories, or other objects to be copied. If you enter only a name, the system assumes that the object is in the current directory. Enter an absolute or relative path to copy an object in another directory.
<code><copy-destination></code>	The location for the copy. If you specify only a name, the system places the copy in the current directory. Specify an absolute or relative path to place the copy in another directory.

Example

This example copies the file `bettysform.txt` from the grid directory `/Shares/Slaterock/` to the directory `c:/Newforms` on the local file system.

```
$ avaki cp --localdst /Shares/Slaterock/bettysform.txt
c:/Newforms/bettysform.txt
```

This example copies the file `bettysform.txt` from the grid directory `/Shares/Slaterock` to the grid directory `/Shares/Newforms`.

```
$ avaki cp /Shares/Slaterock/bettysform.txt
/Shares/Newforms/bettysform.txt
```

This example copies the file `bettysform.txt` from the local file system directory `c:/Slaterock` to the local file system directory `c:/Newforms`.

```
$ avaki cp --localsrc --localdst c:/Slaterock/bettysform.txt
c:/Newforms/bettysform.txt
```

avaki dataservice --add-schedule

Add a caching schedule for a data service.

To use this command, you must have write permission on the cache service, and you must have read and execute permission on the data service.

Syntax

```
avaki dataservice --add-schedule
  {--periodic-schedule=<period-spec> |
  --cron-schedule=<cron-spec> |
  --one-time-schedule=<time-stamp>}
  [--start-time=<time-stamp>] [[--end-time=<time-stamp>] |
  [--max-iterations=<num-iterations>]]
  [--time-zone=<time-zone>] [--exclusion=<exclusion-name>]+
  [--exclusion-action={ drop | sooner | later }]
  [--reschedule-offset=<offset-spec>] <cache-service-path>
  <qualified-ds-name> [<param-name>:<param-value>]+
```

```
--periodic-  
schedule=  
<period-spec>
```

Specify a periodic schedule for running the data service (which is identified in the <qualified-ds-name> argument). Use the following syntax:

```
--periodic-schedule="  
  <number-of-  
  units> <unit-type>"
```

<number-of-units> An integer from 1 to 2147483648

<unit-type> seconds, minutes, hours, days,
 weeks, or months

In the following example, the periodic schedule is set to 30 seconds:

```
--periodic-schedule="30 seconds"
```

```
--cron-  
schedule=  
<cron-spec>
```

Specify how often an advanced schedule based on a cron expression will run the data service (which is identified in the <qualified-ds-name> argument). Use the following syntax:

```
--cron-schedule="  
  <seconds> <minutes> <hours>  
  <days-of-month> <months> <days-of-week>  
  [<years>]"
```

For details on cron expression syntax, see [Appendix C, “Configuring cron schedules”](#).

```
--one-time-  
schedule=  
<time-stamp>
```

Specify a single time and date for running the data service (which is identified in the <qualified-ds-name> argument). Use the following syntax:

```
--one-time-schedule="{<hh>:<mm>:<ss> {am | pm}  
<mm>-<dd>-<yyyy> | now}"
```

```
--start-time=  
<time-stamp>
```

Specify the time and date that the data service schedule will start running, in this format:

```
{<hh>:<mm>:<ss> {am | pm} <mm>-<dd>-<yyyy> |  
now}
```

<code>--end-time= <time-stamp></code>	Specify the time and date that the data service schedule will stop running, in this format: {<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}																																
<code>--max-iterations= <num-iterations></code>	Specify the number of times a schedule will repeat. <num-iterations> An integer from 1 to 2147483648. If no value is specified, the schedule repeats indefinitely.																																
<code>--time-zone= <time-zone></code>	Specify the time zone for the schedule. The schedule can be specified according to the cache server's time zone or relative to Greenwich Mean Time (GMT). If no value is specified for the time zone, the server's local time zone is used. The format for the time zone is as follows: <table> <tr><td>GMT+12:00 Fiji</td><td>GMT+02:00 Cairo</td></tr> <tr><td>GMT+11:00 Noumea</td><td>GMT+01:00 Paris</td></tr> <tr><td>GMT+10:00 Sydney</td><td>GMT London</td></tr> <tr><td>GMT+9:30 Adelaide</td><td>GMT-01:00 Azores</td></tr> <tr><td>GMT+09:00 Tokyo</td><td>GMT-02:00 Mid-Atlantic</td></tr> <tr><td>GMT+08:00 Hong Kong</td><td>GMT-03:00 Rio de Janeiro</td></tr> <tr><td>GMT+07:00 Bangkok</td><td>GMT-03:30 Newfoundland</td></tr> <tr><td>GMT+06:30 Rangoon</td><td>GMT-04:00 Caracas</td></tr> <tr><td>GMT+06:00 Dacca</td><td>GMT-05:00 New York</td></tr> <tr><td>GMT+05:45 Katmandu</td><td>GMT-06:00 Chicago</td></tr> <tr><td>GMT+05:30 Calcutta</td><td>GMT-07:00 Denver</td></tr> <tr><td>GMT+05:00 Karachi</td><td>GMT-08:00 San Francisco</td></tr> <tr><td>GMT+04:30 Kabul</td><td>GMT-09:00 Juneau</td></tr> <tr><td>GMT+04:00 Baku</td><td>GMT-10:00 Hawaii</td></tr> <tr><td>GMT+03:30 Tehran</td><td>GMT-11:00 Samoa</td></tr> <tr><td>GMT+03:00 Moscow</td><td></td></tr> </table>	GMT+12:00 Fiji	GMT+02:00 Cairo	GMT+11:00 Noumea	GMT+01:00 Paris	GMT+10:00 Sydney	GMT London	GMT+9:30 Adelaide	GMT-01:00 Azores	GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic	GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro	GMT+07:00 Bangkok	GMT-03:30 Newfoundland	GMT+06:30 Rangoon	GMT-04:00 Caracas	GMT+06:00 Dacca	GMT-05:00 New York	GMT+05:45 Katmandu	GMT-06:00 Chicago	GMT+05:30 Calcutta	GMT-07:00 Denver	GMT+05:00 Karachi	GMT-08:00 San Francisco	GMT+04:30 Kabul	GMT-09:00 Juneau	GMT+04:00 Baku	GMT-10:00 Hawaii	GMT+03:30 Tehran	GMT-11:00 Samoa	GMT+03:00 Moscow	
GMT+12:00 Fiji	GMT+02:00 Cairo																																
GMT+11:00 Noumea	GMT+01:00 Paris																																
GMT+10:00 Sydney	GMT London																																
GMT+9:30 Adelaide	GMT-01:00 Azores																																
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic																																
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro																																
GMT+07:00 Bangkok	GMT-03:30 Newfoundland																																
GMT+06:30 Rangoon	GMT-04:00 Caracas																																
GMT+06:00 Dacca	GMT-05:00 New York																																
GMT+05:45 Katmandu	GMT-06:00 Chicago																																
GMT+05:30 Calcutta	GMT-07:00 Denver																																
GMT+05:00 Karachi	GMT-08:00 San Francisco																																
GMT+04:30 Kabul	GMT-09:00 Juneau																																
GMT+04:00 Baku	GMT-10:00 Hawaii																																
GMT+03:30 Tehran	GMT-11:00 Samoa																																
GMT+03:00 Moscow																																	

<code>--exclusion= <exclusion-name></code>	The name of an exclusion to this schedule. For information about creating schedule exclusions, see avaki scheduleexclusion --create --daily (page 201), avaki scheduleexclusion --create --weekly (page 205), avaki scheduleexclusion --create --monthly (page 203), avaki scheduleexclusion --create --yearly (page 207), and avaki scheduleexclusion --create --custom (page 199).
<code>--exclusion- action={drop sooner later}</code>	Specify the action to take if the data service is scheduled to run during an exclusionary period.
drop	Don't run the data service during the exclusionary period.
sooner	Reschedule the data service to run before the exclusionary period
later	Reschedule the data service to run after the exclusionary period.
<code>--reschedule- offset= <offset-spec></code>	Specify the amount of time before or after the exclusionary period to perform a rescheduled data service run. Use the following syntax:
<code>--reschedule-offset=" <number-of-units> <unit-type>"</code>	
<code><number-of-units></code>	An integer from 1 to 2147483648
<code><unit-type></code>	Seconds, minutes, hours, days, weeks, or months
<code><cache-service- path></code>	The data catalog path of a cache service.
<code><qualified- ds-name></code>	The name of a data service, in this format: [<domain-name>.]<data-service-name>
<code><param-name>: <param-value>+</code>	Optional. One or more parameters required by the data service and one or more parameter values.

Example

This example adds a periodic schedule in which the data service `FlintstoneDataService` is run every 30 seconds and the time zone is GMT-05:00 New York.

```
$ avaki dataservice --add-schedule --periodic-schedule=
"30 seconds" --time-zone="GMT-05:00 New York"
/System/LocalDomain/Services/CacheServices/
BEDROCK.sybase.com FlintstoneDataService
```

avaki dataservice --create

Create a data service.

To use this command, you must be a member of the `DataProviders` group.

Syntax

```
avaki dataservice --create [{--server | -s}=<grid-server>]
[--localsrc] [--expiration=<expiration-time-in-secs>]
[--run-user={<qualified-user-name> | null}]
<descriptor-path> <data-service-name>
```

<code>--server -s}=<grid-server></code>	Optional. Specify the name of the Avaki server on which the data service will be created. For example, <code>Bedrock</code> . If you omit this option, the data service will be created on the same server as the parent directory.
<code>--localsrc</code>	Optional. Specify the <code>--localsrc</code> option if the <code><descriptor-path></code> is a path in the local file system. If <code>--localsrc</code> is not specified, <code><descriptor-path></code> is assumed to be a grid path.
<code>--expiration= <expiration- time-in-secs></code>	Optional. Specify the interval, in seconds, after which the results of this data service expire from the on-demand cache. If you set this option to <code>-1</code> , the results never expire. If you set it to <code>0</code> , the data is not cached. Default: <code>0</code> (results are not cached) Range: <code>0</code> to <code>2147483647</code>


```
--run-user=
{<qualified-user-
name> | null}
```

Specify the qualified name of the user to create the data service as. The format is:

```
<user>@<authservice>.<authservicetype>.  
<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user to create the data service as.
<authservice>	The name of the user's authentication service.
<authservice-type>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

If you specify null, the data service is run as the current user.

```
<descriptor-path>
```

Optional. Specify the path to an XML file that contains a data service descriptor. See the *Sybase Avaki EII Provisioning and Advanced Data Integration Guide* for a description of the XML schema that specifies the contents of a data service.

```
<data-service-name>
```

A name for the data service. The name must not contain any period characters (.) or spaces.

avaki dataservice --delete

Delete a data service.

To use this command, you must have execute permission on the data service.

Syntax

```
avaki dataservice --delete <qualified-ds-name>
```

<code><qualified-ds-name></code>	The name of a data service, in this format: [<code><domain-name>.</code>] <code><data-service-name></code>
--	---

avaki dataservice --delete-schedule

Delete a caching schedule for a data service.

To use this command, you must have write permission on the cache service.

Syntax

```
avaki dataservice --delete-schedule <cache-service-path>
  <qualified-ds-name> [<param-name>:<param-value>]+
  <schedule-id>
```

<code><cache-service-path></code>	The data catalog path of a cache service.
---	---

<code><qualified-ds-name></code>	The name of a data service, in this format: [<code><domain-name>.</code>] <code><data-service-name></code>
--	---

<code><param-name>: <param-value>+</code>	Optional. One or more parameters required by the data service and one or more parameter values.
---	---

<code><schedule-id></code>	The numeric ID of the schedule to delete. To determine a schedule's ID, use avaki dataservice --list-schedules (page 58).
----------------------------------	---

avaki dataservice --depends

Show the dependencies to and from a data service.

To use this command, you must have read permission on the message service on the grid server where the data service is located, and you must have read permissions on the data service.

Syntax

```
avaki dataservice --depends <qualified-ds-name>
```

<qualified-
ds-name>

The name of a data service, in this format:

[<domain-name>.]<data-service-name>

avaki dataservice --execute

Run a data service.

To use this command, you must have execute permission on the data service.

Syntax

```
avaki dataservice --execute <qualified-ds-name>  
  [<param-name>:<param-value>]+ [--generate-schema]
```

<code><qualified-ds-name></code>	The name of a data service, in this format: <code>[<domain-name>.]<data-service-name></code>
<code><param-name>: <param-value>+</code>	Optional. Specify one or more input or output parameters and associated values required by the data service.
<code>--generate-schema</code>	Optional. Generate an XSD schema file for the data service. The schema file describes the structure of the output of the data service. The generated file will be stored in the following path and file: <code>/Metadata/DataService/<data-service-name>/ <data-service-name>.xsd</code>

avaki dataservice --generate-sql-view

Generate a SQL view from a data service. The SQL view is placed in the following data catalog directory:

System/Domains/<domain-name>/Services/DataServices/VirtualSQLViews

Note You must generate a schema before generating a SQL view. You can use [avaki dataservice --execute \(page 56\)](#) to generate a schema.

Syntax

```
avaki dataservice --generate-sql-view
  [--provision-as=<table-name>]
  [--param-column=<input-param-name>:<column-name>] +
  <qualified-ds-name>
```

<pre>--provision-as= <table-name></pre>	<p>Optional. Specify a name for the SQL view. If no name is specified, the SQL view will have the same name as the data service.</p>
<pre>--param-column= <input-param- name>:<column- name></pre>	<p>Optional. Specify column names for the input parameters, if any. When Avaki generates a SQL view for a data service, it must create table column names for any input parameters to the data service (in addition to using the column names for the output data specified in the schema file) so that the input parameters can be specified in a SQL query on that SQL view.</p> <p>If you omit this option, Avaki uses the name(s) of the data service input parameter(s) to create column names, first removing illegal characters such as spaces.</p>
<pre><qualified-ds- name></pre>	<p>The name of the data service for which you are generating a SQL view, in this format:</p> <pre>[<domain-name>.]<data-service-name></pre>

avaki dataservice --info

Display the XML data service descriptor.

Syntax

```
avaki dataservice --info <qualified-ds-name>
```

<code><qualified-ds-name></code>	The name of a data service, in this format: [<domain-name>.]<data-service-name>
--	--

avaki dataservice --list-schedules

List the caching schedules for a data service.

Syntax

```
avaki dataservice --list-schedules <cache-service-path>
  <qualified-ds-name> [<param-name>:<param-value>]+
```

<code><cache-service-path></code>	The grid path of a cache service.
---	-----------------------------------

<code><qualified-ds-name></code>	The name of a data service, in this format: [<domain-name>.]<data-service-name>
--	--

<code><param-name>: <param-value>+</code>	Optional. Specify one or more input or output parameters and associated values required by the data service.
---	--

avaki dataservice --update

Modify the configuration of a data service.

Syntax

```
avaki dataservice --update <qualified-ds-name> [[--localsrc]
  <descriptor-path> [--expiration=<expiration-time-in-secs>]
  [--run-user={<qualified-user-name> | null}]
```

<qualified-
ds-name>

The name of a data service, in this format:

[<domain-name>.]<data-service-name>

--localsrc

Optional. Use the --localsrc option if the <descriptor-path> is a path in the local file system. If --localsrc is not specified, <descriptor-path> is assumed to be a grid path.

<descriptor-path>

Optional. Specify the path to an XML file that contains a data service descriptor. See the *Sybase Avaki EII Provisioning and Advanced Data Integration Guide* for a description of the XML schema that specifies the contents of a data service.

--expiration=
<expiration-
time-in-secs>

Optional. Specify the interval, in seconds, after which the results of this data service expire from the on-demand cache. If you set this option to -1, the results never expire. If you set it to 0, the data is not cached.

Default: 0 (results are not cached)

Range: 0 to 2147483647

```
--run-user={<qualified-user-name> | null}
```

Specify the qualified name of the user to modify the data service as. The format is:

```
<user>@<authservice>.<authservice-type>.<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-user-name> are as follows:

- | | |
|--------------------|--|
| <user> | The name of the user to modify the data service as. |
| <authservice> | The name of the user's authentication service. |
| <authservice-type> | The type of authentication service (Grid, Nis, or Ldap). |
| <domain> | The name of the user's Avaki domain. |

If you specify null, the database operation is run as the current user.

avaki dbconn --allow-dbop-creation

Allow a user to create database operations that use a designated database connection.

Syntax

```
avaki dbconn --allow-dbop-creation <dbconn-name>
    {<qualified-user-name>}
```

<dbconn-name> The name of the database connector.

<qualified-user-name> Specify the qualified name of the user who can create a database operation that uses the database connection. The format is:

```
<user>@<authservice>.<authservicetype>.<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-user-name> are as follows:

<user> The name of the user who can create a database operation that uses the database connection.

<authservice> The name of the user's authentication service.

<authservice-type> The type of authentication service (Grid, Nis, or Ldap).

<domain> The name of the user's Avaki domain.

avaki dbconn --delete

Delete an Avaki database connector.

To use this command, you must be a member of the DatabaseAdministrators group.

Syntax

```
avaki dbconn --delete <dbconn-name>
```

<dbconn-name> The name of the database connector to be deleted.

avaki dbconn --disallow-dbop-creation

Prevent a user from creating database operations that use a designated database connector.

Note This command is intended to reverse the effect of **avaki dbconn --allow-dbop-creation** (page 61); it might not reverse the effects of permissions that have been set manually.

Syntax

```
avaki dbconn --disallow-dbop-creation <dbconn-name>
    {<qualified-user-name>}
```

<dbconn-name> The name of the database connector.

<qualified-user-name> Specify the qualified name of the user to prevent from creating database operations that uses the database connection. The format is:
 <user>@<authservice>.<authservicetype>.<domain>

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user to prevent from creating a database operation that use the database connection.
<authservice>	The name of the user's authentication service.
<authservice-type>	The type of authentication service (Grid, Nis, or Ldap).
<domain-name>	The user's grid domain.

avaki dbconn --execute

Execute an ad-hoc SQL statement on a database connector.

To use this command, you must be a member of the DatabaseAdministrators group.

Syntax

```
avaki dbconn --execute [{--xmlinputfile | --localxmlinputfile}]
  [--prepared | --callable] [--batch] [--params={in | out |
  inout}:<parameter-types>] <dbconn-name> <sql>
  [<parameter-value>+]
```

<code>--xmlinputfile</code>	Optional. If the query requires an input file, use this option to specify the path to the file. For details about the format for an XML input file, see the <i>Sybase Avaki EII Provisioning and Advanced Data Integration Guide</i>
<code>--localxmlinputfile</code>	Optional. If the query requires an input file, use this option to specify the local path to the file. For details about the format for an XML input file, see the <i>Sybase Avaki EII Provisioning and Advanced Data Integration Guide</i>
<code>--prepared</code>	Optional. Specify this option to use a JDBC prepared statement to execute the SQL.
<code>--callable</code>	Optional. Specify this option if the SQL statement <sql> invokes a stored procedure in the underlying database.
<code>--batch</code>	Optional. Specify this option if the SQL statement <sql> is an update that can be used with JDBC batch mode.

```
--params={in |
out | inout}:
<parameter-
types>
```

Optional. If the query accepts input parameters or produces output parameters, use this option to specify the data type for each input and output. There is no default.

Use in for input parameters, out for output parameters, and inout for parameters that can be both. If you specify multiple parameters, use semicolons (;) to separate them.

You can use the following <parameter-types>:

BIGINT	LONGVARBINARY
BINARY	LONGVARCHAR
BIT	NUMERIC
BOOLEAN	ORACLE_CURSOR
BLOB	OTHER
CHAR	REAL
CLOB	SMALLINT
DATE	TIME
DECIMAL	TIMESTAMP
DOUBLE	TINYINT
FLOAT	VARBINARY
INTEGER	VARCHAR

```
<dbconn-name>
```

The name of a database connector.

```
<sql>
```

Enter the SQL statement that the database connector will execute. Enclose the statement in quotation marks.

```
<parameter-
value>+
```

Optional. One or more parameter values required by the target database connector. If a database query is associated with the database connector, the parameter values will be processed in the order specified in the database operation's query. Any parameter value that includes a space must be enclosed in quotation marks. Any string at the end of the command that is enclosed in quotation marks will be interpreted as a single value. (That is, enclose each parameter value in its own pair of quotation marks.)

avaki dbconn --info

Display information about an Avaki database connector, including its JDBC driver, connection string, user name, connection timeout, and whether it is using Avaki connection pooling.

To use this command, you must be a member of the DatabaseAdministrators group.

Syntax

```
avaki dbconn --info <dbconn-name>
```

<dbconn-name> The name of a database connector.

avaki dbconn --jdbc

Create or modify an Avaki database connector. A database connector allows database operations and their views to run queries on a relational database. In addition, a database connector can be used to provision SQL views.

To use this command, you must be a member of the DatabaseAdministrators group.

Syntax

```
avaki dbconn --jdbc {--create | --update}
  --driver=<jdbc-driver-class> --connUrl=<jdbc-connection-url>
  [--user=<jdbc-user-name>] [--password=<jdbc-password>]
  [--pool={true | false}] [--adhoc={true | false}]
  [--properties=<user-jdbc-properties>]
  [{--server | -s}=<grid-server>]
  [--description=<dbconn-description-text>]
  [--mapping-file=<grid-path>]
  [--database-name=<database-name>] [--admin-name=<DBA-name>]
  [--admin-org=<DBA-org>] [--admin-email=<DBA-email>]
  [--admin-phone=<DBA-phone>]
  [--xa-datasource-class=<class-name>]
  [--xa-property="<name=value>" +] <dbconn-name>
```

<code>--create</code>	Set up a new database connector.
<code>--update</code>	Modify an existing database connector.
<code>--driver= <jdbc-driver-class></code>	Specify the JDBC driver that this database connector will use.
<code>--connUrl=<jdbc- connection-url></code>	Specify the connect string for the JDBC driver that connects to your database.
<code>--user= <jdbc-user-name></code>	Optional. Specify the name of a user account in the target database.
<code>--password= <jdbc-password></code>	Optional. Specify the password for <jdbc-user-name>. If you're executing this command on a Unix system and the password contains non-alphanumeric characters (! or &, for instance), enclose the password in single quotation marks to prevent the shell from interpreting the special characters.
<code>--pool= {true false}</code>	Optional. Indicate whether this database connector will use (true) or not use (false) Avaki connection pooling. When connection pooling is enabled, database connections can be reused, which typically improves performance for JDBC applications. Set this option to true if your JDBC driver does <i>not</i> have built-in connection pooling. Default: false (no pooling)
<code>--adhoc= {true false}</code>	Optional. Set this option to true to enable ad hoc queries, which let you directly query a database in SQL. Otherwise, set to false. Default: false
<code>--user-property= "<name=value>"</code>	Optional. Enter any driver- or database-specific properties that are needed to connect to the target database. For example, the command line syntax might look like this: <pre>--user-property="User=fred" --user-property="Password=mypassword" --user-property="ArbitraryStringProperty= hello" --user-property="ArbitraryInt Property=1"</pre>

<pre>--server -s}= <grid-server></pre>	<p>Optional. Specify the name of the database connector’s grid server. For example, Bedrock. If you omit this option, the database connector is created or modified on the current grid server.</p>
<pre>--description= <dbconn-description -text></pre>	<p>Optional. Enter details about this database connector; enclose the entire text string in quotation marks. This description is shown to users who display information about the database connector (see avaki dbconn --info on page 66). Consider including an explanation of what this database connector is for or where to get more information. Length limit: 255 characters</p>
<pre>--mapping-file= <grid-path></pre>	<p>Optional. Provide the name and path to an XML file in the data catalog in which you specify mappings between SQL data types and the data types used by the Avaki query engine. The mappings will be used for this database connector only.</p> <p>For information on the format and content of the mapping file, see Appendix D, “Data type mapping files”.</p>
<pre>--database-name= <database-name></pre>	<p>Optional. Enter the name of the database that this database connector links to; enclose the entire text string in quotation marks. This text is shown to users who display information about the database connector (see avaki dbconn --info on page 66). Length limit: 255 characters</p>
<pre>--admin-name= <DBA-name></pre>	<p>Optional. Enter the name of the database administrator for this dbconn’s database; enclose the entire text string in quotation marks. This text is shown to users who display information about the database connector (see avaki dbconn --info on page 66). Length limit: 255 characters</p>
<pre>--admin-org= <DBA-org></pre>	<p>Optional. Enter the organization to which the database administrator belongs; enclose the entire text string in quotation marks. This text is shown to users who display information about the database connector (see avaki dbconn --info on page 66). Length limit: 255 characters</p>

- `--admin-email=
<DBA-email>` Optional. Enter the e-mail address of the database administrator for this dbconn's database; enclose the text in quotation marks. This text is shown to users who display information about the database connector (see [avaki dbconn --info on page 66](#)). Length limit: 255 characters
- `--admin-phone=
<DBA-phone>` Optional. Enter the database administrator's telephone number. This text is shown to users who display information about the database connector (see [avaki dbconn --info on page 66](#)). Length limit: 255 characters
- `--xa-datasource-
class=<class-name>` Optional. If you plan to use this database connector for distributed transactions, enter the fully qualified class name of the XADataSource implementation provided by the database vendor. This class must reside in the JAR file for the database vendor's JDBC driver.
- For details on configuring your XA-compliant JDBC driver, see the documentation for the driver.

`--xa-property=`
`"<name=value>"`

Enter a string (enclosed in quotation marks) that contains the name and value of a connection property to be set on the XADataSource. This string should be in the form “<property-name>=<property-value>.” The property names must correspond to the name of a setter method to be called on the XADataSource, with the “set” prefixes excluded. If the value of the property is an integer, it may be specified as a string and will be converted to an integer before the setter method is called. To specify multiple properties, repeat the `--xa-property` option.

For example, if the following calls should be made on the XADataSource:

```
source.setUser("fred");
source.setPassword("mypassword");
source.setArbitraryStringProperty("hello");
source.setArbitraryIntProperty(1);
```

Then the command line syntax might look like this:

```
--xa-property="User=fred"
--xa-property="Password=mypassword"
--xa-property="ArbitraryStringProperty=hello"
--xa-property="ArbitraryIntProperty=1"
```

See also the [Examples](#) section, below. For details on the particular XA connection properties to use with your XA-compliant JDBC driver, refer to the documentation for the driver.

`<dbconn-name>`

A name for the database connector. The name must not exceed 64 characters in length and must not contain any period characters (.) or spaces.

Examples

This command creates an XA-enabled database connector called myXAconn for a Sybase ASE database called dino. (XA-enabled database connectors can support distributed transactions.)

```
$ avaki dbconn --jdbc --create
--driver=com.sybase.jdbc3.jdbc.SybDriver
--connUrl=jdbc:sybase:tds:myDBserver:15000/dino
--user=barney --password=hobby
--xa-datasource-class=com.sybase.jdbc3.jdbc.SybXADataSource
--xa-property="ServerName=myDBserver"
--xa-property="PortNumber=5000"
```

```

--xa-property="ResourceManagerType=2"
--xa-property="ResourceManagerName=connection"
--xa-property="DatabaseName=dino"
--xa-property="NetworkProtocol=Tds" --xa-property="User=barney"
--xa-property="Password=hobby" myXAconn

```

This command creates an XA-enabled database connector called conn2pebbles for a MySQL database called Pebbles whose listening port is 17089.

```

$ avaki dbconn --jdbc --create
--driver=com.mysql.jdbc.Driver
--connUrl=jdbc:mysql://dbserver:17089/Pebbles
--user=barney --password=hobby
--xa-datasource-class=com.mysql.jdbc.jdbc2.optional.MysqlXA
DataSource
--xa-property="URL=jdbc:mysql://dbserver/Pebbles"
--xa-property="User=barney" --xa-property="Password=hobby"
conn2pebbles

```

This command creates a database connector called gazoo-conn for a DB2 database called gazoo whose listening port is 6789.

```

$ avaki dbconn --jdbc --create
--driver=COM.ibm.db2.jdbc.net.DB2Driver
--connUrl=jdbc:db2://localhost:6789/gazoo --user=barney
--password=hobby gazoo-conn

```

avaki dbconn --provision-tables

Provision a SQL view from a database connector.

Syntax

```
avaki dbconn --provision-tables <dbconn-name>
  --spec=<provision-spec>+
```

<code><dbconn-name></code>	The name of the database connector from which you are provisioning a SQL view.
<code>--spec=</code> <code><provision-spec>+</code>	Specify one or more parameters for provisioning the SQL view; enclose parameters in double quotation marks (for example, <code>--spec="table=foo;type=TABLE"</code>). The <code><provision-spec></code> options are as follows:
<code>provision-as=</code> <code><provisioned-</code> <code>SQL-view-name></code>	The name of the SQL view to provision.
<code>[/catalog=</code> <code><database-</code> <code>catalog>]</code>	Optional. The name of the database catalog to provision.
<code>[/schema=</code> <code><database-</code> <code>schema>]</code>	Optional. The name of the database schema to provision.
<code>[/table=</code> <code><table-name>]</code>	Optional. The name of the table to provision.
<code>[/description=</code> <code><description>]</code>	Optional. Enter details about the SQL view; enclose the entire text string in single quotation marks. Consider including an explanation of what this SQL view is for, how to use it, or where to get more information. Length limit: 255 characters
<code>[/type=</code> <code><table-type>]*</code>	Optional. The type of table supported by the native database. The valid types are TABLE and VIEW.

avaki dbconn --test

Test an Avaki database connector. The test connects to the target database and displays the JDBC driver version, database version, and database name. To use this command, you must be a member of the DatabaseAdministrators group.

Syntax

```
avaki dbconn --test <dbconn-name>
```

<dbconn-name> The name of the database connector to be tested.

avaki dbconn --show-tables

Display SQL views that have been provisioned from a database connector.

Syntax

```
avaki dbconn --show-tables <dbconn-name>
  [--catalog=<database-catalog>] [--schema=<database-schema>]
  [--table=<table-name>] [--type=<table-type>] *
```

<dbconn-name> The name of the database connector whose provisioned SQL views you want to view.

--catalog=
<database-
catalog> Optional. Restrict the display to SQL views from the specified catalog.

--schema=
<database-schema> Optional. Restrict the display to SQL views from the specified database schema.

<code>--table= <table-name></code>	Optional. Restrict the display to SQL views from the specified table.
<code>--type= <table-type></code>	Optional. Restrict the display to SQL views from the specified table type. The valid types are TABLE and VIEW.

avaki dbop --add-schedule

Add a caching schedule for a database operation.

To use this command, you must have write permission on the cache service, and you must have read and execute permission on the database operation.

Syntax

```
avaki dbop --add-schedule {--periodic-schedule=<period-spec> |
  --cron-schedule=<cron-spec> |
  --one-time-schedule=<time-stamp>}
  [--start-time=<time-stamp>] [--end-time=<time-stamp>] |
  [--max-iterations=<num-iterations>]
  [--time-zone=<time-zone>] [--exclusion=<exclusion-name>]+
  [--exclusion-action={ drop | sooner | later }]
  [--reschedule-offset=<offset-spec>] <cache-service-path>
  <qualified-dbop-name> [<parameter-value>]+
```

<code>--periodic- schedule= <period-spec></code>	Specify a periodic schedule for running the database operation (which is identified in the <code><qualified-dbop-name></code> argument). Use the following syntax:
--	--

```
--periodic-schedule="<number-of-  
units> <unit-type>"
```

<code><number-of- units></code>	An integer from 1 to 2147483648
---	---------------------------------

<code><unit-type></code>	seconds, minutes, hours, days, weeks, or months
--------------------------------	---

In the following example, the periodic schedule is set to 30 seconds:

```
--periodic-schedule="30 seconds"
```

```
--cron-  
schedule=  
<cron-spec>
```

Specify how often a cron-based schedule for running the database operation (which is identified in the <qualified-dbop-name> argument) will run the database operation. Use the following syntax:

```
--cron-schedule="<seconds> <minutes> <hours>  
<days-of-month> <months> <days-of-week>  
[<years>]"
```

For details about the <cron-spec> syntax, see [Appendix C, “Configuring cron schedules”](#).

```
--one-time-  
schedule=  
<time-stamp>
```

Specify a single time and date for running the database operation (which is identified in the <qualified-dbop-name> argument). Use the following syntax:

```
--one-time-schedule="{<hh>:<mm>:<ss> {am | pm}  
<mm>-<dd>-<yyyy> | now}"
```

```
--start-time=  
<time-stamp>
```

Specify the time and date that the database operation will start running, in this format:

```
{<hh>:<mm>:<ss> {am | pm} <mm>-<dd>-<yyyy> |  
now}
```

```
--end-time=  
<time-stamp>
```

Specify the time and date that the database operation will stop running, in this format:

```
{<hh>:<mm>:<ss> {am | pm} <mm>-<dd>-<yyyy> |  
now}
```

```
--max-iterations=  
<num-iterations>
```

Specify the number of times a schedule will repeat.

```
<num-iterations>      An integer from 1 to 2147483648.  
If no value is specified, the sched-  
ule repeats indefinitely.
```

--time-zone=
<time-zone>

Specify the time zone for the schedule. The schedule can be specified according to the cache server's time zone or relative to Greenwich Mean Time (GMT). If no value is specified for the time zone, the server's local time zone is used. The format for the time zone is as follows:

GMT+12:00 Fiji	GMT+02:00 Cairo
GMT+11:00 Noumea	GMT+01:00 Paris
GMT+10:00 Sydney	GMT London
GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii
GMT+03:30 Tehran	GMT-11:00 Samoa
GMT+03:00 Moscow	

--exclusion=
<exclusion-name>

The name of an exclusion to this schedule. For information about creating schedule exclusions, see [avaki scheduleexclusion](#) **--create --daily** (page 201), [avaki scheduleexclusion --create --weekly](#) (page 205), [avaki scheduleexclusion --create --monthly](#) (page 203), [avaki scheduleexclusion --create --yearly](#) (page 207), and [avaki scheduleexclusion --create --custom](#) (page 199).

<code>--exclusion-action={drop sooner later}</code>	Specify the action to take if the database operation is scheduled to run during an exclusionary period.
<code>drop</code>	Don't run the database operation during the exclusionary period.
<code>sooner</code>	Reschedule the database operation to run before the exclusionary period.
<code>later</code>	Reschedule the database operation to run after the exclusionary period.
<code>--reschedule-offset=<offset-spec></code>	Specify the amount of time before or after the exclusionary period to perform a rescheduled database operation run. Use the following syntax:
	<code>--reschedule-offset="<number-of-units> <unit-type>"</code>
<code><number-of-units></code>	An integer from 1 to 2147483648
<code><unit-type></code>	Seconds, minutes, hours, days, weeks, or months
<code><cache-service-path></code>	The grid path of a cache service.
<code><qualified-dbop-name></code>	The name of a database operation, in this format: [<domain-name>.]<database-connector-name>.<database-operation-name>
<code><parameter-value>+</code>	Optional. One or more parameter values required by the target database operation. The parameter values will be processed in the order specified in the database operation's query. Any parameter value that includes a space must be enclosed in quotation marks. Any string at the end of the command that is enclosed in quotation marks will be interpreted as a single value. (That is, enclose each parameter value in its own pair of quotation marks.)

Example

This example adds a periodic schedule in which the database operation MyDBOP is run every 30 seconds and the time zone is GMT-05:00 New York.

```
$ avaki dbop --add-schedule --periodic-schedule="30 seconds"
--time-zone="GMT-05:00 New York" /System/LocalDomain/
Services/CacheServices/MAGADIEUXP.
sybase.com MyDBConnector.MyDBOP
```

avaki dbop --delete

Remove an Avaki database operation.

To use this command, you must be a member of the DatabaseAdministrators group.

Syntax

```
avaki dbop --delete <dbconn-name> <dbop-name>
```

<code><dbconn-name></code>	The name of the Avaki database connector that this database operation uses.
<code><dbop-name></code>	The name of the database operation to be deleted.

avaki dbop --delete-schedule

Delete a caching schedule for an Avaki database operation.

To use this command, you must have write permission on the cache service.

Syntax

```
dbop --delete-schedule <cache-service-path>
  <qualified-dbop-name> [<parameter-value>]+ <schedule-id>
```

<cache-service-path>	The grid path of a cache service.
<qualified-dbop-name>	The name of a database operation, in this format: [<domain-name>.]<database-connector-name>.<database-operation-name>
<parameter-value>+	Optional. One or more parameter values required by the target database operation. The parameter values must match the values that were specified when the schedule was created. The parameter values will be processed in the order specified in the database operation's query. Any parameter value that includes a space must be enclosed in quotation marks. Any string at the end of the command that is enclosed in quotation marks will be interpreted as a single value. (That is, enclose each parameter value in its own pair of quotation marks.)
<schedule-id>	The numeric ID of the schedule to delete. To determine a schedule's ID, use avaki dbop --list-schedules (page 91).

avaki dbop --depends

Show the dependencies to and from a database operation.

To use this command, you must have read permission on the message service on the grid server where the database operation is located, and you must have read permissions on the database operation.

Syntax

```
avaki dbop --depends {VirtualDb | <dbconn-name>} <dbop-name>
```

VirtualDb |
<dbconn-name>

If you are showing the dependencies for a virtual SQL view that does a query against the query engine, enter the string VirtualDb. Otherwise, specify the name of the Avaki database connector for which you are showing the dependencies.

The name of the Avaki virtual database or database connector that the database operation uses.

<dbop-name>

The name of the database operation for which dependencies will be shown.

avaki dbop --execute

Run an Avaki database operation.

To use this command, you must have execute permission on the database operation.

Syntax

```
avaki dbop --execute [{--xmlinputfile=<grid-file-path> |
  --localxmlinputfile=<local-file-path>}] [--generate-schema]
  {VirtualDb | [<dbconn-name>] <dbop-name> [<parameter-value>]+
```

`--xmlinputfile=<grid-file-path>` Optional. If the database operation requires an input file, use this option to specify the grid path to the file. For details about the format for an XML input file, see the *Sybase Avaki EII Provisioning and Advanced Data Integration Guide*

`--localxmlinputfile=<local-file-path>` Optional. If the database operation requires an input file, use this option to specify the local path to the file. For details about the format for an XML input file, see the *Sybase Avaki EII Provisioning and Advanced Data Integration Guide*

`--generate-schema` Optional. Generate an XSD schema file for the database operation and suppress output. The schema file describes the structure of the output of a database operation. The generated file will be stored in the following path and file:

```
/Metadata/Database/<dbconn-name>/<dbop-name>/
<dbopname>.xsd
```

`VirtualDb | <dbconn-name>` If you are running a virtual database operation that does a query against the query engine, enter the string VirtualDb.Otherwise, specify the name of the Avaki database connector for which you are running a database operation.

<dbop-name>	The name of the database operation to run.
<parameter-value>+	Optional. One or more parameter values required by the target database operation. The parameter values will be processed in the order specified in the database operation's query. Any parameter value that includes a space must be enclosed in quotation marks. Any string at the end of the command that is enclosed in quotation marks will be interpreted as a single value. (That is, enclose each parameter value in its own pair of quotation marks.)

Examples

This command runs a database operation called emp-info. It uses the database connector sybase1 and the parameter value "barney rubble".

```
$ avaki dbop --execute sybase1 emp-info "barney rubble"
```

This command runs a database operation called updatesals. It uses the database connector sybase1 and the input file c:\newsalaries.xml in the local file system.

```
$ avaki dbop --execute
--localxmlinputfile=c:\newsalaries.xml sybase1 updatesals
```

avaki dbop --generate-sql-view

Generate a SQL view from a database operation. The SQL view is placed in the following data catalog directory:

```
System/Domains/<domain-name>/Services/DatabaseServices/<database-connector-name>/VirtualSQLViews/
```

Note You must generate a schema before generating a SQL view. You can use **[avaki dbop --execute](#)** (page 81) to generate a schema.

Syntax

```
avaki dbop --generate-sql-view [--provision-as=<table-name>]
  {VirtualDb | <dbconn-name>} <dbop-name>
```

`--provision-as=<table-name>` Optional. Specify a name for the SQL view. If no name is specified, the SQL view will have the same name as `<dbop-name>`.

`VirtualDb | <dbconn-name>` If you create a virtual SQL view that does a query against the query engine, enter the string `VirtualDb`. Otherwise, specify the name of the Avaki database connector for which you are generating a SQL view.

`<dbop-name>` The name of the database operation for which you are generating a SQL view.

avaki dbop --info

Display information about a database operation, including the name, description, SQL statement, parameters, and other options set when the database operation was created.

To use this command, you must be a member of the DatabaseAdministrators group.

Syntax

```
avaki dbop --info <dbconn-name> <dbop-name>
```

`<dbconn-name>` The name of the Avaki database connector that the target database operation uses.

`<dbop-name>` The name of a database operation.

avaki dbop --jdbc

Create or modify an Avaki database operation. (A database connector for this database operation must previously have been created, using either the web UI or the command **avaki dbconn --jdbc --create** (page 66)).

To use this command, you must be a member of the DatabaseAdministrators group.

Syntax

```
avaki dbop --jdbc {--create | --update}
  [--callable={true | false}] [--batch={true | false}]
  [--expiration=<expiration-time-in-secs>]
  [--params={in | out | inout}:<parameter-types>]
  [--description=<description-text>]
  [--run-user={<user-grid-path> | <qualified-user-name> |
  null}] [--sqlinputfile] <sql> <dbconn-name> <dbop-name>
```

<code>--create</code>	Create a new database operation.
<code>--update</code>	Modify an existing database operation.
<code>--callable={true false}</code>	Optional. Set this option to true if the SQL statement <code><sql></code> invokes a stored procedure in the underlying database. Otherwise, set to false. Default: false
<code>--batch={true false}</code>	Optional. Set this option to true if the SQL statement <code><sql></code> is an update that can be used with JDBC batch mode. Otherwise, set to false. Default: false
<code>--expiration=<expiration-time-in-secs></code>	Optional. Specify the interval, in seconds, after which the results of this database operation expire from the on-demand cache. If you set this option to <code>-1</code> , the results never expire. If you set it to <code>0</code> , the data is not cached. Default: 0 (results are not cached) Range: 0 to 2147483647


```
--params={in | out |
inout}:<parameter-
types>
```

Optional. If the database operation accepts input parameters or produces output parameters, use this option to specify the data type for each input and output. There is no default.

Use in for input parameters, out for output parameters, and inout for parameters that can be both. If you specify multiple parameters, use semicolons (;) to separate them.

You can use the following <parameter-types>:

BIGINT	LONGVARBINARY
BINARY	LONGVARCHAR
BIT	NUMERIC
BOOLEAN	ORACLE_CURSOR
BLOB	OTHER
CHAR	REAL
CLOB	SMALLINT
DATE	TIME
DECIMAL	TIMESTAMP
DOUBLE	TINYINT
FLOAT	VARBINARY
INTEGER	VARCHAR

```
--description=
<description-
text>
```

Optional. Enter details about this database operation; enclose the entire text string in quotation marks. The description is shown to users who display information about the database operation (see [avaki dbop --info on page 83](#)). Consider including an explanation of what this database operation is for, how to use it, or where to get more information. Length limit: 255 characters

```
--run-user={<user-  
grid-path> |  
<qualified-user-  
name> | null}
```

Specify the user to create or modify the database operation as. You can specify the data catalog path to the directory of the user account. For example:

```
/System/Domains/Bedrock/Services/  
AuthServices/Grid/  
DefaultAuthService/Users/wilma
```

Alternatively, you can specify the qualified user name using the following format:

```
<user>@<authservice>.<authservicetype>.  
<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user to create or modify the database operation as.
<authservice>	The name of the user's authentication service.
<authservicetype>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

If you specify null, the database operation is run as the current user.

```
--sqlinputfile
```

Optional. If --sqlinputfile is present, it indicates that the value specified in the <sql> argument is a path to a local file that contains a SQL statement that the database operation will execute. If --sqlinputfile is not present, <sql> consists of the SQL statement that the database operation will execute.

<sql>	<p>Enter one of the following:</p> <ul style="list-style-type: none"> • The SQL statement that the virtual database operation will execute. Enclose the statement in quotation marks, as shown in the examples below. • The path to a local file containing the SQL statement. If you choose this option, include <code>--sqlinputfile</code> in the command.
<dbconn-name>	The name of the Avaki database connector that will give this database operation access to a relational database.
<dbop-name>	The name of the database operation. The name must not exceed 64 characters in length and may not contain any period characters (.) or spaces.

Examples

This command creates a database operation called `gazoobirthdays`, which uses the database connector `gazoconn` to extract birth dates from an employee database.

```
$ avaki dbop --jdbc --create "select * from emp" gazoconn
gazoobirthdays
```

In the next example, the `avaki dbop --jdbc` command creates a database operation called `employee`, which uses the database connector `sybase1` to extract employee records from a database. The `--params` option specifies that the database operation takes two input parameters, a varchar (the employee's name) and a numeric (the employee number).

The command `avaki dbop --execute` ([page 81](#)) runs the `employee` database operation, looking for the record of an employee named Barney whose employee number is 1234.

```
$ avaki dbop --jdbc --create --params=in:VARCHAR;in:NUMERIC
"select * from emp where ename=? AND empno=?" sybase1 employee

$ avaki dbop --execute sybase1 employee Barney 1234
```

avaki dbop --jdbc --create-virtual-dbop

Create an Avaki virtual database operation.

Syntax

```
avaki dbop --jdbc --create-virtual-dbop
  [--expiration=<expiration-time-in-secs>]
  [--params={in | out | inout}:<parameter-types>]
  [--description=<description-text>]
  [--run-user={<qualified-user-name> | null}] [--sqlinputfile]
  [--virtual-database=<virtual-database-path>]
  [{--server | -s}=<grid-server>] <sql> <dbop-name>
```

`--expiration=<expiration-time-in-secs>` Optional. Specify the interval in seconds after which the results of this virtual database operation expire from the on-demand cache. If you set this option to `-1`, the results never expire. If you set it to `0`, the data is not cached. Default: `0` (results are not cached) Range: `0` to `2147483647`

`--params={in | out | inout}:<parameter-types>` Optional. If the virtual database operation accepts input parameters or produces output parameters, use this option to specify the data type for each input and output. There is no default.

Use `in` for input parameters, `out` for output parameters, and `inout` for parameters that can be both. If you specify multiple parameters, use semicolons (`;`) to separate them.

You can use the following `<parameter-types>`:

BIGINT	LONGVARBINARY
BINARY	LONGVARCHAR
BIT	NUMERIC
BOOLEAN	ORACLE_CURSOR
BLOB	OTHER
CHAR	REAL
CLOB	SMALLINT
DATE	TIME
DECIMAL	TIMESTAMP
DOUBLE	TINYINT
FLOAT	VARBINARY
INTEGER	VARCHAR

```
--description=  
<description-  
text>
```

Optional. Enter details about this virtual database operation; enclose the entire text string in quotation marks. The description is shown to users who display information about the database operation (see **avaki dbop --info** on page 83). Consider including an explanation of what this virtual database operation is for, how to use it, or where to get more information. Length limit: 255 characters

```
--run-user={<user-  
grid-path> |  
<qualified-user-  
name> | null}
```

Specify the user to create or modify the virtual database operation as. You can specify the data catalog path to the directory of the user account. For example:

```
/System/Domains/Bedrock/Services/  
AuthServices/Grid/  
DefaultAuthService/Users/wilma
```

Alternatively, you can specify the qualified user name using the following format:

```
<user>@<authservice>.<authservicetype>.  
<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user to create or modify the virtual database operation as.
<authservice>	The name of the user's authentication service.
<authservicetype>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

If you specify null, the virtual database operation is run as the current user.

<code>--sqlinputfile</code>	Optional. If <code>--sqlinputfile</code> is present, it indicates that the value specified in the <code><sql></code> argument is a path to a local file that contains a SQL statement that the database operation will execute. If <code>--sqlinputfile</code> is not present, <code><sql></code> consists of the SQL statement that the database operation will execute.
<code>--virtual-database= <virtual-database- path></code>	Optional. The data catalog path to the virtual database service, in this format: <code>/System/LocalDomain/Services/ VirtualDatabaseServices/Services/ <grid-server></code>
<code>{--server -s}= <grid-server></code>	Optional. Specify the name of the Avaki server on which the virtual database operation will be created. For example, Bedrock. If you omit this option, the virtual database operation is created on the current server.
<code><sql></code>	Enter one of the following: <ul style="list-style-type: none"> • The SQL statement that the virtual database operation will execute. Enclose the statement in quotation marks. • The path to a local file containing the SQL statement. If you choose this option, include <code>--sqlinputfile</code> in the command.
<code><dbop-name></code>	The name of the virtual database operation. The name must not exceed 64 characters in length and may not contain any period characters (.) or spaces.

avaki dbop --list-schedules

List the caching schedules for a database operation.

Syntax

```
avaki dbop --list-schedules <cache-service-path>
  <qualified-dbop-name> [<parameter-value>]+
```

<cache-service-path>

The data catalog path of a cache service.

<qualified-dbop-name>

The name of a database operation, in this format:

```
[<domain-name>.]<database-connector-name>.<database-operation-name>
```

<parameter-value>+

Optional. One or more parameter values required by the target database operation. The parameter values will be processed in the order specified in the database operation's query. Any parameter value that includes a space must be enclosed in quotation marks. Any string at the end of the command that is enclosed in quotation marks will be interpreted as a single value. (That is, enclose each parameter value in its own pair of quotation marks.)

avaki dgas --add-admission-policy

Add a rule to the admission policy table for a data grid access server. To use this command, you must have execute permission on the target DGAS. (Use **avaki dgas --show-admission-policies** (page 114) to display the admission policies currently in effect. Use **avaki dgas --delete-admission-policy** (page 98) to remove rules from the policy table.)

An admission policy table is made up of an ordered list of rules that control which NFS and CIFS client machines may access the Avaki domain through this DGAS. The DGAS decides to permit or deny access by a client machine by matching the client's address with the policy rules. If the client address (either host or network) does not match the first rule, the DGAS looks at the next rule, continuing down the list until a match is found or the list ends. If a match is found (either admitting or denying the client), the DGAS acts accordingly. If no match is found in the rules but a default rule is defined, the client is admitted. If no match is found and no default is defined, the client is denied access. Note that rules are interpreted and applied in the order in which they appear in the table, so later rules can supersede earlier ones. However, deny rules always override allow rules.

Syntax

Two syntaxes perform the same actions:

```
dgas --add-admission-policy <dgas-name> <position> [--replace]
  {--deny --client-host=<machine> | --allow
  --client-host=<machine> --auth-service-path=<path>}
```

```
dgas --add-admission-policy <dgas-name> <position> [--replace]
  {--deny <client-machine> | --allow <client-machine>
  <auth-service-path>}
```

<code><dgas-name></code>	The name of a data grid access server.
<code><position></code>	The row number in the policy table at which you are adding the new rule. The rules are searched in row-number order. Row numbers start with 0.
<code>--replace</code>	Optional. When you use <code>--replace</code> , the new policy replaces the policy at the row number specified by <code><position></code> . If you don't use <code>--replace</code> , the new policy is added at the end of the last policy.

--deny
<client-machine>
or
--deny --client-host=<machine>

--allow
<client-machine>
or
--allow --client-host=<machine>

<auth-service-path>
or
--auth-service-path=<path>

Deny access to the specified client machine or network. <client-machine> or <machine> are the IP address of the client or network in dotted decimal format (for example, 192.37.22.11 for a client or 192.37.22.0/24 for a network.).

Allow the specified client machine or network to access data through this DGAS. <client-machine> or <machine> are the IP address of the client or network in dotted decimal format (for example, 192.37.22.11 for a client or 192.37.22.0/24 for a network.).

Required when you use --allow. Specify the Avaki directory of the authentication service to be used for users and groups accessing this Avaki domain from the client(s) at the address specified in the --allow option. For example, /System/Domains/Bedrock/Services/AuthServices/Ldap/BedrockLdap.

avaki dgas --add-group-mapping

Map a group in your Avaki domain to a group ID (GID) on an NFS or CIFS client machine. This mapping is specific to the target DGAS and supersedes mappings configured in authentication services.

To use this command, you must have write permission on the target DGAS.

Syntax

```
avaki dgas --add-group-mapping <dgas-name> <gid> <client-host>
    <group-path>
```

<code><dgas-name></code>	The name of a data grid access server.
<code><gid></code>	The numeric group ID of the group on the client.
<code><client-host></code>	The IP address of the client machine in IPv4 dotted decimal format. (For example, 192.37.22.12.)
<code><group-path></code>	The data catalog path of the Avaki group.

avaki dgas --add-user-mapping

Map a user in your Avaki domain to a user ID (UID) on an NFS or CIFS client machine. This mapping is specific to the target DGAS and supersedes mappings configured in authentication services.

To use this command, you must have write permission on the target DGAS.

Syntax

```
avaki dgas --add-user-mapping <dgas-name> <uid> <client-host>
    <user-path>
```

<dgas-name>	The name of a data grid access server.
<uid>	The numeric user ID of the user on the client.
<client-host>	The IP address of the client machine in IPv4 dotted decimal format. (For example, 192.37.22.12.)
<user-path>	The data catalog path of the Avaki user.

avaki dgas --cifs-share-info

Display information about a CIFS share. To use this command, you must have read permission on the target CIFS share.

Syntax

```
avaki dgas --cifs-share-info <dgas-name> <share-name>
```

<dgas-name>	The name of a data grid access server.
<share-name>	The name of a CIFS share.

avaki dgas --clear-cached-credentials

Use this command when changes have occurred in authentication services, user accounts, group membership, user mappings, or group mappings.

To use this command, you need execute permission on the target DGAS.

Syntax

```
avaki dgas --clear-cached-credentials <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --create-cifs-share

Create a CIFS share to make data in the Avaki domain available to client machines.

To use this command, you must have execute permission on the target DGAS.

Caution Because the Windows file system on CIFS clients is case-insensitive, the <share-name> you assign to a CIFS share will be converted to all upper-case letters. Do not assign a share name that relies on case for uniqueness.

Syntax

```
avaki dgas --create-cifs-share <dgas-name> <share-name>
  <source-path> [--max-connections=<num>]
  [--enabled=<true | false>] [--comment=<string>]
```

<dgas-name>	The name of the data grid access server on which the CIFS share will reside.
<share-name>	A name for the CIFS share you are creating.
<source-path>	The path to the data catalog directory you are sharing out.
--max-connections=<num>	Optional. The maximum number of users that will be permitted to access this share at one time. To allow unlimited connections, set max connections to -1. Default: -1
--enabled=<true false>	Optional. Turn this share on (true) or off (false). Default: true (the share is enabled)
--comment=<string>	Optional. Provide useful information such as a description of this CIFS share, an explanation of why this CIFS share exists, or contact information for the person responsible for it. This description will be visible to users on the CIFS client.

avaki dgas --delete-admission-policy

On the target data grid access server, remove the rule at the specified position in the admission policy table, or remove all the rules in the table.

To use this command, you must have execute permission on the target DGAS.

Use ***avaki dgas --show-admission-policies*** (page 114) to display the admission policy table.

Syntax

```
avaki dgas --delete-admission-policy <dgas-name> {<position> |  
--all}
```

<dgas-name>	The name of a data grid access server.
<position>	The row number of the rule to be removed.
--all	Remove all the rules in the admission policy table.

avaki dgas --delete-cache

Remove a specified file or directory from the internal cache of a data grid access server.

To use this command, you must have execute permission on the target DGAS.

Syntax

```
avaki dgas --delete-cache <dgas-name> [--target-path=<path>]
  [-r | --recursive]
```

<dgas-name>	The name of a data grid access server.
--target-path=<path>	Optional. Specify the data catalog path of the file or directory to be removed from the cache. If you omit this option, everything is removed from the cache.
-r --recursive	Optional. If --target-path specifies a directory, recursively remove the directory and all its contents.

avaki dgas --delete-cifs-share

Remove a CIFS share.

To use this command, you must have delete permission on the target CIFS share.

Syntax

```
avaki dgas --delete-cifs-share <dgas-name> <share-name>
```

<dgas-name>	The name of a data grid access server.
<share-name>	The name of the CIFS share to be deleted.

avaki dgas --delete-group-mapping

Delete the mapping between the specified group ID on the specified client and the Avaki group it is mapped to. (That is, delete a mapping created by [avaki dgas --add-group-mapping](#) (page 94). This command does not affect mappings created in authentication services.) Use [avaki dgas --list-group-mappings](#) (page 106) to display a list of group mappings for a DGAS.

To use this command, you must have write permission on the target DGAS.

Syntax

```
avaki dgas --delete-group-mapping <dgas-name> [<client-machine>
  [<gid>]]
```

<code><dgas-name></code>	The name of a data grid access server.
<code><gid></code>	Optional if you use the <code><client-machine></code> option. Specify the group ID of the group on the client whose Avaki group mapping should be deleted. If you omit this option, the command deletes all group mappings for the specified client.
<code><client-machine></code>	The IP address of the client machine in IPv4 dotted decimal format. (For example, 192.37.22.12.)

avaki dgas --delete-user-mapping

Delete the mapping between the specified user ID on the specified client and the Avaki user it is mapped to. (That is, delete a mapping created by [avaki dgas --add-user-mapping](#) (page 95). This command does not affect mappings created in authentication services.) Use [avaki dgas --list-user-mappings](#) (page 107) to display a list of user mappings for a DGAS.

To use this command, you must have write permission on the target DGAS.

Syntax

```
avaki dgas --delete-user-mapping <dgas-name> [<client-machine>
  [<uid>]]
```

<dgas-name>	The name of a data grid access server.
<uid>	Optional if you use the <client-machine> option. Specify the user ID of the user on the client. If you omit this option, the command deletes all user mappings for the specified client.
<client-machine>	The IP address of the client machine in IPv4 dotted decimal format. (For example, 192.37.22.12.)

avaki dgas --disconnect-cifs-client

Disconnect a CIFS client from a data grid access server. If the client has more than one connection to the share, all connections are dropped.

To use this command, you must have execute permission on the target DGAS.

Syntax

```
avaki dgas --disconnect-cifs-client <dgas-name> <share-name>
  <client-host>
```

<dgas-name>	The name of the data grid access server on which the CIFS share was created.
<share-name>	The name of the CIFS share to be disconnected.
<client-host>	The DNS name or IP address of the client from which you are disconnecting the CIFS share.

avaki dgas --get-cache-size

Display the current size in bytes of the internal cache for the specified data grid access server.

To use this command, you must have read permission on the target DGAS.

Syntax

```
avaki dgas --get-cache-size <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --get-cache-statistics

Display statistics for the internal cache of the specified data grid access server. Statistics include cache hit and miss totals.

To use this command, you must have read permission on the target DGAS.

Syntax

```
avaki dgas --get-cache-statistics <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --get-free-disk-space

Display the amount of free disk space on the device on which the specified data grid access server's cache directory resides.

Syntax

```
avaki dgas --get-free-disk-space <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --get-properties

For the specified data grid access server, list all DGAS properties and their values.

To use this command, you must have read permission on the target DGAS.

To change DGAS property values, use [avaki dgas --set-property](#) (page 114).

Syntax

```
avaki dgas --get-properties <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --get-property

For the specified data grid access server, display the value of the specified property.

To use this command, you must have read permission on the target DGAS.

To change DGAS property values, use [avaki dgas --set-property](#) (page 114).

Syntax

```
avaki dgas --get-property <dgas-name> <prop-name>
```

<dgas-name> The name of a data grid access server.

<prop-name> The name of a DGAS property.

avaki dgas --get-property-list

Display descriptions of all DGAS properties, including default values and other information.

To use this command, you must have read permission on the target DGAS.

Syntax

```
avaki dgas --get-property-list <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --initialize

Initialize credentials for a new data grid access server. (For an outline of the tasks involved in setting up a DGAS, see [avaki server --dgas --connect on page 224](#).)

Syntax

```
avaki dgas --initialize <dgas-name>
  [--properties-file=<filename>]
```

<dgas-name>	A name for the new data grid access server.
--properties-file= <filename>	Optional. Specify the local path to the properties file the new DGAS will use. If you omit this option, the DGAS uses default values for all properties. Use avaki dgas --get-property-list (above) to display descriptions and default values of DGAS properties.

avaki dgas --list-cifs-clients

Display a list of addresses of client machines that are connected to the specified CIFS share.

To use this command, you must have read permission on the target DGAS.

Syntax

```
avaki dgas --list-cifs-clients <dgas-name> <share-name>
```

<dgas-name>	The name of the data grid access server on the which the CIFS share was created.
<share-name>	The name of the CIFS share whose connected clients you want to list.

avaki dgas --list-cifs-shares

Display a list of the CIFS shares that reside on the specified data grid access server.

To use this command, you must have read permission on the target DGAS.

Syntax

```
avaki dgas --list-cifs-shares <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --list-group-mappings

Display a list of the mappings between grid groups and client GIDs that are configured for the specified data grid access server. (Mappings configured in authentication services are not included.)

To use this command, you must have read permission on the target DGAS.

Syntax

```
avaki dgas --list-group-mappings <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --list-user-mappings

Display a list of the mappings between grid user accounts and client UIDs that are configured for the specified data grid access server. (Mappings configured in authentication services are not included.)

To use this command, you must have read permission on the target DGAS.

Syntax

```
avaki dgas --list-user-mappings <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --modify-cifs-share

Change a CIFS share.

To use this command, you must have write permission on the target CIFS share.

Syntax

```
avaki dgas --modify-cifs-share <dgas-name> <share-name>
  [--name=<string>] [--max-connections=<num>]
  [--enabled=<true | false>] [--comment=<string>]
```

<dgas-name> The name of the data grid access server on which the CIFS share resides.

<share-name> The name of the CIFS share you are modifying.

--name=<string> Optional. If you are changing the name of the CIFS share, use this option to specify a new name.

<code>--max-connections= <num></code>	Optional. The maximum number of users that will be permitted to access this share at one time. A setting of <code>-1</code> allows unlimited connections. Default: <code>-1</code>
<code>--enabled= <true false></code>	Optional. Turn this share on (<code>true</code>) or off (<code>false</code>). Default: <code>true</code> (the share is enabled)
<code>--comment=<string></code>	Optional. Provide useful information such as a description of this CIFS share, an explanation of why this CIFS share exists, or contact information for the person responsible for it. This description will be visible to users on the CIFS client.

avaki dgas --read-log-properties

Force the target data grid access server to reread its `log4j` properties file, which is located in the local directory where the DGAS was started (that is, where the `dgas --start` command was executed).

To use this command, you must have execute permission on the target DGAS.

Syntax

```
avaki dgas --read-log-properties <dgas-name>
```

`<dgas-name>` The name of a data grid access server.

avaki dgas --reset-cache-statistics

Reset all statistics counters for the internal cache of the specified data grid access server.

To use this command, you must have write permission on the target DGAS. (Use [avaki dgas --get-cache-statistics \(page 102\)](#) to display cache statistics.)

Syntax

```
avaki dgas --reset-cache-statistics <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --save-cache

Save a copy of a data grid access server's internal cache to a specified local file.

To use this command, you must have execute permission on the target DGAS.

Syntax

```
avaki dgas --save-cache <dgas-name> <local-dest>
```

<dgas-name> The name of a data grid access server.

<local-dest> A directory name for the saved copy of the cache. The system creates the directory in the local cache directory of the target DGAS. The destination directory must not exist; the command creates it.

Example

Suppose the cache directory of the DGAS called Dinosdgas is /home/avaki/dgas/cache. If you enter this command, a copy of the cache is saved to /home/avaki/dgas/cache/2003-6-11:

```
$ avaki dgas --save-cache Dinosdgas 2003-6-11
```

avaki dgas --self-map

Creates a temporary dynamic mapping between the Avaki user who issues the command and the UID of the Unix account under which the issuing user is logged in. When a user is mapped to a UID, the reference count for that mapping is set to 1. The reference count is incremented each time a user is mapped to that UID. The reference count keeps track of how many jobs or processes are using the mapping. Unless you specify otherwise, dynamic mappings expire after 6 hours. Mappings created with this command are effective only on the DGAS specified by `<dgas-name>`.

User mappings allow local users to access data in the Avaki domain as mapped Avaki users. The dynamic mappings created with this command are especially useful when you want to use Avaki compute grid to run jobs that access data through the Avaki domain. Dynamic mappings are also useful in environments where IP addresses are assigned dynamically.

You can map the same Avaki user to several UIDs as long as the UIDs are on different machines. Two local accounts on the same machine cannot be mapped to the same Avaki user.

This command is available only on Unix systems.

Note You can't create a dynamic user mapping for a UID for which an Avaki administrator has already configured a DGAS-specific user mapping.

Note We recommend against using this command on a machine with more than one IP address. The mapping uses the machine's IP address to ensure that the UID is unique. On a multihomed machine, it is not possible to predict which IP address the mapping will use, which can lead to unexpected results.

Syntax

```
avaki dgas --self-map <dgas-name> [--time-to-live=<minutes>]
```

<code><dgas-name></code>	The name of a data grid access server.
<code>--time-to-live=<minutes></code>	Optional. Specify the number of minutes this user-to-UID mapping will last. To set up a mapping that lasts indefinitely, set <code>--time-to-live=-1</code> . Default: 360 (6 hours)

avaki dgas --self-unmap

Removes (or decrements the reference count for) dynamic user mappings for the issuing user created with **avaki dgas --self-map** (above).

You can configure several compute grid jobs to use a single dynamic user mapping. To do this, execute **avaki dgas --self-map** once for each job. Then include an **avaki dgas --self-unmap** command at the end of each job. As the jobs finish running, each one decrements the reference count, until the last job removes the user mapping.

Syntax

```
avaki dgas --self-unmap <dgas-name> [--force]
```

<code><dgas-name></code>	The name of a data grid access server.
<code>--force</code>	Optional. Removes the mapping for the Avaki user who issues the command even if the reference count for the mapping is greater than 1.

avaki dgas --set-admission-policy-domain

Configure Windows domain information for the LDAP authentication service associated with an admission policy allow rule. (The target rule must not be a deny rule, because deny rules do not have associated authentication services.)

The DGAS supports pass-through authentication for LDAP authentication services. (That is, the DGAS forwards user names and passwords to a specified Windows domain controller for authentication.) If you do not set an admission policy domain, the DGAS uses the LDAP server for this grid domain's LDAP authentication service in place of the Windows domain controller.

To use this command, you must have execute permission on the target DGAS. To remove Windows domain information from an admission policy rule, use ***avaki dgas --unset-admission-policy-domain*** (page 117). To display Windows domain information for an admission policy rule, use ***avaki dgas --show-admission-policy-domain*** (page 115).

Syntax

```
avaki dgas --set-admission-policy-domain <dgas-name> <position>
    <domain-name> <domain-controller>+
```

<code><dgas-name></code>	The name of a data grid access server.
<code><position></code>	The row number in the admission policy table of the allow rule for which you're configuring domain information. (All allow rules are associated with grid authentication services.)
<code><domain-name></code>	A Windows domain name.
<code><domain-controller>+</code>	The names of one or more Windows domain controllers for the specified Windows <code><domain-name></code> .

avaki dgas --set-default-admission-policy

Add or replace a default rule in the admission policy table for the specified data grid access server. The default rule is always the last one to be processed; it allows all clients who do not match any of the other rules to access the Avaki domain. (If you want access to be denied to all clients who are not specifically allowed by other rules, do not configure a default rule.) To use this command, you must have execute permission on the target DGAS. You can use [avaki dgas --delete-admission-policy](#) (page 98) to remove the default rule or any other rule from the policy table. Use [avaki dgas --show-admission-policies](#) (page 114) to display all the rules.

Syntax

```
avaki dgas --set-default-admission-policy <dgas-name>  
      <auth-service-path>
```

<dgas-name> The name of a data grid access server.

<auth-service-path> The path to the grid authentication service that will provide user and group mappings for users of clients admitted by this rule.

avaki dgas --set-property

Set a property on a data grid access server.

To use this command, you must have write permission on the target DGAS.

Use [avaki dgas --get-property-list \(page 104\)](#) to display a list of DGAS properties and their values.

Syntax

```
avaki dgas --set-property <dgas-name> <prop-name> <value>
```

<dgas-name> The name of a data grid access server.

<prop-name> The name of a DGAS property.

<value> The value you are assigning to <prop-name> for this DGAS.

avaki dgas --show-admission-policies

Display, in the configured order, the admission policies for the specified data grid access server. To use this command, you must have execute permission on the target DGAS.

Syntax

```
avaki dgas --show-admission-policies <dgas-name>
```

<dgas-name> The name of a data grid access server.

avaki dgas --show-admission-policy-domain

Display Windows domain information for an authentication service in an admission policy rule for a data grid access server. The target rule must be a “deny” rule.

To use this command, you must have read permission on the target DGAS.

Syntax

```
avaki dgas --show-admission-policy-domain <dgas-name>  
    <position>
```

<dgas-name> The name of a data grid access server.

<position> The row number in the admission policy table of the deny rule for which you’re displaying domain information.

avaki dgas --sync-cache

Check whether the DGAS's cached copy of the target file or directory is the most recent version. If the cached copy is not current or if no copy exists in the DGAS internal cache, the DGAS fetches the most recent version and caches it. (If a cache service for this DGAS has been configured, the DGAS tries to get the file or directory from the cache. If there is no cache or the cache does not have a current copy of the file or directory, the DGAS fetches it from the source file system.)

To use this command, you must have execute permission on the target DGAS and read permission on the target files and directories.

Syntax

```
avaki dgas --sync-cache <dgas-name> <target-path> [-r |  
--recursive]
```

<dgas-name>	The name of a data grid access server.
<target-path>	The grid path of a file or directory to be fetched into the cache.
-r --recursive	Optional. If the target path is a directory, recursively fetch and cache all contents.

avaki dgas --unset-admission-policy-domain

Remove Windows domain information from an admission policy rule for a data grid access server. The target rule must be a “deny” rule.

To use this command, you must have execute permission on the target DGAS.

Syntax

```
avaki dgas --unset-admission-policy-domain <dgas-name>  
    <position>
```

<dgas-name> The name of a data grid access server.

<position> The row number in the admission policy table of the deny rule from which you’re removing domain information.

avaki dgas --unset-property

Unset the specified data grid access server property (that is, change the value to “unset”). When a DGAS property is unset, its value returns to the default setting.

Syntax

```
avaki dgas --unset-property <dgas-name> <prop-name>
```

<dgas-name> The name of a data grid access server.

<prop-name> The name of the DGAS property that you are unsetting.

avaki directory --add-schedule

Add a new pin or cache schedule to a currently cached Avaki directory or directory service. For information about tagging an Avaki directory or a directory hierarchy for caching within a cache service, see [avaki directory --cache \(page 122\)](#).

To use this command, you must have write permission on the cache service, and you must have read permission on the directory.

Syntax

```
avaki directory --add-schedule --action={directory-only |
tree-on-demand | tree-pinned}
{--periodic-schedule=<period-spec> |
--cron-schedule=<cron-spec> |
--one-time-schedule=<time-stamp>}
[--start-time=<time-stamp>] [[--end-time=<time-stamp>] |
[--max-iterations=<num-iterations>]]
[--time-zone=<time-zone>] [--exclusion=<exclusion-name>]+
[--exclusion-action={ drop | sooner | later }]
[--reschedule-offset=<offset-spec>] <cache-service-path>
<directory-path>+
```

<code>--action</code>	Specify how to cache the directory. The formats are as follows:
<code><directory-only></code>	Pin a directory but not its contents for scheduled caching.
<code><tree-on-demand></code>	Tag a directory and its contents for on-demand caching, in which files or file portions can be cached as needed.
<code><tree-pinned></code>	Pin a directory and its contents for scheduled caching.

`--periodic-schedule=<period-spec>` Specify a periodic schedule for refreshing the cached directory data. Use the following syntax:

```
--periodic-schedule="<number-of-units>
<unit-type>"
```

`<number-of-units>` An integer from 1 to 2147483648

`<unit-type>` seconds, minutes, hours, days, weeks, or months

In the following example, the periodic schedule is set to 30 seconds:

```
--periodic-schedule="30 seconds"
```

`--cron-schedule=<cron-spec>` Create an advanced schedule using a cron expression that tells the cache service how often to refresh cached directory data. Use the following syntax:

```
--cron-schedule="<seconds> <minutes> <hours>
<days-of-month> <months> <days-of-week>
[<years>] "
```

For details about the `<cron-spec>` syntax, see [Appendix C, “Configuring cron schedules”](#).

`--one-time-schedule=<time-stamp>` Specify a single time and date for refreshing the cached directory data. Use the following syntax:

```
--one-time-schedule="{<hh>:<mm>:<ss> {am | pm}
<mm>-<dd>-<yyyy> | now}"
```

`--start-time=<time-stamp>` Specify the time and date that the cache service will start refreshing cached directory data, in this format:

```
{<hh>:<mm>:<ss> {am | pm} <mm>-<dd>-<yyyy> |
now}
```

`--end-time=<time-stamp>` Specify the time and date that the cache service will stop refreshing cached directory data, in this format:

```
{<hh>:<mm>:<ss> {am | pm} <mm>-<dd>-<yyyy> |
now}
```

<code>--max-iterations= <num-iterations></code>	Specify the number of times a schedule will repeat. <num-iterations> An integer from 1 to 2147483648. If no value is specified, the schedule repeats indefinitely.																																
<code>--time-zone= <time-zone></code>	Specify the time zone for the schedule. The schedule can be specified according to the cache server's time zone or relative to Greenwich Mean Time (GMT). If no value is specified for the time zone, the server's local time zone is used. The format for the time zone is as follows: <table> <tr><td>GMT+12:00 Fiji</td><td>GMT+02:00 Cairo</td></tr> <tr><td>GMT+11:00 Noumea</td><td>GMT+01:00 Paris</td></tr> <tr><td>GMT+10:00 Sydney</td><td>GMT London</td></tr> <tr><td>GMT+9:30 Adelaide</td><td>GMT-01:00 Azores</td></tr> <tr><td>GMT+09:00 Tokyo</td><td>GMT-02:00 Mid-Atlantic</td></tr> <tr><td>GMT+08:00 Hong Kong</td><td>GMT-03:00 Rio de Janeiro</td></tr> <tr><td>GMT+07:00 Bangkok</td><td>GMT-03:30 Newfoundland</td></tr> <tr><td>GMT+06:30 Rangoon</td><td>GMT-04:00 Caracas</td></tr> <tr><td>GMT+06:00 Dacca</td><td>GMT-05:00 New York</td></tr> <tr><td>GMT+05:45 Katmandu</td><td>GMT-06:00 Chicago</td></tr> <tr><td>GMT+05:30 Calcutta</td><td>GMT-07:00 Denver</td></tr> <tr><td>GMT+05:00 Karachi</td><td>GMT-08:00 San Francisco</td></tr> <tr><td>GMT+04:30 Kabul</td><td>GMT-09:00 Juneau</td></tr> <tr><td>GMT+04:00 Baku</td><td>GMT-10:00 Hawaii</td></tr> <tr><td>GMT+03:30 Tehran</td><td>GMT-11:00 Samoa</td></tr> <tr><td>GMT+03:00 Moscow</td><td></td></tr> </table>	GMT+12:00 Fiji	GMT+02:00 Cairo	GMT+11:00 Noumea	GMT+01:00 Paris	GMT+10:00 Sydney	GMT London	GMT+9:30 Adelaide	GMT-01:00 Azores	GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic	GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro	GMT+07:00 Bangkok	GMT-03:30 Newfoundland	GMT+06:30 Rangoon	GMT-04:00 Caracas	GMT+06:00 Dacca	GMT-05:00 New York	GMT+05:45 Katmandu	GMT-06:00 Chicago	GMT+05:30 Calcutta	GMT-07:00 Denver	GMT+05:00 Karachi	GMT-08:00 San Francisco	GMT+04:30 Kabul	GMT-09:00 Juneau	GMT+04:00 Baku	GMT-10:00 Hawaii	GMT+03:30 Tehran	GMT-11:00 Samoa	GMT+03:00 Moscow	
GMT+12:00 Fiji	GMT+02:00 Cairo																																
GMT+11:00 Noumea	GMT+01:00 Paris																																
GMT+10:00 Sydney	GMT London																																
GMT+9:30 Adelaide	GMT-01:00 Azores																																
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic																																
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro																																
GMT+07:00 Bangkok	GMT-03:30 Newfoundland																																
GMT+06:30 Rangoon	GMT-04:00 Caracas																																
GMT+06:00 Dacca	GMT-05:00 New York																																
GMT+05:45 Katmandu	GMT-06:00 Chicago																																
GMT+05:30 Calcutta	GMT-07:00 Denver																																
GMT+05:00 Karachi	GMT-08:00 San Francisco																																
GMT+04:30 Kabul	GMT-09:00 Juneau																																
GMT+04:00 Baku	GMT-10:00 Hawaii																																
GMT+03:30 Tehran	GMT-11:00 Samoa																																
GMT+03:00 Moscow																																	

<pre>--exclusion= <exclusion-name></pre>	<p>The name of an exclusion to this schedule. For information about creating schedule exclusions, see avaki scheduleexclusion --create --daily (page 201), avaki scheduleexclusion --create --weekly (page 205), avaki scheduleexclusion --create --monthly (page 203), avaki scheduleexclusion --create --yearly (page 207), and avaki scheduleexclusion --create --custom (page 199).</p>						
<pre>--exclusion- action={drop sooner later}</pre>	<p>Specify the action to take if the cache service is scheduled to refresh cached directory data during an exclusionary period.</p> <table border="0"> <tr> <td data-bbox="618 487 672 517">drop</td> <td data-bbox="958 487 1358 586">Don't refresh the cached directory data during the exclusionary period.</td> </tr> <tr> <td data-bbox="618 638 696 661">sooner</td> <td data-bbox="958 638 1339 737">Reschedule the cache service to refresh the cached directory data before the exclusionary period.</td> </tr> <tr> <td data-bbox="618 782 672 805">later</td> <td data-bbox="958 782 1339 881">Reschedule the cache service to refresh the cached directory data after the exclusionary period.</td> </tr> </table>	drop	Don't refresh the cached directory data during the exclusionary period.	sooner	Reschedule the cache service to refresh the cached directory data before the exclusionary period.	later	Reschedule the cache service to refresh the cached directory data after the exclusionary period.
drop	Don't refresh the cached directory data during the exclusionary period.						
sooner	Reschedule the cache service to refresh the cached directory data before the exclusionary period.						
later	Reschedule the cache service to refresh the cached directory data after the exclusionary period.						
<pre>--reschedule- offset= <offset-spec></pre>	<p>Specify the amount of time before or after the exclusionary period to perform a rescheduled cache refresh. Use the following syntax:</p> <pre>--reschedule-offset="<number-of-units> <unit-type>"</pre> <table border="0"> <tr> <td data-bbox="618 1164 886 1187"><number-of-units></td> <td data-bbox="958 1164 1350 1187">An integer from 1 to 2147483648</td> </tr> <tr> <td data-bbox="618 1246 793 1269"><unit-type></td> <td data-bbox="958 1246 1315 1303">Seconds, minutes, hours, days, weeks, or months</td> </tr> </table>	<number-of-units>	An integer from 1 to 2147483648	<unit-type>	Seconds, minutes, hours, days, weeks, or months		
<number-of-units>	An integer from 1 to 2147483648						
<unit-type>	Seconds, minutes, hours, days, weeks, or months						
<pre><cache-service- path></pre>	The grid path of a cache service.						
<pre><directory-path>+</pre>	The grid path of one or more directories to which a pin schedule will be added.						

This example pins the directory Water Buffalo Lodge and its contents for scheduled caching. The cached directory data is refreshed every 30 seconds and the time zone is GMT-05:00 New York.

```
$ avaki directory --add-schedule --action=tree-pinned
--periodic-schedule="30 seconds" --time-zone="GMT-05:00 New
York" /System/LocalDomain/Services/CacheServices/
MAGADIEUXP.sybase.com "/Shares/Water Buffalo Lodge"
```

avaki directory --cache

Tag an Avaki directory or a directory hierarchy for caching within a cache service.

Syntax

```
avaki directory --cache --action={directory-only |
tree-on-demand | tree-pinned} [--offline-expiration=<value>]
[[--periodic-schedule=<period-spec> |
--cron-schedule=<cron-spec> |
--one-time-schedule=<time-stamp>}]
[--start-time=<time-stamp>] [--end-time=<time-stamp>] |
[--max-iterations=<num-iterations>]]
[--time-zone=<time-zone>] [--exclusion=<exclusion-name>]+
[--exclusion-action={ drop | sooner | later } ]
[--reschedule-offset=<offset-spec>] <cache-service-path>
<directory-path>+
```

<code>--action</code>	Specify how to cache the directory. The formats are as follows:
<code><directory-only></code>	Pin a directory but not its contents for scheduled caching.
<code><tree-on-demand></code>	Tag a directory and its contents for on-demand caching, in which files or file portions can be cached as needed.
<code><tree-pinned></code>	Pin a directory and its contents for scheduled caching.

<pre>--offline- expiration=<value></pre>	<p>Optional. If the back-end source directory is offline, the cache service continues to return the cached data for the directory for the duration (in seconds) specified by the offline expiration value. A value of 0 indicates that the cache service is not allowed to return cached data if the source is offline. A value of -1 indicates that the cached data will never expire when the back-end source is offline. If the offline expiration value is not specified, the cache service default (0) is used.</p> <p>Range: -1 to 2147483647</p>				
<pre>--periodic- schedule= <period-spec></pre>	<p>Specify a periodic schedule for refreshing the cached directory data. Use the following syntax:</p> <pre>--periodic-schedule="<number-of-units> <unit-type>"</pre> <table border="0" style="margin-left: 40px;"> <tr> <td style="padding-right: 20px;"><code><number-of-units></code></td> <td>An integer from 1 to 2147483648</td> </tr> <tr> <td style="padding-right: 20px;"><code><unit-type></code></td> <td>Seconds, minutes, hours, days, weeks, or months</td> </tr> </table>	<code><number-of-units></code>	An integer from 1 to 2147483648	<code><unit-type></code>	Seconds, minutes, hours, days, weeks, or months
<code><number-of-units></code>	An integer from 1 to 2147483648				
<code><unit-type></code>	Seconds, minutes, hours, days, weeks, or months				
<pre>--cron- schedule= <cron-spec></pre>	<p>Create a a cron schedule that tells the cache service how often to refresh cached directory data. Use the following syntax:</p> <pre>--cron-schedule="<seconds> <minutes> <hours> <days-of-month> <months> <days-of-week> [<years>]"</pre> <p>For details about the <code><cron-spec></code> syntax, see Appendix C, “Configuring cron schedules”.</p>				
<pre>--one-time- schedule= <time-stamp></pre>	<p>Specify a single time and date for refreshing the cached directory data. Use the following syntax:</p> <pre>--one-time-schedule="{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}"</pre>				
<pre>--start-time= <time-stamp></pre>	<p>Specify the time and date that the cache service will start refreshing cached directory data, in this format:</p> <pre>{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}</pre>				

<code>--end-time= <time-stamp></code>	Specify the time and date that the cache service will stop refreshing cached directory data, in this format: {<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}																																
<code>--max-iterations= <num-iterations></code>	Specify the number of times a schedule will repeat. <num-iterations> An integer from 1 to 2147483648. If no value is specified, the schedule repeats indefinitely.																																
<code>--time-zone= <time-zone></code>	Specify the time zone for the schedule. The schedule can be specified according to the cache server's time zone or relative to Greenwich Mean Time (GMT). If no value is specified for the time zone, the server's local time zone is used. The format for the time zone is as follows: <table> <tr><td>GMT+12:00 Fiji</td><td>GMT+02:00 Cairo</td></tr> <tr><td>GMT+11:00 Noumea</td><td>GMT+01:00 Paris</td></tr> <tr><td>GMT+10:00 Sydney</td><td>GMT London</td></tr> <tr><td>GMT+9:30 Adelaide</td><td>GMT-01:00 Azores</td></tr> <tr><td>GMT+09:00 Tokyo</td><td>GMT-02:00 Mid-Atlantic</td></tr> <tr><td>GMT+08:00 Hong Kong</td><td>GMT-03:00 Rio de Janeiro</td></tr> <tr><td>GMT+07:00 Bangkok</td><td>GMT-03:30 Newfoundland</td></tr> <tr><td>GMT+06:30 Rangoon</td><td>GMT-04:00 Caracas</td></tr> <tr><td>GMT+06:00 Dacca</td><td>GMT-05:00 New York</td></tr> <tr><td>GMT+05:45 Katmandu</td><td>GMT-06:00 Chicago</td></tr> <tr><td>GMT+05:30 Calcutta</td><td>GMT-07:00 Denver</td></tr> <tr><td>GMT+05:00 Karachi</td><td>GMT-08:00 San Francisco</td></tr> <tr><td>GMT+04:30 Kabul</td><td>GMT-09:00 Juneau</td></tr> <tr><td>GMT+04:00 Baku</td><td>GMT-10:00 Hawaii</td></tr> <tr><td>GMT+03:30 Tehran</td><td>GMT-11:00 Samoa</td></tr> <tr><td>GMT+03:00 Moscow</td><td></td></tr> </table>	GMT+12:00 Fiji	GMT+02:00 Cairo	GMT+11:00 Noumea	GMT+01:00 Paris	GMT+10:00 Sydney	GMT London	GMT+9:30 Adelaide	GMT-01:00 Azores	GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic	GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro	GMT+07:00 Bangkok	GMT-03:30 Newfoundland	GMT+06:30 Rangoon	GMT-04:00 Caracas	GMT+06:00 Dacca	GMT-05:00 New York	GMT+05:45 Katmandu	GMT-06:00 Chicago	GMT+05:30 Calcutta	GMT-07:00 Denver	GMT+05:00 Karachi	GMT-08:00 San Francisco	GMT+04:30 Kabul	GMT-09:00 Juneau	GMT+04:00 Baku	GMT-10:00 Hawaii	GMT+03:30 Tehran	GMT-11:00 Samoa	GMT+03:00 Moscow	
GMT+12:00 Fiji	GMT+02:00 Cairo																																
GMT+11:00 Noumea	GMT+01:00 Paris																																
GMT+10:00 Sydney	GMT London																																
GMT+9:30 Adelaide	GMT-01:00 Azores																																
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic																																
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro																																
GMT+07:00 Bangkok	GMT-03:30 Newfoundland																																
GMT+06:30 Rangoon	GMT-04:00 Caracas																																
GMT+06:00 Dacca	GMT-05:00 New York																																
GMT+05:45 Katmandu	GMT-06:00 Chicago																																
GMT+05:30 Calcutta	GMT-07:00 Denver																																
GMT+05:00 Karachi	GMT-08:00 San Francisco																																
GMT+04:30 Kabul	GMT-09:00 Juneau																																
GMT+04:00 Baku	GMT-10:00 Hawaii																																
GMT+03:30 Tehran	GMT-11:00 Samoa																																
GMT+03:00 Moscow																																	

<code>--exclusion= <exclusion-name></code>	The name of an exclusion to this schedule. For information about creating schedule exclusions, see avaki scheduleexclusion --create --daily (page 201), avaki scheduleexclusion --create --weekly (page 205), avaki scheduleexclusion --create --monthly (page 203), avaki scheduleexclusion --create --yearly (page 207), and avaki scheduleexclusion --create --custom (page 199).						
<code>--exclusion- action={drop sooner later}</code>	Specify the action to take if the cache service is scheduled to refresh cached directory data during an exclusionary period. <table> <tr> <td data-bbox="619 487 672 517">drop</td> <td data-bbox="958 487 1380 552">Don't refresh the cached directory data during the exclusionary period.</td> </tr> <tr> <td data-bbox="619 604 701 633">sooner</td> <td data-bbox="958 604 1380 703">Reschedule the cache service to refresh the cached directory before the exclusionary period.</td> </tr> <tr> <td data-bbox="619 748 672 777">later</td> <td data-bbox="958 748 1380 847">Reschedule the cache service to refresh the cached directory data after the exclusionary period.</td> </tr> </table>	drop	Don't refresh the cached directory data during the exclusionary period.	sooner	Reschedule the cache service to refresh the cached directory before the exclusionary period.	later	Reschedule the cache service to refresh the cached directory data after the exclusionary period.
drop	Don't refresh the cached directory data during the exclusionary period.						
sooner	Reschedule the cache service to refresh the cached directory before the exclusionary period.						
later	Reschedule the cache service to refresh the cached directory data after the exclusionary period.						
<code>--reschedule- offset= <offset-spec></code>	Specify the amount of time before or after the exclusionary period to perform a rescheduled cache refresh. Use the following syntax: <pre>--reschedule-offset="<number-of-units> <unit-type>"</pre> <table> <tr> <td data-bbox="619 1095 886 1124"><number-of-units></td> <td data-bbox="958 1095 1380 1124">An integer from 1 to 2147483648</td> </tr> <tr> <td data-bbox="619 1177 793 1206"><unit-type></td> <td data-bbox="958 1177 1380 1237">Seconds, minutes, hours, days, weeks, or months</td> </tr> </table>	<number-of-units>	An integer from 1 to 2147483648	<unit-type>	Seconds, minutes, hours, days, weeks, or months		
<number-of-units>	An integer from 1 to 2147483648						
<unit-type>	Seconds, minutes, hours, days, weeks, or months						
<code><cache-service- path></code>	The grid path of a cache service.						
<code><directory-path>+</code>	The grid path of one or more directories to cache.						

avaki directory --delete-schedule

Delete a caching schedule for a directory.

To use this command, you must have write permission on the cache service.

Syntax

```
avaki directory --delete-schedule <cache-service-path>  
  <directory-path> <schedule-id>
```

<cache-service-path> The grid path of a cache service.

<directory-path> The grid path of the directory whose caching schedule will be deleted.

<schedule-id> The numeric ID of the schedule to delete. To determine a schedule's ID, use [avaki directory --list-schedules](#) (page 127).

avaki directory --do-not-cache

Mark a cached directory for no caching.

Syntax

```
avaki directory --do-not-cache <cache-service-path>  
  <directory-path>+
```

<cache-service-path> The data catalog path of a cache service.

<directory-path> The data catalog path of the directory to be marked for no caching.

avaki directory --list-schedules

List the schedules for a cached directory.

Syntax

```
avaki directory --list-schedules <cache-service-path>
    <directory-path>
```

<cache-service-path> The data catalog path of a cache service.

<directory-path> The data catalog path of the directory whose schedule will be listed.

avaki domain --create

Create an Avaki domain.

Syntax

```
avaki domain --create <domain> [<host>] [<gdc-port>] | <host>:
    <gdc-port>]
```

<domain> Specify the name of the Avaki domain to create. The name can be up to 30 characters long.

<host> Optional. The DNS name or IP address of the machine on which you want to create the new domain's GDC.

<gdc-port> Optional. The connect port for the new grid domain controller. (The GDC uses the connect port to communicate with other Avaki objects; see the [connect port](#) entry in the Glossary, [page 310](#), for more information.) If you omit this option, the system assumes that the connect port is 3099.

avaki domain --disconnect

Disconnect an Avaki domain from another domain.

Syntax

```
avaki domain --disconnect <other-domain>
```

<other-domain> The name of the Avaki domain to disconnect from.

avaki domain --info

View a description of an Avaki domain, including lists of all servers in the domain and any connected domains.

Syntax

```
avaki domain --info
```

– This command has no options.

avaki domain --interconnect

Connect an Avaki domain to another Avaki domain.

Syntax

```
avaki domain --interconnect <other-domain> <machine>:<port>
```

<other-domain> The name of the Avaki domain to connect with.

<machine>:<port> For <machine>, specify the DNS name or IP address of the GDC machine in the other Avaki domain. For port, specify the remote GDC machine's connect port number.

avaki executionservice --info

Display information about an execution service, including its name, the maximum number of concurrent data services, and any other grid servers whose execution services are pooled with this execution service.

Syntax

```
avaki executionservice --info <server-path>
```

<server-path> The data catalog path to the Avaki server on which the execution service resides.

avaki executionservice --set

Configure an execution service.

Syntax

```
avaki executionservice --set  
  [--max-concurrent=<maximum-data-services>]  
  --pool=<server-path>* <server-path>
```

<code>--max-concurrent= <maximum-data- services></code>	Specify an interval that represents the maximum number of concurrent data services.
<code>--pool= <server-path>*</code>	The path to one or more grid servers that the execution service is associated with. To specify more than one server, include the <code>--pool</code> argument for each server. The format is as follows: <code>--pool=<serverpath1> --pool=<serverpath2></code>
<code><server-path></code>	The data catalog path to the grid server on which the execution service resides.

avaki file --add-schedule

Add a new pin or cache schedule to a currently cached file. (For information about tagging an Avaki directory or a directory hierarchy for caching within a cache service, see [avaki file --cache-on-demand](#) (page 135) and [avaki file --pin](#) (page 138)).

To use this command, you must have write permission on the cache service, and you must have read permission on the file.

Syntax

```
avaki file --add-schedule {--periodic-schedule=<period-spec> |
  --cron-schedule=<cron-spec> |
  --one-time-schedule=<time-stamp>}
  [--start-time=<time-stamp>] [--end-time=<time-stamp>] |
  [--max-iterations=<num-iterations>]
  [--time-zone=<time-zone>] [--exclusion=<exclusion-name>]+
  [--exclusion-action={ drop | sooner | later }]
  [--reschedule-offset=<offset-spec>] <cache-service-path>
  <file-path>+
```

--periodic-
schedule=
<period-spec>

Specify a periodic schedule for refreshing the cached file data.
Use the following syntax:

```
--periodic-schedule="<number-of-  
units> <unit-type>"
```

<number-of-units> An integer from 1 to 2147483648

<unit-type> seconds, minutes, hours, days,
 weeks, or months

In this example, the periodic schedule is set to 30 seconds:

```
--periodic-schedule="30 seconds"
```

<code>--cron-schedule= <cron-spec></code>	<p>Specify a periodic schedule for refreshing the cached file data. Use the following syntax:</p> <pre>--cron-schedule="<code><seconds> <minutes> <hours> <days-of-month> <months> <days-of-week> [<years>]</code>"</pre> <p>For details about the <code><cron-spec></code> syntax, see Appendix C, “Configuring cron schedules”.</p>						
<code>--one-time-schedule= <time-stamp></code>	<p>Specify a single time and date for refreshing the cached file data. Use the following syntax:</p> <pre>--one-time-schedule="{<code><hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now</code>}"</pre>						
<code>--start-time= <time-stamp></code>	<p>Specify the time and date that the cache service will start refreshing cached file data, in this format:</p> <pre>{<code><hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now</code>}</pre>						
<code>--end-time= <time-stamp></code>	<p>Specify the time and date that the cache service will stop refreshing cached file data, in this format:</p> <pre>{<code><hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now</code>}</pre>						
<code>--max-iterations= <num-iterations></code>	<p>Specify the number of times a schedule will repeat.</p> <pre><num-iterations></pre> <p>An integer from 1 to 2147483648. If no value is specified, the schedule repeats indefinitely.</p>						
<code>--time-zone= <time-zone></code>	<p>Specify the time zone for the schedule. The schedule can be specified according to the cache server’s time zone or relative to Greenwich Mean Time (GMT). If no value is specified for the time zone, the server’s local time zone is used. The format for the time zone is as follows:</p> <table border="0" style="margin-left: 40px;"> <tr> <td>GMT+12:00 Fiji</td> <td>GMT+02:00 Cairo</td> </tr> <tr> <td>GMT+11:00 Noumea</td> <td>GMT+01:00 Paris</td> </tr> <tr> <td>GMT+10:00 Sydney</td> <td>GMT London</td> </tr> </table>	GMT+12:00 Fiji	GMT+02:00 Cairo	GMT+11:00 Noumea	GMT+01:00 Paris	GMT+10:00 Sydney	GMT London
GMT+12:00 Fiji	GMT+02:00 Cairo						
GMT+11:00 Noumea	GMT+01:00 Paris						
GMT+10:00 Sydney	GMT London						

GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii
GMT+03:30 Tehran	GMT-11:00 Samoa
GMT+03:00 Moscow	

`--exclusion=
<exclusion-name>`

The name of an exclusion to this schedule. For information about creating schedule exclusions, see [avaki scheduleexclusion --create --daily](#) (page 201), [avaki scheduleexclusion --create --weekly](#) (page 205), [avaki scheduleexclusion --create --monthly](#) (page 203), [avaki scheduleexclusion --create --yearly](#) (page 207), and [avaki scheduleexclusion --create --custom](#) (page 199).

`--exclusion-
action={drop |
sooner | later}`

Specify the action to take if the cache service is scheduled to refresh cached file data during an exclusionary period.

drop	Don't refresh the cached file data during the exclusionary period.
sooner	Reschedule the cache service to refresh the cached file data before the exclusionary period.
later	Reschedule the cache service to refresh the cached file data after the exclusionary period.

<code>--reschedule-offset=<offset-spec></code>	Specify the amount of time before or after the exclusionary period to perform a rescheduled cache refresh. Use the following syntax: <code>--reschedule-offset="<number-of-units><unit-type>"</code>
<code><number-of-units></code>	An integer from 1 to 2147483648
<code><unit-type></code>	Seconds, minutes, hours, days, weeks, or months
<code><cache-service-path></code>	The data catalog path of a cache service.
<code><file-path>+</code>	The data catalog path of the file to which a pin schedule will be added.

Example

This example pins the file `fredsform.txt` for scheduled caching. The cached directory data is refreshed every 30 seconds and the time zone is GMT-05:00 New York.

```
$ avaki file --add-schedule --periodic-schedule="30 seconds"
--time-zone="GMT-05:00 New York" /System/LocalDomain/
Services/CacheServices/MAGADIEUXP.sybase.com "/Shares/Water
Buffalo Lodge/fredsform.txt"
```

avaki file --cache-on-demand

Tags a file within a cache service for on-demand caching.

Syntax

```
avaki file --cache-on-demand [--coherence-window=<value>]
  [--offline-expiration=value] <cache-service-path>
  <file-path>+
```

--coherence-
window=<value>

Optional. The duration (in seconds) during which the cache service assumes the cached file content to be fresh after it last inspected the back-end source file for updates. The valid values are -1 and 0 to 2147483647 seconds. If the coherence window is not specified, the cache service default, 86400 (1 day), is used. Set this attribute to -1 when the cache includes static files that should never expire, or set it to 0 to indicate that they should never be cached.

--offline-
expiration=<value>

Optional. If the back-end file is offline, the cache service continues to return the cached data for the file for the duration (in seconds) specified by the offline-expiration value. A value of 0 indicates that the cache service is not allowed to return cached data after it has determined that the source is offline. A value of -1 indicates that the cached data never expires when the back-end source is offline. If the offline-expiration value is not specified, the cache service default (0) is used.
Range: -1 to 2147483647

<cache-service-
path>

The data catalog path of a cache service.

<file-path>+

The data catalog path of one or more files to tag for on-demand caching.

avaki file --delete-schedule

Delete a pin schedule from a cached file.

To use this command, you must have write permission on the cache service.

Syntax

```
avaki file --delete-schedule <cache-service-path> <file-path>
    <schedule-id>
```

<cache-service-path> The data catalog path of a cache service.

<file-path> The data catalog path of the file whose pin schedule will be deleted.

<schedule-id> The numeric ID of the schedule to delete. To determine a schedule's ID, use [avaki file --list-schedules](#) (page 137).

avaki file --do-not-cache

Mark a file for no caching.

Syntax

```
avaki file --do-not-cache <cache-service-path> <file-path>+
```

<cache-service-path> The data catalog path of a cache service.

<file-path>+ The data catalog path of one or more files to be marked for no caching.

avaki file --list-schedules

List the schedules for a cached file.

Syntax

```
avaki file --list-schedules <cache-service-path> <file-path>
```

<cache-service-path> The data catalog path of a cache service.

<file-path> The data catalog path of the file whose schedule will be listed.

avaki file --pin

Marks (“pins”) a file within a cache service for scheduled caching.

Syntax

```
avaki file --pin [--offline-expiration=<value>]
  [{--periodic-schedule=<period-spec> |
  --cron-schedule=<cron-spec> |
  --one-time-schedule=<time-stamp>}]
  [--start-time=<time-stamp>] [--end-time=<time-stamp>] |
  [--max-iterations=<num-iterations>]
  [--time-zone=<time-zone>] [--exclusion=<exclusion-name>] +
  [--exclusion-action={ drop | sooner | later }]
  [--reschedule-offset=<offset-spec>] <cache-service-path>
  <file-path>+
```

--offline-expiration=<value> Optional. If the back-end source file is offline, the cache service continues to return the cached data for the file for the duration (in seconds) specified by the offline-expiration value. A value of 0 indicates that the cache service is not allowed to return cached data after it has determined that the source is offline. A value of -1 indicates that the cached data never expires when the back-end source is offline. If the offline-expiration value is not specified, the cache service default (0) is used.
Range: -1 to 2147483647

--periodic-schedule=<period-spec> Specify a periodic schedule for refreshing the cached file data. Use the following syntax:

```
--periodic-schedule="<number-of-units>
<unit-type>"
```

<number-of-units> An integer from 1 to 2147483648

<unit-type> Seconds, minutes, hours, days, weeks, or months

<pre>--cron- schedule= <cron-spec></pre>	<p>Specify a periodic schedule for refreshing the cached file data. Use the following syntax:</p> <pre>--cron-schedule="<seconds> <minutes> <hours> <days-of-month> <months> <days-of-week> [<years>]"</pre> <p>For details about the <cron-spec> syntax, see Appendix C, “Configuring cron schedules”.</p>								
<pre>--one-time- schedule= <time-stamp></pre>	<p>Specify a single time and date for refreshing the cached file data. Use the following syntax:</p> <pre>--one-time-schedule="{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}"</pre>								
<pre>--start-time= <time-stamp></pre>	<p>Specify the time and date that the cache service will start refreshing cached file data, in this format:</p> <pre>{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}</pre>								
<pre>--end-time= <time-stamp></pre>	<p>Specify the time and date that the cache service will stop refreshing cached file data, in this format:</p> <pre>{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}</pre>								
<pre>--max-iterations= <num-iterations></pre>	<p>Specify the number of times a schedule will repeat.</p> <pre><num-iterations></pre> <p>An integer from 1 to 2147483648. If no value is specified, the schedule repeats indefinitely.</p>								
<pre>--time-zone= <time-zone></pre>	<p>Specify the time zone for the schedule. The schedule can be specified according to the cache server’s time zone or relative to Greenwich Mean Time (GMT). If no value is specified for the time zone, the server’s local time zone is used. The format for the time zone is as follows:</p> <table> <tr> <td data-bbox="619 1376 808 1402">GMT+12:00 Fiji</td> <td data-bbox="953 1376 1165 1402">GMT+02:00 Cairo</td> </tr> <tr> <td data-bbox="619 1433 865 1459">GMT+11:00 Noumea</td> <td data-bbox="953 1433 1158 1459">GMT+01:00 Paris</td> </tr> <tr> <td data-bbox="619 1491 855 1517">GMT+10:00 Sydney</td> <td data-bbox="953 1491 1108 1517">GMT London</td> </tr> <tr> <td data-bbox="619 1548 858 1574">GMT+9:30 Adelaide</td> <td data-bbox="953 1548 1172 1574">GMT-01:00 Azores</td> </tr> </table>	GMT+12:00 Fiji	GMT+02:00 Cairo	GMT+11:00 Noumea	GMT+01:00 Paris	GMT+10:00 Sydney	GMT London	GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+12:00 Fiji	GMT+02:00 Cairo								
GMT+11:00 Noumea	GMT+01:00 Paris								
GMT+10:00 Sydney	GMT London								
GMT+9:30 Adelaide	GMT-01:00 Azores								

GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii
GMT+03:30 Tehran	GMT-11:00 Samoa
GMT+03:00 Moscow	

```
--exclusion=  
<exclusion-name>
```

The name of an exclusion to this schedule. For information about creating schedule exclusions, see [avaki scheduleexclusion --create --daily](#) (page 201), [avaki scheduleexclusion --create --weekly](#) (page 205), [avaki scheduleexclusion --create --monthly](#) (page 203), [avaki scheduleexclusion --create --yearly](#) (page 207), and [avaki scheduleexclusion --create --custom](#) (page 199).

```
--exclusion-  
action={drop |  
sooner | later}
```

Specify the action to take if the cache service is scheduled to refresh cached file data during an exclusionary period.

drop	Don't refresh the cached file data during the exclusionary period.
sooner	Reschedule the cache service to refresh the cached file data before the exclusionary period.
later	Reschedule the cache service to refresh the cached file data after the exclusionary period.

<code>--reschedule- offset= <offset-spec></code>	Specify the amount of time before or after the exclusionary period to perform a rescheduled cache refresh. Use the following syntax: <code>--reschedule-offset=" <number-of-units> <unit-type>"</code> <code><number-of-units></code> An integer from 1 to 2147483648 <code><unit-type></code> Seconds, minutes, hours, days, weeks, or months
<code><cache-service- path></code>	The data catalog path of a cache service.
<code><file-path>+</code>	The grid paths of one or more files to tag for scheduled caching.

avaki group --add --user

Add a user to a group.

Syntax

```
avaki group {--add | -a} {--user | -u} {<qualified-user-name>}
           {<group-grid-path> | <qualified-group-name>}
```

`<qualified-user-name>` Specify the qualified name of the user that you are adding to the group. The format is:

```
<user>@<authservice>.<authservicetype>.<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of `<qualified-user-name>` are as follows:

`<user>` The name of the user to add to the group.

`<authservice>` The name of the user's authentication service.

	<code><authservicetype></code>	The type of authentication service (Grid, Nis, or Ldap).
	<code><domain></code>	The name of the user's Avaki domain.
<code><group-grid-path> <qualified-group-name></code>		Specify the group to which you are adding the user. You can specify the data catalog path to the group directory. For example: <code>/System/Domains/Bedrock/Services/AuthServices/Grid/DefaultAuthService/Groups/Development</code> Alternatively, you can specify the qualified group name using the following format: <code><group>@<authservice>.<authservicetype>.<domain></code> For example: <code>Development@DefaultAthService.Grid.Bedrock</code> The elements of <code><qualified-group-name></code> are as follows:
	<code><group></code>	The name of the group to which you are adding the user.
	<code><authservice></code>	The name of the group's authentication service.
	<code><authservicetype></code>	The type of the group's authentication service (Grid, Nis, or Ldap).
	<code><domain></code>	The name of the group's Avaki domain.

avaki group --create

Create a new grid group.

Syntax

```
avaki group --create [--description=<description-text>]  
  <grid-group>+
```

--description=
<description-
text>

Optional. Enter details about this group; enclose the entire text string in quotation marks. The description is shown to web UI users who display information about the group. Consider including an explanation of what this group is for or where to get more information.

Note: If you create or import multiple groups with one command, the same description is applied to every group.

Length limit: 255 characters

<grid-group>+

The name(s) of the group or groups to create.

avaki group --delete

Delete a group.

Note To list the groups in a authentication service, enter a command of the form `avaki ls <auth-service-directory>/Groups`. For example, `avaki ls /System/Domains/Bedrock/Services/AuthServices/Ldap/testldap/Groups`.

Syntax

```
avaki group {--delete | -d} <auth-service-path> {--all | -a} |
  <group-name>+}
```

`<auth-service-path>` The data catalog path to the authentication service from which you are deleting the group. For example:
 /System/LocalDomain/Services/AuthServices/Grid/Default-AuthService/AuthService

`--all | -a` Delete all groups from the specified authentication service.

`<group-name>` The name of the group you are deleting. To specify more than one group, enter a space-delimited list of group names. For example:

```
Development QA
```

avaki group --delete --user

Delete a user from a group.

Syntax

```
avaki group {--delete | -d} {--user | -u}
  {<qualified-user-name>} {<group-grid-path> |
  <qualified-group-name>}
```

<qualified-
user-name>

Specify the qualified name of the user that you are deleting from the group. The format is:

```
<user>@<authservice>.<authservice-
type>.<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user to delete from the group.
<authser- vice>	The name of the user's authentication service.
<authser- vicetype>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

<group-grid-path> |
<qualified-
group-name>

Specify the qualified name of the group from which you are deleting the user. The format is:

```
<group>@<authservice>.<authservicetype>.  
<domain>
```

For example:

```
Development@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-group-name> are as follows:

<code><group></code>	The path or name of the group from which you are removing the user.
<code><authservice></code>	The name of the group's authentication service.
<code><authservicetype></code>	The type of the group's authentication service (Grid, Nis, or Ldap).
<code><domain></code>	The name of the group's Avaki domain.

avaki group --info

Display a group's description.

Syntax

```
avaki group --info {<group-grid-path> |
  <qualified-group-name>}}
```

<group-grid-path> | The name of the group that you are obtaining information
<qualified- about.
group-name>

You can specify the data catalog path to the directory of the group. For example:

```
/System/Domains/Bedrock/Services/AuthServices/Grid/DefaultAuthService/Groups/Development
```

Alternatively, you can specify the qualified group name using the following format:

```
<group>@<authservice>.<authservice-type>.<domain>
```

For example:

```
Development@DefaultAthService.Grid.Bedrock
```

The elements of <qualified-group-name> are as follows:

<group>	The name of the group for which you are obtaining information.
<authservice>	The name of the group's authentication service.
<authservice-type>	The type of the group's authentication service (Grid, Nis, or Ldap).
<domain>	The name of the group's Avaki domain.

avaki group --list-user

List all users in a group.

Syntax

```
avaki group {--list-user | -l} {<group-grid-path> |
  <qualified-group-name>}
```

<group-
grid-path> |
<qualified-
group-name>

The grid path of the group for which you want to obtain a list of users. You can specify the path to the grid directory of the group. For example:

```
/System/LocalDomain/Services/AuthServices/Grid/DefaultAuthService/Groups/Development
```

Alternatively, you can specify the qualified group name by using the following format:

```
<group>@<authservice>.<authservicetype>.  
<domain>
```

For example:

```
Development@DefaultAthService.Grid.Bedrock
```

The elements of <qualified-group-name> are as follows:

<group>	The name of the group for which you want to obtain a list of users.
<authservice>	The name of the group's authentication service.
<authservice- type>	The type of the group's authentication service (Grid, Nis, or Ldap).
<domain>	The name of the group's Avaki domain.

avaki help

View a list of Avaki commands and the Avaki release version number, or view usage information for a particular command.

Syntax

```
avaki help [<cmd>] [--verbose | -v]
```

<code><cmd></code>	Optional. The command for which you want to view usage information. If no command is specified, the system displays a list of all Avaki commands.
<code>--verbose -v</code>	Optional. Display the commands' syntax and descriptions.

avaki id

Download an interconnection ID for a user or group.

Syntax

```
avaki id --user={<full-user-path> | <full-group-path>}
        [--filename=<output-filename>]
```

<code>--user=</code> {<full-user-path> <full-group-path>}	The grid directory of the user or group for which you want to obtain an interconnection ID.
<code>--filename=</code> <output-filename>	Optional. Specify a filename and path in the local file system for the interconnection ID that this command generates.

avaki ldap --add-schedule

Add a schedule according to which users will be imported automatically. When users are imported according to a schedule, the list of currently imported users is refreshed and any new users are imported automatically.

To use this command, you need write permission on the LDAP authentication service.

Note If you want to import different groups on different schedules, configure a cron specification for the `avaki ldap --add-schedule` command and use it in conjunction with [avaki ldap --import](#) (page 155).

Syntax

```
avaki ldap --add-schedule <auth-service>
  {--periodic-schedule=<period-spec> |
  --cron-schedule=<cron-spec> |
  --one-time-schedule=<time-stamp>}
  [--start-time=<time-stamp>] [--end-time=<time-stamp>] |
  [--max-iterations=<num-iterations>]]
  [--time-zone=<time-zone>] [--exclusion=<exclusion-name>]+
  [--exclusion-action={ drop | sooner | later }]
  [--reschedule-offset=<offset-spec>]
```

<code><auth-service></code>	The name of the authentication service for which the schedule will be added.
<code>--periodic-schedule=<period-spec></code>	Specify a periodic schedule for importing users automatically. Use the following syntax: <code>--periodic-schedule="<number-of-units> <unit-type>"</code>
<code><number-of-units></code>	An integer from 1 to 2147483648
<code><unit-type></code>	seconds, minutes, hours, days, weeks, or months

In the following example, the periodic schedule is set to 30 seconds:

```
--periodic-schedule="30 seconds"
```

```
--cron-  
schedule=  
<cron-spec>
```

Specify how often a cron-based scheduler for importing users will import users. Use the following syntax:

```
--cron-schedule="<seconds> <minutes> <hours>  
<days-of-month> <months> <days-of-week>  
[<years>]"
```

For details about the <cron-spec> syntax, see [Appendix C, “Configuring cron schedules”](#).

```
--one-time-  
schedule=  
<time-stamp>
```

Specify a single time and date for importing the users. Use the following syntax:

```
--one-time-schedule="{<hh>:<mm>:<ss> {am | pm}  
<mm>-<dd>-<yyyy> | now}"
```

```
--start-time=  
<time-stamp>
```

Specify the time and date that the users will be imported, in this format:

```
{<hh>:<mm>:<ss> {am | pm} <mm>-<dd>-<yyyy> |  
now}
```

```
--end-time=  
<time-stamp>
```

Specify the time and date that the users will be imported, in this format:

```
{<hh>:<mm>:<ss> {am | pm} <mm>-<dd>-<yyyy> |  
now}
```

```
--max-iterations=  
<num-iterations>
```

Specify the number of times a schedule will repeat.

```
<num-iterations>
```

An integer from 1 to 2147483648. If no value is specified, the schedule repeats indefinitely.

`--time-zone=
<time-zone>`

Specify the time zone for the schedule. The schedule can be specified according to the grid server's time zone or relative to Greenwich Mean Time (GMT). If no value is specified for the time zone, the server's local time zone is used. The format for the time zone is as follows:

GMT+12:00 Fiji	GMT+02:00 Cairo
GMT+11:00 Noumea	GMT+01:00 Paris
GMT+10:00 Sydney	GMT London
GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii
GMT+03:30 Tehran	GMT-11:00 Samoa
GMT+03:00 Moscow	

`--exclusion=
<exclusion-name>`

The name of an exclusion to this schedule. For information about creating schedule exclusions, see [avaki scheduleexclusion](#) **--create --daily** (page 201), [avaki scheduleexclusion --create --weekly](#) (page 205), [avaki scheduleexclusion --create --monthly](#) (page 203), [avaki scheduleexclusion --create --yearly](#) (page 207), and [avaki scheduleexclusion --create --custom](#) (page 199).

<code>--exclusion-action={drop sooner later}</code>	Specify the action to take if the authentication service is scheduled to update users automatically during an exclusionary period.
<code>drop</code>	Don't run the import task during the exclusionary period.
<code>sooner</code>	Reschedule the import task to run before the exclusionary period.
<code>later</code>	Reschedule the import task to run after the exclusionary period.
<code>--reschedule-offset=<offset-spec></code>	Specify the amount of time before or after the exclusionary period to perform a user import. Use the following syntax:
	<code>--reschedule-offset="<number-of-units><unit-type>"</code>
<code><number-of-units></code>	An integer from 1 to 2147483648
<code><unit-type></code>	Seconds, minutes, hours, days, weeks, or months

Example

This example add a schedule in which users will be imported automatically into the MyLDAP authentication service every 2 hours. In this example, the time zone is GMT-05:00 New York.

```
$ avaki ldap --add-schedule MyLDAP --periodic-schedule="2 hours" --time-zone="GMT-05:00 New York"
```

avaki ldap --delete

Delete an LDAP authentication service or delete users from an authentication service.

Syntax

To delete an authentication service:

```
avaki ldap {--delete | -d} <auth-service>
```

To delete users:

```
avaki ldap {--delete | -d} {--user | -u}  
  {<user-name>+ | {--all | -a}} <auth-service>
```

- | | |
|-----------------------------------|--|
| <code><auth-service></code> | The name of the authentication service to delete. |
| <code>--user -u</code> | Delete one or more users from the specified authentication service. |
| <code><user-name>+</code> | The name or names of user accounts to import. To specify more than one user, enter a space-delimited list of user IDs. |
| <code>--all -a</code> | Import all users or groups from the specified authentication service. |

avaki ldap --delete-schedule

Delete a schedule that specifies when users will be imported automatically.

To use this command, you must have write permission on the LDAP authentication service.

Syntax

```
avaki ldap --delete-schedule <auth-service> <schedule-id>
```

<auth-service> The name of the authentication service for which the schedule will be deleted.

<schedule-id> The numeric ID of the schedule to delete. To determine a schedule's ID, use [avaki ldap --list-schedules](#) (page 160).

avaki ldap --import

Import user accounts or groups from an LDAP authentication service into an Avaki grid domain.

Note When you import a user, you import all the groups to which the user belongs. When you import a group, you import all users in the group and all groups to which those users belong. Exceptions are related to the Domain Users group in Active Directory: importing an AD Domain Users group has no effect. And when you import a member of an AD Domain Users group, the Domain Users group is not imported into the grid domain. (Other groups that the user belongs to will be imported as usual.)

The import uses the searchbase you entered when you integrated the LDAP directory service.

Syntax

To import user accounts:

```
avaki ldap {--import | -i} {--user | -u} <user-name>+
  <auth-service>
```

```
avaki ldap {--import | -i} {--user | -u} --dn <user-name>+
  <auth-service>
```

```
avaki ldap {--import | -i} {--user | -u} {--all | -a}
  <auth-service>
```

To import groups:

```
avaki ldap {--import | -i} {--group | -g} {<group-name>+ |
  [--all | -a]} <auth-service>
```

<code>--user -u</code>	Import users.
<code><user-name>+</code>	The name or names of user accounts to import. To specify more than one user, enter a space-delimited list of user IDs. If you use the <code>--dn</code> option, each name must be an LDAP distinguished name (for example, <code>CN=Joe Smith, OU=Burlington, DC=Avaki, DC=com</code>). Enclose each DN string in quotation marks and separate the strings with spaces.
<code>--dn</code>	Indicates that you are providing LDAP distinguished names in the <code><users></code> argument.
<code>--all -a</code>	Import all users or groups from the specified authentication service.
<code>--group -g</code>	Import groups.
<code><group-name>+</code>	The name or names of the groups to import. To specify more than one group, enter a space-delimited list of group names.
<code><auth-service></code>	The name of the authentication service to import.

Examples

This command uses LDAP distinguished names to import user accounts for Fred Flintstone and Barney Rubble from the testldap authentication service.

```
$ avaki ldap --import -u --dn "CN=Fred Flintstone,  
OU=Bedrock, DC=Slate Quarry, DC=com" "CN=Barney Rubble,  
OU=Bedrock, DC=Slate Quarry, DC=com" testldap
```

This command imports all groups from the testldap authentication service.

```
$ avaki ldap --import -g -a testldap
```

avaki ldap --info

Display information about an LDAP authentication service.

Syntax

```
avaki ldap --info <auth-service>
```

<auth-service> The name of the authentication service from which to add or delete the search base.

avaki ldap --integrate

Integrate an Active Directory or other LDAP-compliant authentication service into an Avaki domain.

Syntax

```
avaki ldap --integrate <auth-service>
  {--host | -h}=<ldap-host>[:<port-number>]
  --auth-type=<type> --auth-dn=<search-base>
  --auth-pw=<password> --user-searchbase=<search-base>
  --group-searchbase=<search-base>
  --user-account-attr=<user-name-attribute-name>
  --group-account-attr=<group-name-attribute-name>
  [--uid-attr=<uid-attribute-name>]
  [--refresh-on-login] [--enabled-query=<query-string>]
```

<auth-service>	Devise a name for the new service.
{--host -h}= <ldap-host>[:<port- number>]	The name or IP address of the host machine on which the LDAP authentication service resides. Optionally, you can include the port number on the LDAP host. By default, Avaki uses the default LDAP host port, 389. If your LDAP host uses another port, use --host or -h to specify it.
--auth-type= <type>	The authentication server's type. For <type>, specify simple.
--auth-dn= <search-base>	Enter one or more LDAP name components in the Search Base field. The string you enter here must include values for CN, OU, DC—for example, CN=Joe Smith, OU=Burlington, DC=Avaki, DC=com. If you do not know what to enter here, consult your LDAP administrator.
--auth-pw= <password>	The password for the authentication server. If you're executing this command on a Unix system and the password contains non-alphanumeric characters (! or &, for instance), enclose the password in single quotation marks to prevent the shell from interpreting the special characters.

<code>--user-searchbase= <search-base></code>	The top level of the first LDAP directory context within which the system will search for users. For example, you might enter OU=Burlington, DC=Avaki, DC=com. To specify more than one searchbase, use commas to separate the searchbase names. The system searches on the LDAP name components to find users that you can import. If you do not know what to enter here, consult your LDAP administrator.
<code>--group-searchbase= <search-base></code>	The top level of the first LDAP directory context within which the system will search for groups. For example, you might enter OU=Groups, DC=Avaki, DC=com. To specify more than one searchbase, use commas to separate the searchbase names. The system searches on the LDAP name components to find groups that you can import. If you do not know what to enter here, consult your LDAP administrator.
<code>--user-account-attr= <user-name-attribute-name></code>	Enter the name of the LDAP attribute that contains the user name that a user from this authentication service will use to log in to this Avaki domain. For example, enter “sAMAccountName” for an out-of-the-box Active Directory installation, or “uid” for an out-of-the-box iPlanet installation.
<code>--group-account-attr= <group-name-attribute-name></code>	Enter the name of the LDAP attribute that contains the group name that users from this authentication service will use to log in to this Avaki domain. For example, enter “cn” for out-of-the-box Active Directory and iPlanet installations.
<code>--uid-attr= <uid-attribute-name></code>	Optional. Enter the name of an LDAP attribute that contains UID information.
<code>[--refresh-on-login]</code>	Optional. Refresh a user’s information (such as group memberships) from the external directory each time the user logs in.
<code>[--enabled-query= <query-string></code>	Optional. Specify a query that will be sent to the LDAP server to determine whether a user is enabled. For information about the format for LDAP queries, see your LDAP documentation.

avaki ldap --list-schedules

Display a list of schedules that specify when users will be imported automatically.

Syntax

```
avaki ldap --list-schedules <auth-service>
```

<auth-service> The name of the authentication service for which the schedule will be displayed.

avaki ldap --searchbase

Add or delete a search base for an Active Directory or other LDAP-compliant authentication service that has been integrated into an Avaki domain.

Syntax

```
avaki ldap --searchbase {--add | --delete} {--user | -u} |
  {--group | -g} <auth-service> <search-base>
```

--add | --delete Specify --add to add a search base or --delete to delete one.

--user | -u Specify --user or -u to add or delete a user search base.

--group | -g Specify --group or -g to add or delete a group search base.

<auth-service> The name of the authentication service from which to add or delete the search base.

<search-base> The LDAP directory context within which to add or delete the search base. For example, you might enter OU=Burlington, DC=Avaki, DC=com. If you do not know what to enter here, consult your LDAP administrator.

avaki ldap --update

Update an Active Directory or other LDAP-compliant authentication service that has been integrated into an Avaki domain.

Syntax

```
ldap --update <auth-service> {--host | -h}=<ldap-host>
--auth-type=<type> --auth-dn=<search-base>
--auth-pw=<password> --user-searchbase=<search-base>
--group-searchbase=<search-base>
--user-account-attr=<user-name-attribute-name>
--group-account-attr=<group-name-attribute-name>
[--uid-attr=<uid-attribute-name>]
[--enabled-query=<query-string>]
```

<code><auth-service></code>	The name of the authentication service.
<code>{--host -h}=<ldap-host></code>	The name or IP address of the host machine on which the LDAP authentication service resides.
<code>--auth-type=<type></code>	The authentication server's type. For <code><type></code> , specify simple.
<code>--auth-dn=<search-base></code>	Enter one or more LDAP name components in the Search Base field. The string you enter here must include values for CN, OU, DC—for example, CN=Joe Smith, OU=Burlington, DC=Avaki, DC=com. If you do not know what to enter here, consult your LDAP administrator.
<code>--auth-pw=<password></code>	The password for the authentication server. If you're executing this command on a Unix system and the password contains non-alphanumeric characters (! or &, for instance), enclose the password in single quotation marks to prevent the shell from interpreting the special characters.

<code>--user-searchbase= <search-base></code>	The top level of the first LDAP directory context within which the system will search for users. For example, you might enter OU=Burlington, DC=Avaki, DC=com. To specify more than one searchbase, use commas to separate the searchbase names. The system searches on the LDAP name components to find users that you can import. If you do not know what to enter here, consult your LDAP administrator.
<code>--group- searchbase= <search-base></code>	The top level of the first LDAP directory context within which the system will search for groups. For example, you might enter OU=Groups, DC=Avaki, DC=com. To specify more than one searchbase, use commas to separate the searchbase names. The system searches on the LDAP name components to find groups that you can import. If you do not know what to enter here, consult your LDAP administrator.
<code>--user-account- attr=<user-name- attribute-name></code>	Enter the name of the LDAP attribute that contains the user name that a user from this authentication service will use to log in to this Avaki grid domain. For example, enter “sAMAccountName” for an out-of-the-box Active Directory installation, or “uid” for an out-of-the-box iPlanet installation.
<code>--group-account- attr=<group-name- attribute-name></code>	Enter the name of the LDAP attribute that contains the group name that users from this authentication service will use to log in to this Avaki grid domain. For example, enter “cn” for out-of-the-box Active Directory and iPlanet installations.
<code>--uid-attr= <uid-attribute-name></code>	Optional. Enter the name of an LDAP attribute that contains UID information.
<code>[--enabled-query= <query-string></code>	Optional. Specify a query that will be sent to the LDAP server to determine whether a user is enabled. For information about the format for LDAP queries, see your LDAP documentation.

avaki ln

Create a link between a grid file or directory and another location in the data catalog. Similar to the Unix **ln** command.

You can create *hard links* or *soft links*. A hard link (the default type) is an alternate name for a file or directory; changes to the object's other names have no effect on the hard link. A soft link is a pointer to a particular location (name) in the data catalog. If the object at that location is moved, deleted, or renamed, the soft link breaks. If you need to create a link to an Avaki generated view, use a soft link.

Syntax

```
avaki ln [-s] <target-grid-path> {<link-name> |
  <directory-target>}
```

-s Soft link (optional).

<target-grid-path> The path to the grid file or directory for which you are creating a link.

<link-name> If you want to give a link a different name than the object to which it refers, specify a link name or the directory in which the link will be placed. The path may be relative or absolute.

<directory-target> The directory in which the link will be placed. The link will have the same name as the object to which the link refers.

avaki locks --clear

Clear any file locks that are held on a file.

Syntax

```
avaki locks {--clear | -c} <grid-file>
```

<grid-file> The grid path of the locked file.

avaki locks --list

Lists the file locks that are currently held on a file.

Syntax

```
avaki locks {--list | -l} <grid-file>
```

<grid-file> The grid path of the file for which you want lock information.

avaki login

Log in to an Avaki domain.

Syntax

```
avaki login {--user=<user> | <user>} {--password=<password> |
  <password>} [--auth-service=<auth-service>]
  [--type=<service-type>] [--domain=<domain>]
```

--user=<user> |
<user>

Your Avaki user name. You can specify your user name by itself (for example, wilma), or you can specify your qualified user name. The format for the qualified user name is as follows:

```
<user>@<authservice>.<authservice-
type>.<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

If you specify a qualified user name, do not specify <auth-service>, <service-type>, or <domain> separately.

The elements of <qualified-user-name> are as follows:

<user> Your user name.

<authservice> Your authentication service.

<authservice-
type> The type of authentication service (Grid,
Nis, or Ldap).

<domain> Your Avaki domain name.

--password=
<password> |
<password>

Your Avaki password. If you're executing this command on a Unix system and your password contains nonalphanumeric characters (! or &, for instance), enclose your password in single quotation marks to prevent the shell from interpreting the special characters.

<code>--auth-service</code> <code>=<auth-service></code>	Optional. The name of your Avaki authentication service.
<code>--type=</code> <code><service-type></code>	Optional. Your authentication service type (grid, NIS, or LDAP).
<code>--domain=<domain></code>	Optional. The name of the Avaki domain to log in to.

avaki logout

Log out of an Avaki domain.

Syntax

```
avaki logout
```

- This command has no options.

avaki ls

Display a list of objects (such as files) in the specified grid directory, or if no directory is specified, list the objects in the current working directory. Similar to the Unix **ls** command.

Syntax

```
avaki ls [-r] [<grid-path>+]
```

-r Recursive mode (optional). If you use the **-r** option, the system displays a list of objects in the specified grid directory and any of the directory's subdirectories.

<grid-path>+ Optional. One or more files or grid directories to be listed.

avaki mkdir

Create one or more grid directories. Similar to the Unix **mkdir** command.

Syntax

```
avaki mkdir [-p] [--server | -s]=<grid-server> <grid-path>+
```

-p Optional. If the parent directory in **<grid-path>** doesn't exist, the parent directory will be created.

{--server | -s}=<grid-server> Optional. Specify the name of the Avaki server on which the directory will be created. For example, Bedrock. If you omit this option, the directory will be created on the same server as the parent directory.

<grid-path>+ The grid path of the new directories to create.

avaki monitor --add

Add a monitor test to a monitor service.

Syntax

```
avaki monitor <monitor-name> --add --target=<target-grid-path>
  [--grid-server | --share-server | --proxy-server |
  --dgas-server --secondary-gdc-server] --type=<task>
  --retries=<retries> --period=<period-in-seconds>
```

<monitor-name>	The name of the monitor service to which the monitor test will be added.
--target= <target-grid- path>	The grid directory of the grid server or share server that you want to monitor.
--grid-server --share-server --proxy-server --dgas-server --secondary- gdc-server	Optional. Specify the server to which the monitor test will be added. If you do not specify the server, you must use the --target option to specify full path of the object to monitor. If you do specify the server, make sure the --target option contains only the name of the object, not the full path to the object.
--grid-server	Grid server
--share-server	Share server
--proxy server	Proxy server
--dgas-server	Data grid access server
--secondary- gdc-server	Secondary GDC
--type=<task>	The type of monitor test to set up.

PingTask This test sends one or more IP datagrams to a specified destination host and requests a reply, then measures the round-trip time.

MessageTask This test sends one or more JMS messages to a specified destination host and listens to a response message, then measures the round-trip time.

--retries=
<retries>

The number of times the monitor service will attempt to repeat the test before declaring it a failure. (For example, 1 retry, the default, means that two consecutive pings must fail in order to generate a test failure in the log.)

--period=
<period-in-
secs>

The time (in milliseconds) that the test waits for a response before declaring a failure. Timeout defaults to 10,000 (10 seconds).

avaki monitor --clear

Remove tests from a monitor service. To remove one test, you must specify exactly the same options used when the test was added. (Use [avaki monitor --list](#) (page 172) to display the option values.)

Syntax

To remove all tests from the specified monitor service:

```
avaki monitor <monitor-name> --clear
```

To remove a particular test:

```
avaki monitor <monitor-name> --clear
  [--target=<target-grid-path> [--grid-server | --share-server
  | --proxy-server | --dgas-server --secondary-gdc-server]
  --type=<task> --retries=<retries>
  --period=<period-in-seconds>]
```

<code><monitor-name></code>	The name of the monitor service from which the monitor test will be removed.
<code>--target= <target-grid- path></code>	The grid directory of the Avaki server whose monitor test you want to remove.
<code>--grid-server --share-server --proxy-server --dgas-server --secondary- gdc-server</code>	Optional. Specify the server from which the monitor test will be removed.
<code>--grid-server</code>	Grid server
<code>--share-server</code>	Share server
<code>--proxy server</code>	Proxy server
<code>--dgas-server</code>	Data grid access server

<code>--secondary-gdc-server</code>	Secondary GDC
<code>--type=<task></code>	The type of monitor test to remove.
<code>PingTask</code>	This test sends one or more IP datagrams to a specified destination host and requests a reply, then measures the round-trip time.
<code>MessageTask</code>	This test sends one or more JMS messages to a specified destination host and listens to a response message, then measures the round-trip time.
<code>--retries=<retries></code>	The number of times the monitor service will attempt to repeat the test before declaring it a failure.
<code>--period=<period-in-secs></code>	The time (in milliseconds) that the test waits for a response before declaring a failure. Timeout defaults to 10,000 (10 seconds).

avaki monitor --create

Create a monitor service.

Syntax

```
avaki monitor <monitor-name> --create
  [{--server | -s}=<grid-server>]
```

`<monitor-name>` The name of the monitor service you are creating.

`--{server | -s}=<grid-server>` Optional. Specify the name of the Avaki server on which the monitor service will be created. For example, Bedrock.

avaki monitor --delete

Delete a monitor service from the grid.

Syntax

```
avaki monitor <monitor-name> --delete
```

<monitor-name> The name of the monitor service to delete.

avaki monitor --list

Display a list of all active monitor tests. The list includes details about each test: the test type, target grid path, number of retries, and period.

Syntax

```
avaki monitor <monitor-name> --list
```

<monitor-name> The name of the monitor service from which to obtain a list of active tests.

avaki monitor --result

View the results of all tests of type <task> that have been added to a given monitor service.

Syntax

```
avaki monitor <monitor-name> --result --type=<task>
```

<monitor-name> The name of the monitor service from which to obtain test results.

--type=<task> The type of monitor test to view.

PingTask This test sends one or more IP datagrams to a specified destination host and requests a reply, then measures the round-trip time.

MessageTask This test sends one or more JMS messages to a specified destination host and listens to a response message, then measures the round-trip time.

avaki monitor --start

Restart all tests that have been added to a given monitor service.

Syntax

```
avaki monitor <monitor-name> --start
  [--target=<target-grid-path> [--grid-server | --share-server
  | --proxy-server | --dgas-server --secondary-gdc-server]
  --type=<task> --retries=<retries>
  --period=<period-in-seconds>]
```

<code><monitor-name></code>	The name of the monitor service on which monitor tests will be restarted.
<code>--target=<target-grid-path></code>	The grid directory of the grid server or share server that you want to monitor.
<code>--grid-server --share-server --proxy-server --dgas-server --secondary-gdc-server</code>	Optional. Specify the type of server on which the monitor test will be restarted.
<code>--grid-server</code>	Grid server
<code>--share-server</code>	Share server
<code>--proxy-server</code>	Proxy server
<code>--dgas-server</code>	Data grid access server
<code>--secondary-gdc-server</code>	Secondary GDC

<code>--type=<task></code>	The type of monitor test to restart.
<code>PingTask</code>	This test sends one or more IP datagrams to a specified destination host and requests a reply, then measures the round-trip time.
<code>MessageTask</code>	This test sends one or more JMS messages to a specified destination host and listens to a response message, then measures the round-trip time.
<code>--retries=<retries></code>	The number of times the monitor service will attempt to repeat the test before declaring it a failure.
<code>--period=<period-in-secs></code>	The time (in milliseconds) that the test waits for a response before declaring a failure. Timeout defaults to 10,000 (10 seconds).

avaki monitor --stop

Temporarily stop all active tests that have been added to a given monitor service.

Syntax

```
avaki monitor <monitor-name> --stop
  [--target=<target-grid-path> [--grid-server | --share-server
  | --proxy-server | --dgas-server --secondary-gdc-server]
  --type=<task> --retries=<retries>
  --period=<period-in-seconds>]
```

<monitor-name>	The name of the monitor service on which monitor tests will be stopped.
--target= <target-grid- path>	The grid directory of the grid server or share server that you want to stop monitoring.
--grid-server --share-server --proxy-server --dgas-server --secondary- gdc-server	Optional. Specify the type of server on which the monitor test will be stopped.
--grid-server	Grid server
--share-server	Share server
--proxy server	Proxy server
--dgas-server	Data grid access server
--secondary- gdc-server	Secondary GDC
--type=<task>	The type of monitor test to stop.

PingTask	This test sends one or more IP datagrams to a specified destination host and requests a reply, then measures the round-trip time.
MessageTask	This test sends one or more JMS messages to a specified destination host and listens to a response message, then measures the round-trip time.
<code>--retries= <retries></code>	The number of times the monitor service will attempt to repeat the test before declaring it a failure.
<code>--period= <period-in- secs></code>	The time (in milliseconds) that the test waits for a response before declaring a failure. Timeout defaults to 10,000 (10 seconds).

avaki mv

Move or rename an Avaki directory or a file in an Avaki directory. Similar to the Unix `mv` command.

Syntax

```
avaki mv [--force | -f] <source-grid-path> <target-grid-path>
```

`--force | -f` If you specify this option, the directory or file will be moved even if the `forcedDelete` attribute for a directory or file is set to true. By default, the `forcedDelete` attribute is set to true for system-related directories and files that should be moved only with caution.

`<source-grid-path>` The path to the grid directory or file that will be moved.

`<target-grid-path>` The grid directory to which the file or grid directory will be moved. By default, the file or directory will retain its original name when moved. If you want to rename it, specify the new name in `<target-grid-path>`.

avaki nis --add-schedule

Add a schedule according to which users will be imported automatically.

To use this command, you need write permission on the NIS authentication service.

Syntax

```
avaki nis --add-schedule <auth-service>
  {--periodic-schedule=<period-spec> |
  --cron-schedule=<cron-spec> |
  --one-time-schedule=<time-stamp>}
  [--start-time=<time-stamp>] [[--end-time=<time-stamp>] |
  [--max-iterations=<num-iterations>]]
  [--time-zone=<time-zone>] [--exclusion=<exclusion-name>]+
  [--exclusion-action={ drop | sooner | later }]
  [--reschedule-offset=<offset-spec>]
```

<auth-service> The name of the authentication service for which the schedule will be added.

--periodic-schedule=<period-spec> Specify a periodic schedule for importing users automatically. Use the following syntax:

```
--periodic-schedule="<number-of-units> <unit-type>"
```

<number-of-units> An integer from 1 to 2147483648

<unit-type> seconds, minutes, hours, days, weeks, or months

In the following example, the periodic schedule is set to 30 seconds:

```
--periodic-schedule="30 seconds"
```

<pre>--cron- schedule= <cron-spec></pre>	<p>Specify how often a a cron-based scheduler for importing users will import users. Use the following syntax:</p> <pre>--cron-schedule="<seconds> <minutes> <hours> <days-of-month> <months> <days-of-week> [<years>]"</pre> <p>For details about the <cron-spec> syntax, see Appendix C, “Configuring cron schedules”.</p>						
<pre>--one-time- schedule= <time-stamp></pre>	<p>Specify a single time and date for importing the users. Use the following syntax:</p> <pre>--one-time-schedule="{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}"</pre>						
<pre>--start-time= <time-stamp></pre>	<p>Specify the time and date that the users will be imported, in this format:</p> <pre>{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}</pre>						
<pre>--end-time= <time-stamp></pre>	<p>Specify the time and date that the users will be imported, in this format:</p> <pre>{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}</pre>						
<pre>--max-iterations= <num-iterations></pre>	<p>Specify the number of times a schedule will repeat.</p> <pre><num-iterations></pre> <p>An integer from 1 to 2147483648. If no value is specified, the schedule repeats indefinitely.</p>						
<pre>--time-zone= <time-zone></pre>	<p>Specify the time zone for the schedule. The schedule can be specified according to the grid server’s time zone or relative to Greenwich Mean Time (GMT). If no value is specified for the time zone, the server’s local time zone is used. The format for the time zone is as follows:</p> <table border="0" style="margin-left: 40px;"> <tr> <td>GMT+12:00 Fiji</td> <td>GMT+02:00 Cairo</td> </tr> <tr> <td>GMT+11:00 Noumea</td> <td>GMT+01:00 Paris</td> </tr> <tr> <td>GMT+10:00 Sydney</td> <td>GMT London</td> </tr> </table>	GMT+12:00 Fiji	GMT+02:00 Cairo	GMT+11:00 Noumea	GMT+01:00 Paris	GMT+10:00 Sydney	GMT London
GMT+12:00 Fiji	GMT+02:00 Cairo						
GMT+11:00 Noumea	GMT+01:00 Paris						
GMT+10:00 Sydney	GMT London						

GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii
GMT+03:30 Tehran	GMT-11:00 Samoa
GMT+03:00 Moscow	

```
--exclusion=  
<exclusion-name>
```

The name of an exclusion to this schedule. For information about creating schedule exclusions, see [avaki scheduleexclusion --create --daily](#) (page 201), [avaki scheduleexclusion --create --weekly](#) (page 205), [avaki scheduleexclusion --create --monthly](#) (page 203), [avaki scheduleexclusion --create --yearly](#) (page 207), and [avaki scheduleexclusion --create --custom](#) (page 199).

```
--exclusion-  
action={drop |  
sooner | later}
```

Specify the action to take if the authentication service is scheduled to update users automatically during an exclusionary period.

drop	Don't run the import task during the exclusionary period.
sooner	Reschedule the import task to run before the exclusionary period.
later	Reschedule the import task to run after the exclusionary period.

<code>--reschedule-offset=<offset-spec></code>	Specify the amount of time before or after the exclusionary period to perform a user import. Use the following syntax: <code>--reschedule-offset="<number-of-units><unit-type>"</code>
<code><number-of-units></code>	An integer from 1 to 2147483648
<code><unit-type></code>	Seconds, minutes, hours, days, weeks, or months

Example

This example add a schedule in which users will be imported automatically into the MyNIS authentication service every 2 hours. In this example, the time zone is GMT-05:00 New York.

```
$ avaki nis --add-schedule MyNIS --periodic-schedule="2 hours" --time-zone="GMT-05:00 New York"
```

avaki nis --delete

Delete an NIS authentication service or delete users from an NIS authentication service.

Syntax

To delete an authentication service:

```
avaki nis {--delete | -d} <auth-service>
```

To delete users:

```
avaki nis {--delete | -d} {--user | -u}
  {<user-name>+ | {--all | -a}} <auth-service>
```

`<auth-service>` The name of the authentication service to delete.

`--user | -u` Delete one or more users from the specified authentication service.

<code><user-name>+</code>	The name or names of user accounts to delete. To specify more than one user, enter a space-delimited list of user IDs.
<code>--all -a</code>	Delete all users from the specified authentication service.

avaki nis --delete-schedule

Delete a schedule that specifies when users will be imported automatically.

To use this command, you need write permission on the NIS authentication service.

Syntax

```
avaki nis --delete-schedule <auth-service> <schedule-id>
```

<code><auth-service></code>	The name of the authentication service for which the schedule will be deleted.
<code><schedule-id></code>	The numeric ID of the schedule to delete. To determine a schedule's ID, use avaki nis --list-schedules (page 185).

avaki nis --import

Import user accounts or groups from an NIS authentication service into the current Avaki domain.

Note When you import a user, you import all the groups to which the user belongs. When you import a group, you import all users in the group and all groups to which those users belong.

Syntax

To import user accounts:

```
avaki nis {--import | -i} {--user | -u} {<user-name>+ | [--all  
| -a]} <auth-service>
```

To import groups:

```
avaki nis {--import | -i} {--group | -g} {<group-name>+ |  
[--all | -a]} <auth-service>
```

<code>--user -u</code>	Import user accounts from the specified authentication service.
<code>{<user-name>+ [--all -a]}</code>	The name or names of the user accounts to import. To import all users, specify <code>--all</code> or <code>-a</code> . To specify one or more users, enter a space-delimited list of users.
<code>--group -g</code>	Import groups from the specified authentication service.
<code>{<groups> [--all -a]}</code>	The name or names of the groups to import. To import all groups, specify <code>--all</code> or <code>-a</code> . To specify one or more groups, enter a space-delimited list of groups.
<code><auth-service></code>	The name of the authentication service to import from.

avaki nis --info

Display information about an NIS authentication service.

Syntax

```
avaki nis --info <auth-service>
```

<auth-service> Specify the name of the service.

avaki nis --integrate

Integrate an NIS authentication service.

Syntax

```
avaki nis --integrate <auth-service> --nis-server=<nis-server>  
--nis-domain=<nis-domain> [--refresh-on-login]
```

<auth-service> Devise a name for the new service.

--nis-server=
<nis-server> The name or IP address of the host machine on which the NIS authentication service resides.

--nis-domain=
<nis-domain> The name of the NIS domain for the authentication service you want to integrate.

--refresh-on-
login Optional. Refresh a user's information (such as group memberships) from the external directory each time the user logs in.

avaki nis --list-schedules

Display a list of schedules that specify when users will be imported automatically.

Syntax

```
avaki nis --list-schedules <auth-service>
```

<auth-service> The name of the authentication service for which the schedule will be deleted.

avaki nis --update

Update an NIS authentication service.

Syntax

```
avaki nis --update <auth-service> --nis-server=<nis-server>  
--nis-domain=<nis-domain>
```

<auth-service> Specify the name of the service.

--nis-server=
<nis-server> The name or IP address of the host machine on which the NIS authentication service resides.

--nis-domain=
<nis-domain> The name of the NIS domain for the authentication service you want to modify.

avaki passwd

Change the password for an Avaki user account.

Note Do not use this command on an account imported from an external directory service.

Syntax

```
avaki passwd [--user=<user>] {--old-password | --old}=<old>
  {--new-password | --new}=<new> [--force | -f]
```

<code>--user=<user></code>	Optional. An Avaki account name. If you do not enter an account name on the command line, the system assumes you are changing your own password. To specify a user other than yourself, you must be a member of the Administrators group, and you must use the <code>--force</code> option.
<code>{--old-password --old}=<old></code>	The current password. If you're executing this command on a Unix system and your password contains non-alphanumeric characters (! or &, for instance), enclose your password in single quotation marks to prevent the shell from interpreting the special characters.
<code>{--new-password --new}=<new></code>	The new password. If you're executing this command on a Unix system and your password contains non-alphanumeric characters (! or &, for instance), enclose your password in single quotation marks to prevent the shell from interpreting the special characters.
<code>--force -f</code>	(Administrators only) Do not display the prompt to enter the old password.

avaki permissions

Display read, write, execute, and delete permissions for each user or group in the access control list for an object.

Syntax

```
avaki permissions <grid-path>+
```

<grid-path>+ The path to the grid file or grid directory for which you are viewing permissions. You can enter an absolute path or a relative one.

avaki plugin --generate

Run the data service Plug-in Wizard to generate skeleton Java code and build.xml and manifest files for a data service plug-in. For more information about writing your own data service plug-ins, see the *Sybase Avaki EII Provisioning and Advanced Data Integration Guide*.

Syntax

```
avaki plugin --generate [--plugin-name=<plugin-name>]
  [--package=<package-name>]
  [--impl-class=<implementation-class-name>]
  [--method-name=<implementation-method>]
  [--parameter=<parameter-spec>] * [--input=<stream-spec>] *
  [--output=<stream-spec>] *
  [--target-dir=<output-target-path>]
  [--api-classes=<api-class-location>]
  [--template-dir=<template-path>]
```

--plugin-name=
<plugin-name> Optional. The name of the plug-in. If this option is not specified, the default value is MyPlugin. The plug-in name will be included in the manifest file. The JAR file will be named <plugin-name>.jar.

- `--package=
<package-name>` Optional. The package name of the implementation class. For example, if you specify the value “com.avaki.mypackage” the plug-in implementation class will be placed in the specified package. If this option is not specified, the implementation class will not be put in any packages.
- `--impl-class=
<implementation-
class-name>` Optional. The class name of the plug-in implementation that the tool will generate. If this option is not specified, the class name is the value of `<plugin-name>`.
- `--method-name=
<implementation-
method>>` Optional. The name of the method that implements the plug-in. The plug-in tool creates an empty method, and a developer adds custom code to it. The default value is `plugin`.
- `--parameter=
<parameter-
spec>*` Optional. A parameter specification for this plug-in. If the plug-in accepts more than one parameter, you can specify multiple `--parameter` options. Use the following form to specify a parameter:

`<name>=<value>`

Enclose parameters in double quotation marks (""). For example:

```
--parameter="name=myString;type=VARCHAR"
```

If you specify multiple parameters, use semicolons (;) to separate them.

Specify any of the following names and values:

name A unique name for this parameter

type The type of this parameter, which can be one of the following:

ARRAY	INTEGER
BIGINT	JAVA_OBJECT
BINARY	LONGVARCHAR
BIT	NUMERIC
BOOLEAN	REAL
BLOB	SMALLINT
CHAR	TIME
CLOB	TIMESTAMP
DATE	TINYINT
DECIMAL	VARBINARY
DOUBLE	VARCHAR
FLOAT	

islist Indicates whether this parameter's value is a list or a singleton. The possible values are true or false. The default is false.

The following example generates a plug-in that accepts two parameters:

```
avaki plugin -generate
--parameter="name=param1;type=INTEGER"
--parameter="name=param2;type=
VARCHAR,islist=true"
```

In this example, the type for the second parameter, "param2," is a list of strings.

```
--input=
<stream-spec>
```

Optional. The input specification for the plug-in.

Specify one or more equal-sign-separated <name>=<value> pairs, where <name> is a unique name for the input specification and <value> is its value. If you enter more than one <name>=<value> pair, use a semicolon (;) to separate the pairs.

Specify the following names and values:

name A unique name for this input.

type The stream type of this input data, which can be one of the following:

- XML: Custom data in XML format
- ByteStream: Bytes of raw data
- ResultSet: Database result rowsets

	islist	(Optional) Indicates whether this input is a list or a single item. The possible values are true or false. The default is false
--output= <stream-spec>		Optional. The output destination for the plug-in. Specify one or more equal sign-separated <name>=<value> pairs, where <name> is a unique name for the output specification and <value> is its value. If you enter more than one <name>=<value> pair, use a semicolon (;) to separate the pairs.
		Specify the following names and values:
	name	A unique name for this output
	type	The stream type of this output data, which can be one of the following: <ul style="list-style-type: none"> • XML: Custom data in XML format • ByteStream: Bytes of raw data • ResultSet: Database result rowsets
	islist	(Optional) Indicates whether this output is a list or a singleton. The possible values are true or false. The default is false.
--target-dir= <output- target-path>		Optional. The name of the directory in which the plug-in wizard will store the generated code. The default location is ./plugin_src. If the necessary directories do not exist, they will be created.
--api-classes= <api-class- location>		Optional. The location of the Avaki API classes that are required to compile the generated plug-in code. You can specify a JAR file or the path to the classes. The default location is <avaki-install-dir>/apilib/javaki_plugin_api.jar.
--template-dir= <template- path>		Optional. The location of the directory that contains template files the plug-in wizard reads when generating code. The default location is <avaki-install-dir>/templates/plugins.

Example

```
$ avaki plugin -generate -plugin-name=Example
--package=com.avaki.examples
--impl-class=MyExample --method-name=myRunMethod
--parameter="name=myInt;type=INTEGER"
--parameter="name=myString;type=VARCHAR"
--input="name=input1;type=XML"
--input="name=input2;type=XML"
--output="name=myOutput;type=XML"
```

avaki proxy --add

Add another Avaki domain's proxy server to your Avaki domain's proxy routing table. This is part of the procedure for establishing an interconnection from one domain to another.

Note If anyone modifies your domain's proxy routing table during a CLI session, execute the command **avaki client --connect** (page 46) to force your command client to pick up the new routing information.

Syntax

```
avaki proxy {--add | -a} [--unencrypted] [--compress-data]
  <site-ID> {<proxy-server> <proxy-port> |
  <proxy-server>:<proxy-port>}
```

--unencrypted Optional. Turn off SSL encryption, which is enabled by default. When SSL is disabled, performance will improve, but data is not secure. SSL encryption should not be disabled unless both grid domains are within a larger secure environment, such as behind a corporate firewall.

--compress-data Optional. Use data compression when communicating with this proxy server.

<site-ID> The name of the Avaki domain to connect with.

<code><proxy-server></code>	The IP address or DNS name of the machine on which the other domain's proxy server is running. We recommend that you enter the proxy host's IP address rather than its DNS name.
<code><proxy-port></code>	The proxy server's HTTP or HTTPS listening port. These are commonly: HTTP: 18080 HTTPS: 18443

avaki proxy --delete

Delete a proxy server from your Avaki domain's proxy routing table.

Syntax

```
avaki proxy {--delete | -d} <site-ID> {<proxy-server>
  <proxy-port> | <proxy-server>:<proxy-port>}
```

<code><site-ID></code>	The name of the Avaki domain to which the proxy server belongs.
<code><proxy-server></code>	The IP address or DNS name of the proxy server to be deleted.
<code><proxy-port></code>	The proxy server's HTTP or HTTPS listening port. These are commonly: HTTP: 18080 HTTPS: 18443

avaki proxy --list

List the proxy server entries in the proxy routing table.

Syntax

```
avaki proxy {--list | -l}
```

– This command has no options.

avaki pwd

Display the name of the current Avaki directory. Similar to the Unix **pwd** command.

Syntax

```
avaki pwd
```

– This command has no options.

avaki replica --add

Add a secondary Avaki domain controller to a domain.

Syntax

```
avaki replica {--add | -a} <machine>[:<port>]
```

<machine>[:<port>] Specify the IP address or DNS name of the machine on which the secondary GDC is running. If the secondary GDC is using a connect port other than 3099, the default, specify the connect port number.

avaki replica --config

Configure the refresh interval for a secondary GDC.

Syntax

```
avaki replica {--config | -c}  
  {--interval}=<backup-interval-in-seconds>
```

--interval=
<backup-interval-
in-seconds> How frequently (in seconds) the secondary GDC is backed up.
 We recommend a minimum of 900 seconds (15 minutes).

avaki replica --delete

Delete a secondary grid domain controller from an Avaki domain.

Syntax

```
avaki replica {--delete | -d} <machine>[:<port>]
```

<machine>[:<port>] The IP address or DNS name of the machine on which the secondary GDC is running. If the secondary GDC is using a connect port other than 3099, the default, specify the connect port number.

avaki replica --info

View the names and settings of any secondary GDCs in your grid domain.

Syntax

```
avaki replica --info
```

– This command has no options.

avaki replica --synch

Force an immediate update of a secondary GDC.

Syntax

```
avaki replica --synch
```

– This command has no options.

avaki rm

Delete a file, an Avaki directory, or the contents of an Avaki share.

Caution When you delete a shared file or directory from a grid domain, the change is reflected in the source file system—not in the grid only. If you want to shut down a share without affecting the source files, use [avaki server --share --stop](#) on page 231).

Syntax

```
avaki rm [--recursive | -r] [--force | -f] <grid-path>
```

- | | |
|--------------------------------|---|
| <code>--recursive -r</code> | Optional. If you use <code>-r</code> to delete a directory, the system deletes all the directory's contents, including all subdirectories and their contents. |
| <code>--force -f</code> | Optional. Force the file, grid directory, or contents of a share to be removed from the grid directory. |
| <code><grid-path></code> | The path to the file, directory, or share to be deleted. The file, directory, or share must be in a grid directory. You can enter an absolute path or a relative one. |

avaki schedule --delete

Delete a schedule from the specified server.

Syntax

```
avaki schedule --delete <server-name> <schedule-id>
```

<server-name> The name of the Avaki server from which the schedule will be deleted. For example, Bedrock.sybase.com.

<schedule-id> The numeric ID of the schedule to delete. To determine a schedule's ID, use [avaki schedule --list](#) (page 198).

avaki schedule --info

Display information about the specified schedule.

Syntax

```
avaki schedule --info <server-name> <schedule-id>
```

<server-name> The name of the Avaki server for which information will be displayed. For example, Bedrock.sybase.com.

<schedule-id> The numeric ID of the schedule. To determine a schedule's ID, use [avaki schedule --list](#) (page 198).

avaki schedule --list

List the schedules registered on the specified server.

Syntax

```
avaki schedule --list <server-name>
```

<server-name> The name of the Avaki server for which schedules will be listed. For example, Bedrock.sybase.com.

avaki schedule --print-iterations

List the times when the specified schedule will execute.

Syntax

```
avaki schedule --print-iterations <server-name> <schedule-id>  
  <num-iterations>
```

<server-name> The name of the Avaki server for the execution times will be listed. For example, Bedrock.sybase.com.

<schedule-id> The numeric ID of the schedule. To determine a schedule's ID, use [avaki schedule --list](#) (page 198).

<num-iterations> The number of execution times to display.

avaki scheduleexclusion --create --custom

Create a custom exclusion to a schedule.

Syntax

```
avaki scheduleexclusion --create --custom
  [--description=<description-text>] [--time-zone=<time-zone>]
  <schedule-exclusion-name>
  {<start-time-stamp> <end-time-stamp>}+
```

--description=
<description-text> Specify a description for the exclusion.

--time-zone=
<time-zone> Specify whether the time zone for the excluded period is the grid
server's local time or relative to Greenwich Mean Time (GMT).
The format for the time zone is as follows:

GMT+12:00 Fiji	GMT+02:00 Cairo
GMT+11:00 Noumea	GMT+01:00 Paris
GMT+10:00 Sydney	GMT London
GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii

GMT+03:30 Tehran GMT-11:00 Samoa

GMT+03:00 Moscow

<schedule-
exclusion-name>

Specify a name for the schedule exclusion.

{<start-time-stamp>
<end-time-stamp>}+

Specify one or more ranges of times and dates during which the exclusion will be in effect, in this format:

```
{ "<hh>:<mm>:<ss> {am | pm} <mm>-<dd>-<yyyy>"
  "<hh>:<mm>:<ss> {am | pm} <mm>-<dd>-<yyyy>" }
```

avaki scheduleexclusion --create --daily

Create a daily exclusion to a schedule.

Syntax

```
avaki scheduleexclusion --create --daily
  [--description=<description-text>] [--time-zone=<time-zone>]
  <schedule-exclusion-name> <start-time-of-day>
  <end-time-of-day>
```

--description=
<description-text> Specify a description for the exclusion.

--time-zone=
<time-zone> Specify whether the time zone for the excluded period is the grid
server's local time or relative to Greenwich Mean Time (GMT).
The format for the time zone is as follows:

GMT+12:00 Fiji	GMT+02:00 Cairo
GMT+11:00 Noumea	GMT+01:00 Paris
GMT+10:00 Sydney	GMT London
GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii

GMT+03:30 Tehran GMT-11:00 Samoa

GMT+03:00 Moscow

<schedule-
exclusion-name>

Specify a name for the schedule exclusion.

<start-time-of-day>
<end-time-of-day>

Specify the range of times and dates during which the exclusion will be in effect, in this format:

"<hh>:<mm>:<ss> {am | pm}" "<hh>:<mm>:<ss> {am
| pm}"

avaki scheduleexclusion --create --monthly

Create a monthly exclusion to a schedule.

Syntax

```
avaki scheduleexclusion --create --monthly
  [--description=<description-text>] [--time-zone=<time-zone>]
  {--month-day=<month-day>}+ <schedule-exclusion-name>
  [<start-time-of-day> <end-time-of-day>]
```

--description=
<description-text> Specify a description for the exclusion.

--time-zone=
<time-zone> Specify whether the time zone for the excluded period is the grid
server's local time or relative to Greenwich Mean Time (GMT).
The format for the time zone is as follows:

GMT+12:00 Fiji	GMT+02:00 Cairo
GMT+11:00 Noumea	GMT+01:00 Paris
GMT+10:00 Sydney	GMT London
GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii

GMT+03:30 Tehran GMT-11:00 Samoa

GMT+03:00 Moscow

```
{--month-day=
<month-day>}+
```

Specify the days of the month during which the exclusion will be in effect, in this format:

```
{ 1 - 31 | L }
```

L represents the last day of the month.

```
<schedule-
exclusion-name>
```

Specify a name for the schedule exclusion.

```
[<start-time-of-
day> <end-time-of-
day>]
```

Optional. Specify the range of times and dates during which the exclusion will be in effect, in this format:

```
"<hh>:<mm>:<ss> {am | pm}" "<hh>:<mm>:<ss> {am
| pm}"
```

avaki scheduleexclusion --create --weekly

Create a weekly exclusion to a schedule.

Syntax

```
avaki scheduleexclusion --create --weekly
  [--description=<description-text>] [--time-zone=<time-zone>]
  {--weekday=<weekday-name>}+ <schedule-exclusion-name>
  [<start-time-of-day> <end-time-of-day>]
```

--description= <description-text> Specify a description for the exclusion.

--time-zone= <time-zone> Specify whether the time zone for the excluded period is the grid server's local time or relative to Greenwich Mean Time (GMT). The format for the time zone is as follows:

GMT+12:00 Fiji	GMT+02:00 Cairo
GMT+11:00 Noumea	GMT+01:00 Paris
GMT+10:00 Sydney	GMT London
GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii

GMT+03:30 Tehran GMT-11:00 Samoa

GMT+03:00 Moscow

```
{--weekday=
<weekday-name>}+
```

Specify the days of the week during which the exclusion will be in effect, in this format (the format is not case-sensitive):

```
{Sun | Mon | Tue | Wed | Thu | Fri | Sat}
```

```
<schedule-
exclusion-name>
```

Specify a name for the schedule exclusion.

```
[<start-time-of-
day> <end-time-of-
day>]
```

Optional. Specify the range of times and dates during which the exclusion will be in effect, in this format:

```
"<hh>:<mm>:<ss> {am | pm}" "<hh>:<mm>:<ss> {am
| pm}"
```

avaki scheduleexclusion --create --yearly

Create a yearly exclusion to a schedule.

Syntax

```
avaki scheduleexclusion --create --yearly
  [--description=<description-text>] [--time-zone=<time-zone>]
  {--date=<mm>-<dd>}+ <schedule-exclusion-name>
  [<start-time-of-day> <end-time-of-day>]
```

--description=
<description-text> Specify a description for the exclusion.

--time-zone=
<time-zone> Specify whether the time zone for the excluded period is the grid
server's local time or relative to Greenwich Mean Time (GMT).
The format for the time zone is as follows:

GMT+12:00 Fiji	GMT+02:00 Cairo
GMT+11:00 Noumea	GMT+01:00 Paris
GMT+10:00 Sydney	GMT London
GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii

GMT+03:30 Tehran GMT-11:00 Samoa

GMT+03:00 Moscow

`{--date=<mm>--<dd>}`+ Specify the days of the month during which the exclusion will be in effect, in this format:

`<mm>--{<dd> | L}`

L represents the last day of the month. For example, both of the following specify October 31:

```
--date=10-31
--date=10-L
```

`<schedule-exclusion-name>` Specify a name for the schedule exclusion.

`[<start-time-of-day> <end-time-of-day>]` Optional. Specify the range of times and dates during which the exclusion will be in effect, in this format:

```
"<hh>:<mm>:<ss> {am | pm}" "<hh>:<mm>:<ss> {am | pm}"
```

avaki scheduleexclusion --delete

Delete a schedule exclusion from the grid domain.

Syntax

```
avaki scheduleexclusion --delete <schedule-exclusion-name>
```

`<schedule-exclusion-name>` Specify the name of the schedule exclusion to delete.

avaki scheduleexclusion --info

Display information about a specified schedule exclusion.

Syntax

```
avaki scheduleexclusion --info <schedule-exclusion-name>
```

<code><schedule-exclusion-name></code>	Specify the name of the schedule exclusion about which information will be displayed.
--	---

avaki scheduleexclusion --list

List the names of all schedule exclusions in the grid domain.

Syntax

```
avaki scheduleexclusion --list
```

<code>_</code>	This command has no options.
----------------	------------------------------

avaki search --create

Set up a search service to index attributes for all objects in a particular grid directory.

Syntax

```
avaki search --create [--include=<attribute-name>]
  [--exclude=<attribute-name>] [--interval=<rehash-interval>]
  [{--server | -s}=<grid-server>] [--recursive | -r]
  <search-service-name> <grid-directory-to-index>
```

<code>--include=</code> <code><attribute-name></code>	Optional. Specify which attributes the search service will include in the index.
<code>--exclude=</code> <code><attribute-name></code>	Optional. Specify which attributes the search service will exclude from the index.
<code>--interval=</code> <code><reindex-</code> <code>interval></code>	Optional. How often (in seconds) attribute information will be updated. Use the <code>--interval</code> flag when creating a search service; to change the reindex interval for an existing search service, use the <code>avaki rehash</code> command.
<code>{--server -s}=</code> <code><grid-server></code>	Optional. Specify the name of the Avaki server on which the search service will be created. For example, <code>Bedrock</code> . If you omit this option, the search service is created on the current server.
<code>[--recursive -r]</code>	Optional. Recursively search the specified grid directory and all its contents.
<code><search-</code> <code>service-name></code>	Specify a name for the search service.
<code><grid-directory-</code> <code>to-index></code>	The root directory from which the attributes will be collected. All searches using a particular search service are recursive—that is, they search everything under the <code><grid-directory-to-index></code> .

avaki search --delete

Delete a search service from a grid domain.

Syntax

```
avaki search {--delete | -d} <search-service-name>  
  [{--domain | -D}=<domain>]
```

<search-service-name> The name of the search service you're deleting.

{--domain | -D}=<domain> Optional. The grid domain of the search service to delete. If you do not specify a domain, the system assumes the path is in the current grid domain.

avaki search (execute)

Perform a search for grid objects (such as files and directories) that have a particular attribute or attribute value. To use this command, you must previously have created a search service (see [avaki search --create on page 210](#)).

Note The web UI has more powerful search capabilities than the CLI, including the ability to search for values that are greater than, less than, or not equal to the specified <search-value>. See the *Sybase Avaki EII Administration Guide* for instructions on using the search tools in the web UI.

Syntax

```
avaki search --name=<attribute-name> [--value=<search-value>]
  [{--operation | -op}=<operation>] [{--type | -t}=<type>]
  <search-service-name> [{--domain | -D}=<domain>]
```

`--name=`
`<attribute-name>` Specify the attribute to search for. For example, to search for a file by name, enter `--name=system/name`. (You can use the command [avaki attribute --list \(page 21\)](#) to list other attributes to search for.)

`--value=`
`<search-value>` Optional. Specify a value for the attribute in the `--name` option. If you use this option, the search returns a list of grid objects (files, for example) for which the target attribute is set to `<search-value>`. If you omit this option, the search returns a list of all grid objects that have the target attribute.

You can use a wildcard to indicate a character or set of characters to use in searches. A wildcard can indicate a point in a search string where any character or string is a good match, where any of one or more specified characters is a good match, or where any of the specified characters is not an acceptable match. The following wildcard characters may be used in searches:

`%` or `*` Match any character or string at this position in the search pattern

`_` or `?` Match any single character at this position in the search pattern


```
{--operation |
-op}=<operation>
```

Optional. Specify a logical operator to be used in the search, in this format:

= is equal to

<> is not equal to

> is greater than

>= is greater than or equal to

< is less than

<= is less than or equal to

If --operation is not specified, <operation> is assumed to be = (is equal to).

```
{--type | -t}=
<type>
```

Specify the type of the attribute to search for.

string

integer

real

date

time

timestamp

If --type is not specified, <type> is assumed to be string.

<code><search-service-name></code>	Specify the name of the search service to use for this search.
<code>{--domain -D}= <domain></code>	Optional. The grid domain of the search service to use. If you do not specify a domain, the system assumes the path is in the current grid domain.

avaki search --get-rehash-interval

View the rehash interval for a search service.

Syntax

```
avaki search --get-rehash-interval <search-service-name>
  [{--domain | -D}=<domain>]
```

<code><search-service-name></code>	The name of the search service whose rehash interval you're viewing.
<code>{--domain -D}= <domain></code>	Optional. The grid domain of the search service whose rehash interval you're viewing. If you do not specify a domain, the system assumes the path is in the current grid domain.

avaki search --info

Display information about a search service.

Syntax

```
avaki search {--info | -i} <search-service-name>
  [{--domain | -D}=<domain>]
```

<code><search-service-name></code>	The name of the search service whose information you want to display.
<code>{--domain -D}=<domain></code>	Optional. The grid domain of the search service to display information about. If you do not specify a domain, the system assumes the path is in the current grid domain.

avaki search --rehash

Forces a search service to reindex its contents.

Syntax

```
avaki search --rehash <search-service-name>
  [{--domain | -D}=<domain>]
```

<code><search-service-name></code>	The name of the search service you're rehashing.
<code>{--domain -D}=<domain></code>	Optional. The grid domain of the search service to rehash. If you do not specify a domain, the system assumes the path is in the current grid domain.

avaki search --set-rehash-interval

Configure the rehash interval for a search service.

Syntax

```
avaki search --set-rehash-interval <search-service-name>  
  [{--domain | -D}=<domain>] <new-rehash-interval>
```

<code><search-service-name></code>	The name of the search service whose rehash interval you're configuring.
<code>{--domain -D}=<domain></code>	Optional. The grid domain of the search service whose rehash interval you're configuring. If you do not specify a domain, the system assumes the path is in the current grid domain.
<code><new-rehash-interval></code>	Specify how frequently (in seconds) updates are propagated from the local file system into the shared directory in the grid. Default: 0 (never rehash) Range: 0 to 2147483648 seconds (about 68 years)

avaki security --config

Configure the encryption level for all files in an Avaki directory or Avaki share.

To display the encryption level for a grid object, use [avaki security --info](#) (page 222).

Syntax

```
avaki security --config --encryption-level={clear | encrypted}  
  <target-grid-path>
```

`--encryption-level={clear | encrypted}` Specify the encryption level for the files: clear (encryption is turned off for this directory or share) or encrypted (encryption is turned on).

`<target-grid-path>` The path to the file or directory for which you are modifying the encryption level.

avaki security --default-gid

Configure the default local GID for groups in a specified grid authentication service, or display the current default local GID for the specified authentication service. When no specific mapping exists for a group in that authentication service, the data grid access server allows members of the group to access grid data (subject to grid access controls) and passes the specified GID to the NFS client when the client needs to know the ownership of a file or directory. If you do not configure a default GID, users in the target authentication service who are not covered by specific user or group mappings will not be able to access grid data through the DGAS.

Syntax

To display the current default GID:

```
avaki security --default-gid
```

To configure a default GID:

```
avaki security --default-gid [--auth-service-path=<path>]
    [<gid>]
```

- | | |
|---|---|
| <code>--auth-service-path=<path></code> | Optional. The parent directory of the grid authentication service in which to configure the default GID. For example, <code>/System/Domains/Bedrock/Services/AuthServices/Ldap/BedrockLdap</code> . The default is <code>/System/Domains/<domain-name>/Services/AuthServices/Grid/DefaultAuthService</code> . |
| <code><gid></code> | Optional. Specify the default GID. If you omit this option, the system displays the current default local GID. |

avaki security --default-group

Use this command to configure a DGAS to use the specified group name as the default group from within the specified authentication service path if an unknown or unmapped user accesses the grid from an NFS client.

Syntax

```
avaki security --default-group [--auth-service-path=<path>]
    [<group-name>]
```

- | | |
|---|--|
| <code>--auth-service-path=<path></code> | Optional. The parent directory of the specified or default group's grid authentication service. For example, /System/Domains/Bedrock/Services/AuthServices/Ldap/BedrockLdap. The default is /System/Domains/<domain-name>/Services/AuthServices/Grid/DefaultAuthService. |
| <code><group-name></code> | Optional. The name of the group in the specified authentication service that you are configuring as the default grid group for client groups to map to. If you omit this option, the system displays the default group name. |

avaki security --default-uid

Configure the default local UID for users in a specified grid authentication service, or display the current default local UID for the specified grid authentication service. When no specific mapping exists for a user in that authentication service, the data grid access server allows users to access grid data (subject to grid access controls) and passes the specified UID to the NFS client when the client needs to know the ownership of a file or directory. If you do not configure a default UID, users in the target authentication service who are not covered by specific user or group mappings will not be able to access grid data through the DGAS.

Syntax

To display the current default UID:

```
avaki security --default-uid
```

To configure a default UID:

```
avaki security --default-uid [--auth-service-path=<path>]
    [<uid>]
```

- | | |
|---|---|
| <code>--auth-service-path=<path></code> | Optional. The parent directory of the grid authentication service for which you are configuring a default UID. For example, <code>/System/Domains/Bedrock/Services/AuthServices/Ldap/BedrockLdap</code> . The default is <code>/System/Domains/<domain-name>/Services/AuthServices/Grid/DefaultAuthService</code> . |
| <code><uid></code> | Optional. Specify the default UID. If you omit this option, the system displays the default UID. |

avaki security --default-user

If you want to allow unknown or unmapped users to access the grid from NFS clients, configure a DGAS to use the specified user name as the default user from within the specified authentication service path.

Syntax

```
avaki security --default-user [--auth-service-path=<path>]
    [<user-name>]
```

`--auth-service-path=<path>` Optional. The parent directory of the default user's grid authentication service. For example, `/System/Domains/Bedrock/Services/AuthServices/Ldap/BedrockLdap`. The default is `/System/Domains/<domain-name>/Services/AuthServices/Grid/DefaultAuthService`.

`<user-name>` Optional. The name of the user in the specified authentication service that you are configuring as the default grid user for client users to map to. If you omit this option, the system displays the default user name.

avaki security --gid

Configure the NFS client group ID for the specified grid group, or display the current GID for the specified grid group.

Syntax

To display the current GID:

```
avaki security --gid --group=<path>
```

To configure a GID:

```
avaki security --gid --group=<path> <new-gid>
```

To remove a GID:

```
avaki security --gid --group=<path> --delete
```

<code>--group=<path></code>	The grid directory of the group.
<code>--delete</code>	Optional. Remove the GID configured for the specified grid group.
<code><new-gid></code>	Optional. Specify the GID that you are mapping to the grid group.

avaki security --info

Display the encryption level for a grid directory, file, or share.

To set encryption levels, use [avaki security --config](#) (page 217).

Syntax

```
avaki security --info <target-grid-path>
```

<code><target-grid-path></code>	The path to the grid directory, file, or share for which you want to view the encryption level.
---------------------------------------	---

avaki security --uid

Configure the NFS client user ID for the specified grid user, or display the current UID for the specified grid user.

Syntax

To display the current UID:

```
avaki security --uid --user=<path>
```

To configure a UID:

```
avaki security --uid --user=<path> <new-uid>
```

To remove a UID:

```
avaki security --uid --user=<path> --delete
```

- | | |
|----------------------------------|---|
| <code>--user=<path></code> | The grid directory of the user. For example, /System/LocalDomain/Services/AuthServices/Grid/DefaultAuthService/Users/wilma. |
| <code>--delete</code> | Optional. Remove the UID configured for the specified grid user. |
| <code><new-uid></code> | Optional. Specify the UID that you are mapping to the grid user. |

avaki server --dgas --connect

Connect a data grid access server to a grid domain. To use this command, you must be logged into Avaki as a member of the Administrators group.

These are the tasks you must complete to set up a DGAS:

1. If the DGAS will run on a Windows machine and if you want the DGAS to serve CIFS shares, disable native CIFS Windows services running on ports 445, 137, 138, and 139. See the *Sybase Avaki EII Administration Guide* for details.
2. (Optional.) Create a properties file for the DGAS. (See [“Appendix B, “Setting up DGAS properties files”](#) for more information.)
3. Use ***dgas --start*** to start the DGAS. See the *Sybase Avaki EII Administration Guide* for details.
4. Use ***avaki server --dgas --connect*** (page 224) to connect the DGAS to a grid domain.
5. Use ***avaki dgas --set-property*** (page 114) to set properties for the DGAS. (You can skip this step if you created a properties file in item 2.)
6. Use ***avaki dgas --initialize*** (page 105) to initialize the DGAS’s credentials. If you created a properties file in item 2., specify the location of the properties file in this command.
7. Set up CIFS shares (***avaki dgas --create-cifs-share*** (page 97)).
8. Access data catalog files on CIFS or NFS clients.

Syntax

```
avaki server --dgas --connect <machine>:[<port>] <dgas-name>
```

<code><machine></code>	The IP address or DNS name of the machine hosting the DGAS that you are connecting to the grid domain. If you want the DGAS to use a connect port other than 1399, the default, specify the connect port number.
<code>:[<port>]</code>	
<code><dgas-name></code>	The name you are assigning to the DGAS.

avaki server --dgas --destroy

Destroy a data grid access server and delete its state information and caches. To use this command, you must be logged in to Avaki as a member of the Administrators group.

Caution This command irrevocably destroys the DGAS. To halt the DGAS without destroying it or losing its state and caches, use **avaki server --dgas --stop**, below.

Syntax

```
avaki server --dgas --destroy [-f | --force] <dgas-name>
```

-f | --force Optional. If you specify the **-f** option, the server will be removed from the grid even if the server is not running. If the server is running, any of the server's links to domainwide database operations (such as `/System/LocalDomain/Services/DatabaseServices/MyDBConnector/MyDBOperation`) or data services (such as `/System/LocalDomain/Services/DataServices/MyDataService`) will be removed.

<dgas-name> The name of the DGAS to be destroyed.

avaki server --dgas --stop

Stop a data grid access server, but preserve its state information and caches. This command can be issued from a remote machine. (Use **dgas --start** (page 7) to restart a stopped DGAS.)

Syntax

```
avaki server --dgas --stop <dgas-name>
```

<dgas-name> The name of the DGAS to be stopped.

avaki server --grid --connect

Connect a grid server to an Avaki domain.

Syntax

```
avaki server --grid --connect <machine>:[<port>]
```

--connect
<machine>
:[<port>]

For <machine>, specify the IP address or DNS name of the machine hosting the Avaki server that you are connecting to the grid domain. For <port>, specify connect port number for the machine if the server is using a nondefault connect port.

avaki server --grid --destroy

Disconnect a grid server from an Avaki domain and destroy it.

Caution The **avaki server --grid --destroy** command irrevocably destroys the grid server. To halt the grid server without destroying it or losing its state and caches, use **avaki server --grid --stop** (page 228).

Syntax

```
avaki server --grid --destroy [-f | --force]
    [<machine>[:<port>]]
```

--destroy [-f | --force] [**<machine>[:<port>]**]

If you specify the **-f** option, the grid server will be removed from the Avaki domain even if the grid server is not running. If the server is running, any of the server's links to domainwide database operations (such as `/System/LocalDomain/Services/DatabaseServices/MyDBConnector/MyDBOperation`) or data services (such as `/System/LocalDomain/Services/DataServices/MyDataService`) will be removed.

For **<machine>[:<port>]**, specify one of the following:

- If the grid server to be destroyed is alive, specify the IP address or DNS name of the machine hosting the Avaki server that you are destroying. If the server is using a nondefault connect port, specify the connect port number.
- If the grid server to be destroyed is not alive, specify the server's node name; do not specify a connect port. To determine the node name, execute the **avaki ls** command on the grid path to the Servers directory. For example, if you issue the following command:

```
avaki ls /System/LocalDomain/Servers
```

The result will be similar to the following:

```
Servers (directory)
  . (directory)
  .. (directory)
  slatequarry.avaki.com (directory)
```

In this example, the server's node name is: `slatequarry.avaki.com`

avaki server --grid --stop

Shut down an Avaki grid server. (To restart the grid server, see [grid-server --start](#) on [page 11](#).)

Syntax

```
avaki server --grid --stop <machine>[:<port>]
```

`--stop <machine>[:<port>]` The IP address or DNS name of the machine hosting the Avaki server that you are stopping. If the machine is on a nondefault connect port, specify the connect port number.

avaki server --proxy

Connect an Avaki firewall proxy server to a grid domain, destroy a firewall proxy server, or stop a firewall proxy server.

Caution The `avaki server --proxy --destroy` command irrevocably destroys the proxy server. To halt the proxy server without destroying it or losing its state and caches, use `avaki server --proxy --stop` ([page 228](#)).

Syntax

```
avaki server --proxy {--connect <machine>[:<port>] | --destroy  
| --stop} [-f | --force] <proxy-name>
```

`--connect <machine>[:<port>]` Connect the Avaki proxy server on the specified machine to the current grid domain. For `<machine>`, enter the IP address or DNS name of the machine hosting the proxy server. If the proxy server is using a connect port other than 1199, the default, specify the connect port number.

`--destroy` Delete the specified Avaki proxy server.

<code>--stop</code>	Stop (but do not destroy) a proxy server that is registered as a service. (To restart a proxy server, use proxy-server --start (page 14) or proxy-server --register (page 13).)
<code>-f</code> <code>--force</code>	Optional. If you specify the <code>-f</code> option, the server will be removed from the grid even if the server is not alive. If the server is alive, any of the server's links to domainwide database operations (such as <code>/System/LocalDomain/Services/DatabaseServices/MyDBConnector/MyDBOperation</code>) or data services (such as <code>/System/LocalDomain/Services/DataServices/MyDataService</code>) will be removed.
<code><proxy-name></code>	The grid name of the proxy server.

avaki server --share --connect

Connects a share server to a grid server. You can connect a share server to multiple grid servers.

Syntax

```
avaki server --share --connect <share-server-logical-name>
  <share-server-host-name> [:<port>]
  <grid-server-host-name> [:<port>]
```

<code><share-server-logical-name></code>	The logical name of the new share server.
<code><share-server-host-name> [:<port>]</code>	The machine name or IP address of the new share server. If you want the share server to use a connect port other than 2099, the default, specify the connect port number.
<code><grid-server-host-name> [:<port>]</code>	The DNS name or IP address of the grid server to which to join this share server. If the grid server is using a connect port other than 3099, the default, specify the connect port number.

avaki server --share --disconnect

Disconnects a share server from a grid server. A share server may be connected to multiple grid servers.

Note Before you disconnect a share server, make sure there are no Avaki shares using the share server. You can disable the shares by taking them off line (see [“avaki share --set-status” on page 244](#)) or you can dissociate them from the share server (see [“avaki share --remove-share-servers” on page 241](#))

Syntax

```
avaki server --share --disconnect [-f | --force]
    <share-server-name> <grid-server-name>
```

-f | --force Optional. If you specify the -f option, the server will be removed from the grid even if the server is not alive. If the server is alive, any of the server's links to domainwide database operations (such as /System/LocalDomain/Services/DatabaseServices/MyDBConnector/MyDBOperation) or data services (such as /System/LocalDomain/Services/DataServices/MyDataService) will be removed.

<share-server-name> The logical name of the share server.

<grid-server-name> The DNS name or IP address of the grid server to disconnect from.

avaki server --share --stop

Shut down a share server. When you shut down a share server, the share server's shared directories and file names remain visible in the grid data catalog, but the shared files are inaccessible.

(You can restart the share server using the [share-server --start \(page 17\)](#) or [share-server --register \(page 16\)](#) command on the machine where the share server runs.)

Syntax

```
avaki server --share --stop <machine>[:<port>]
```

`<machine>[:<port>]` The IP address or DNS name of the machine on which you are shutting down the share server. If the share server is using a starting port other than 2099, the default, specify the starting port number.

avaki share --add-rehash-schedule

Add a rehash (refresh) schedule for an Avaki share.

Syntax

```
avaki share --add-rehash-schedule <share-grid-path>
  [{--periodic-schedule=<period-spec> |
  --cron-schedule=<cron-spec> |
  --one-time-schedule=<time-stamp>}]
  [--start-time=<time-stamp>] [--end-time=<time-stamp>] |
  [--max-iterations=<num-iterations>]]
  [--time-zone=<time-zone>] [--exclusion=<exclusion-name>]+
  [--exclusion-action={ drop | sooner | later }]
  [--reschedule-offset=<offset-spec>]
```

<share-grid-path> The path of the root share directory in the data catalog. For example:
/Shares/MyShare

--periodic-schedule=<period-spec> Specify a periodic schedule for running the rehash schedule. Use the following syntax:

```
--periodic-schedule="<number-of-units>
<unit-type>"
```

<number-of-units> An integer from 1 to 2147483648

<unit-type> Seconds, minutes, hours, days, weeks, or months

--cron-schedule=<cron-spec> Specify how often a cron-based scheduler for rehashing schedules will run the rehash schedule. Use the following syntax:

```
--cron-schedule="<seconds> <minutes> <hours>
<days-of-month> <months> <days-of-week> [<years>]"
```

For details about the **<cron-spec>** syntax, see [Appendix C, “Configuring cron schedules”](#).

<pre>--one-time- schedule= <time-stamp></pre>	<p>Specify a single time and date for running the rehash schedule. Use the following syntax:</p> <pre>--one-time-schedule="{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}"</pre>																				
<pre>--start-time= <time-stamp></pre>	<p>Specify the time and date that the rehash schedule will start running, in this format:</p> <pre>{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}</pre>																				
<pre>--end-time= <time-stamp></pre>	<p>Specify the time and date that the rehash schedule will stop running, in this format:</p> <pre>{<hh>:<mm>:<ss> {am pm} <mm>-<dd>-<yyyy> now}</pre>																				
<pre>--max- iterations= <num-iterations></pre>	<p>Specify the number of times a schedule will repeat.</p> <pre><num-iterations></pre> <p>An integer from 1 to 2147483648. If no value is specified, the schedule repeats indefinitely.</p>																				
<pre>--time-zone= <time-zone></pre>	<p>Specify the time zone for the schedule. The schedule can be specified according to the share server's time zone or relative to Greenwich Mean Time (GMT). If no value is specified for the time zone, the server's local time zone is used. The format for the time zone is as follows:</p> <table> <tr> <td>GMT+12:00 Fiji</td> <td>GMT+02:00 Cairo</td> </tr> <tr> <td>GMT+11:00 Noumea</td> <td>GMT+01:00 Paris</td> </tr> <tr> <td>GMT+10:00 Sydney</td> <td>GMT London</td> </tr> <tr> <td>GMT+9:30 Adelaide</td> <td>GMT-01:00 Azores</td> </tr> <tr> <td>GMT+09:00 Tokyo</td> <td>GMT-02:00 Mid-Atlantic</td> </tr> <tr> <td>GMT+08:00 Hong Kong</td> <td>GMT-03:00 Rio de Janeiro</td> </tr> <tr> <td>GMT+07:00 Bangkok</td> <td>GMT-03:30 Newfoundland</td> </tr> <tr> <td>GMT+06:30 Rangoon</td> <td>GMT-04:00 Caracas</td> </tr> <tr> <td>GMT+06:00 Dacca</td> <td>GMT-05:00 New York</td> </tr> <tr> <td>GMT+05:45 Katmandu</td> <td>GMT-06:00 Chicago</td> </tr> </table>	GMT+12:00 Fiji	GMT+02:00 Cairo	GMT+11:00 Noumea	GMT+01:00 Paris	GMT+10:00 Sydney	GMT London	GMT+9:30 Adelaide	GMT-01:00 Azores	GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic	GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro	GMT+07:00 Bangkok	GMT-03:30 Newfoundland	GMT+06:30 Rangoon	GMT-04:00 Caracas	GMT+06:00 Dacca	GMT-05:00 New York	GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+12:00 Fiji	GMT+02:00 Cairo																				
GMT+11:00 Noumea	GMT+01:00 Paris																				
GMT+10:00 Sydney	GMT London																				
GMT+9:30 Adelaide	GMT-01:00 Azores																				
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic																				
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro																				
GMT+07:00 Bangkok	GMT-03:30 Newfoundland																				
GMT+06:30 Rangoon	GMT-04:00 Caracas																				
GMT+06:00 Dacca	GMT-05:00 New York																				
GMT+05:45 Katmandu	GMT-06:00 Chicago																				

	GMT+05:30 Calcutta	GMT-07:00 Denver
	GMT+05:00 Karachi	GMT-08:00 San Francisco
	GMT+04:30 Kabul	GMT-09:00 Juneau
	GMT+04:00 Baku	GMT-10:00 Hawaii
	GMT+03:30 Tehran	GMT-11:00 Samoa
	GMT+03:00 Moscow	
<code>--exclusion= <exclusion- name></code>	The name of an exclusion to this schedule. For information about creating schedule exclusions, see avaki scheduleexclusion --create --daily (page 201), avaki scheduleexclusion --create --weekly (page 205), avaki scheduleexclusion --create --monthly (page 203), avaki scheduleexclusion --create --yearly (page 207), and avaki scheduleexclusion --create --custom (page 199).	
<code>--exclusion- action={drop sooner later}</code>	Specify the action to take if the rehash schedule is scheduled to run during an exclusionary period.	
	drop	Don't run the rehash schedule during the exclusionary period.
	sooner	Reschedule the rehash schedule to run before the exclusionary period.
	later	Reschedule the rehash schedule to run after the exclusionary period.
<code>--reschedule- offset= <offset-spec></code>	Specify the amount of time before or after the exclusionary period to rehash a schedule. Use the following syntax:	
	<code>--reschedule-offset="<number-of-units> <unit-type>"</code>	
	<code><number-of-units></code>	An integer from 1 to 2147483648
	<code><unit-type></code>	Seconds, minutes, hours, days, weeks, or months

avaki share --add-share-servers

Add one or more share servers to an Avaki share.

Syntax

```
avaki share --add-share-servers <share-grid-path>
  {--share-server=<share-server-name>[:<load-balance-factor>]}+
```

<share-grid-path>

The path of the root share directory in the grid. For example:

```
/Shares/MyShare
```

--share-server=
<share-server-name>
[:<load-balance-
factor>]+

Optional. The share servers for the share. The <share-server-name> name can be up to 32 characters long, and is specific to the grid server you are using. If you specify more than one name, you enable failover between the servers. For more information, see the *Sybase Avaki EII Administration Guide*.

If you configure multiple share servers, the <load-balance-factor> specifies the approximate proportion of the load carried by <share-server-name>. The proportion is a fraction, as follows:

- The numerator is the LBF, an integer that specifies the portion of load that a share server owns for a given Avaki share.
- The denominator is the sum of the LBFs for all the share servers serving this Avaki share.

avaki share --create

Create an Avaki share (that is, share a local directory into the data catalog).

By default, a share does not reread the contents of the shared directory from the local file system. For information about adding a rehash (refresh) schedule to an Avaki share directory, see [avaki share --rehash on page 241](#). For information about listing the rehash (refresh) events for an Avaki share, see [avaki share --list-rehash-schedules on page 240](#). For information about deleting a rehash (refresh) schedule, see [avaki share --delete-rehash-schedule on page 238](#).

Note When you create a share from a command client, Avaki looks for the <local-path> on the GDC machine if no share server or grid server is specified.

Syntax

```
avaki share --create <local-path> <share-grid-path>
  [--grid-server=<grid-server>]
  [--share-server=<share-server-name>[:<load-balance-factor>]] +
  [--encryption-level={clear | encrypted}]
  [--background | --bg] [--owner=<initial-owner>]
```

<code><local-path></code>	The path to the directory in the local file system that you are sharing into the data catalog. This path must resolve to a valid name on all share servers associates with the shares. We recommend that you provide an absolute path name to avoid confusion. The <share-server> name can be up to 32 characters long.
<code><share-grid-path></code>	Specify the parent directory and subdirectory for the share. The <share-grid-path> can be up to 255 characters long. A grid directory with this name is placed inside the parent directory you have specified. This subdirectory contains the shared data.
<code>--grid-server=<grid-server></code>	Optional. If you want the new share to use a particular grid server, specify the name of the grid server. For example, Bedrock.


```
--share-server=  
<share-server-name>  
[:<load-balance-  
factor>]+
```

Optional. The share servers for the share. The <share-server-name> name can be up to 32 characters long, and is specific to the grid server you are using. If you specify more than one name, you enable failover between the servers. For more information, see the *Sybase Avaki EII Administration Guide*.

If you configure multiple share servers, the <load-balance-factor> specifies the approximate proportion of the load carried by <share-server-name>. The proportion is a fraction, as follows:

- The numerator is the LBF, an integer that specifies the portion of load that a share server owns for a given Avaki share.
- The denominator is the sum of the LBFs for all the share servers serving this Avaki share.

```
--encryption-  
level={clear |  
encrypted}
```

Optional. Specify the encryption level for the share: clear (encryption is turned off for this share) or encrypted (encryption is turned on).

```
--background | --bg
```

Optional. The share is be built in the background. The command line returns immediately.

```
--owner=<initial-  
owner>
```

Optional. Specify a share's owner. By default, new shares are owned by the user creating the share.

avaki share --delete-rehash-schedule

Delete a rehash (refresh) schedule from an Avaki share.

Syntax

```
avaki share --delete-rehash-schedule <share-grid-path>
    <schedule-id>
```

<share-grid-path> The path to the shared directory in the local file system.

<schedule-id> The numeric ID of the schedule to delete. To determine a schedule's ID, use [avaki share --list-rehash-schedules](#) (page 240).

avaki share --disconnect

Permanently disconnect and shut down an Avaki share and remove any links to that share.

Caution When you disconnect a shared directory from a grid domain, the change is not reflected in the source file system; any related source files remain on the source file system.

Syntax

```
avaki share --disconnect <share-grid-path>
```

<share-grid-path> The path of the root share directory in the grid. For example:
 /Shares/MyShare

avaki share --get-local-path

Get the local path for an Avaki share.

Syntax

```
avaki share --get-local-path <share-grid-path>
```

<share-grid-path> The path of the root share directory in the data catalog. For example:
 /Shares/MyShare

avaki share --get-status

Get the current status for an Avaki share: online or offline.

Syntax

```
avaki share --get-status <share-grid-path>
```

<share-grid-path> The path of the root share directory in the data catalog. For example:
 /Shares/MyShare

avaki share --list-rehash-schedules

List the rehash (refresh) events for an Avaki share.

Syntax

```
avaki --list-rehash-schedules <share-grid-path>
```

<code><share-grid-path></code>	The path of the root share directory in the data catalog. For example: <code>/Shares/MyShare</code>
--------------------------------------	--

avaki share --list-share-servers

List the share servers for an Avaki share.

Syntax

```
avaki share --list-share-servers <share-grid-path>
```

<code><share-grid-path></code>	The path of the root share directory in the data catalog. For example: <code>/Shares/MyShare</code>
--------------------------------------	--

avaki share --rehash

Rehash (refresh) an Avaki share.

Syntax

```
avaki share --rehash <share-grid-path> [--background | --bg]
```

<share-grid-path> The path of the root share directory in the data catalog. For example:
/Shares/MyShare

--background | --bg Optional. The share will be rehashed in the background. The command line returns immediately.

avaki share --remove-share-servers

Remove a share server from an Avaki share.

Syntax

```
avaki share --remove-share-servers <share-grid-path>
  {--share-server=<share-server-name>}+
```

<share-grid-path> The path of the root share directory in the data catalog. For example:
/Shares/MyShare

--share-server=share-server-name The logical name of one or more share servers to remove.

avaki share --set-local-path

Set the local path for an Avaki share.

Syntax

```
avaki --set-local-path <share-grid-path> <local-path>
```

`<share-grid-path>` The path of the root share directory in the data catalog. For example:
 /Shares/MyShare

`local-path` The local path for the shared directory.

Example

```
$ avaki --set-local-path /Shares/MyAvakiShare c:\Bedrock
```

avaki share --set-share-servers

Replace the set of share servers for an Avaki share.

Syntax

```
avaki --set-share-servers <share-grid-path>
  {--share-server=<share-server-name>[:<load-balance-factor>]}+
```

<share-grid-path>

The path of the root share directory in the data catalog. For example:

```
/Shares/MyShare
```

--share-server=
<share-server-name>
[:<load-balance-
factor>]+

The names of one or more share servers that will replace the set of share servers for <share-grid-path>. The <share-server> names can be up to 32 characters long.

If you configure multiple share servers, the <load-balance-factor> specifies the approximate proportion of the load carried by <share-server-name>. The proportion is a fraction, as follows:

- The numerator is the LBF, an integer that specifies the portion of load that a share server owns for a given Avaki share.
- The denominator is the sum of the LBFs for all the share servers serving this Avaki share.

For more information, see the *Sybase Avaki EII Administration Guide*.

avaki share --set-status

Setting the status for an Avaki share: online or offline.

Syntax

```
avaki --set-status <share-grid-path> {online | offline}
```

<share-grid-path> The path of the root share directory in the grid. For example:
 /Shares/MyShare

online | The new status for a shared directory: online or offline.
offline

avaki share --update-share-servers

Modify the load balance factor for a set of share servers.

Syntax

```
avaki --update-share-servers <share-grid-path>
  {--share-server=<share-server-name>:<load-balance-factor>}+
```

<share-grid-path>

The path of the root share directory in the data catalog. For example:

```
/Shares/MyShare
```

--share-server=
<share-server-name>
[:<load-balance-factor>]+

The names of one or more share servers that will replace the set of share servers for <share-grid-path>. The <share-server> names can be up to 32 characters long.

If you configure multiple share servers, the <load-balance-factor> specifies the approximate proportion of the load carried by <share-server-name>. The proportion is a fraction, as follows:

- The numerator is the LBF, an integer that specifies the portion of load that a share server owns for a given Avaki share.
- The denominator is the sum of the LBFs for all the share servers serving this Avaki share.

For more information, see the *Sybase Avaki EII Administration Guide*.

avaki shell

Open an Avaki shell. You can use an Avaki shell to issue troubleshooting commands or to work in an Avaki-only command environment. In the shell, do not enter the **avaki** part of Avaki commands. For example, for **avaki ls** you enter just **ls**.

Note To exit an Avaki shell, type **quit**.

Syntax

```
avaki shell [--debug | -d] [--perf | -p]
```

`--debug | -d` Optional. Toggles debug mode on/off. When debug mode is enabled, you can use it to troubleshoot problems.

`--perf | -p` Optional. Toggles performance tracking on/off. When performance tracking is enabled, you can view how long it takes to run a command that you issue within the shell.

avaki sqlview --delete

Delete a SQL view.

Syntax

```
sqlview --delete --type={dataservice | dbop | provisioned |
  virtual_dbop} <qualified-table-name>
```

`--type=`
`{dataservice | dbop` The name of the data service, database operation, provi-
`| provisioned |` sioned table, or virtual database operation from which the
`virtual_dbop}` SQL view will be deleted. The format for `<quali-`
`<qualified-table-name>` is:

Data service [`<domain-name>`].DATASERVICE.
`<table-name>`

Database operation	[<domain-name>.]<database-connector-name>.<table-name>
Provisioned table	[<domain-name>.]<database-connector-name>.<table-name>
Virtual database operation	[<domain-name>.]VIRTUALDB.<table-name>

avaki sqlview --get-description

Display a description for a SQL view.

Syntax

```
avaki sqlview --get-description --type={dataservice | dbop |
    provisioned | virtual_dbop} <qualified-table-name>
```

--type=
{dataservice | dbop
| provisioned |
virtual_dbop}

The name of the data service, database operation, provisioned table, or virtual database operation for which the SQL view's description will be displayed. The format for <qualified-table-name> is:

Data service	[<domain-name>.]DATASERVICE.<table-name>
Database operation	[<domain-name>.]<database-connector-name>.<table-name>
Provisioned table	[<domain-name>.]<database-connector-name>.<table-name>
Virtual database operation	[<domain-name>.]VIRTUALDB.<table-name>

avaki sqlview --set-description

Modify the description of a SQL view.

Syntax

```
avaki sqlview --set-description --type={dataservice | dbop |
    provisioned | virtual_dbop} <qualified-table-name>
    <description>
```

--type=

```
{dataservice | dbop
 | provisioned |
virtual_dbop}
```

The name of the data service, database operation, provisioned table, or virtual database operation for which the SQL view's description will be modified. The format for <qualified-table-name> is:

```
[<domain-name>.]<database-connector-name>
<database-operation-name>
```

Data service [<domain-name>.]DATASERVICE.
 <table-name>

Database operation [<domain-name>.]<database-
 connector-name>.<table-name>

Provisioned table [<domain-name>.]<database-
 connector-name>.<table-name>

Virtual database operation [<domain-name>.]VIRTUALDB.
 <table-name>

<description>

Specify a description for the SQL view.

avaki status

Display the status of operations currently running on a grid server.

Syntax

```
avaki status --server=<grid-server> [--xml] [<opType>+]
```

`--grid-server=<grid-server>` Optional. Specify the name of the grid server for which the status will be printed. For example, Bedrock.

`--xml` Optional. Print the output in XML format.

`<opType>+` Optional. One or more operation types for which you want to obtain the status.

DataService Data service

DatabaseOperation Database operation

HttpOperation HTTP operation

avaki upgrade

Upgrade Avaki software on the specified grid servers and share servers.

Note See the *Sybase Avaki EII Administration Guide* for instructions on performing upgrades.

Syntax

```
avaki upgrade {--all | --server=<server-name>}
  [--version=<version>]
```

--all | --server=<server-name> Upgrade the specified grid servers and share servers. Enter **avaki upgrade --all** to upgrade all grid servers and share servers in this Avaki domain, or specify the name of a single grid server or share server.

--version=<version> Optional. Specify the version of Avaki software to be upgraded to. The default is the current version of the system code.

avaki upgrade --info

Display the upgrade version for the specified grid servers or share servers.

Syntax

```
avaki upgrade --info [--all | --server=<server-name>]
```

--all | --server=<server-name> List upgrade information for the specified grid servers and share servers. Enter **avaki upgrade --info --all** to display upgrade information on all grid servers and share servers in this Avaki domain, or specify the name of a single grid server or share server.

avaki user

Enable or disable a grid user account on an LDAP or NIS authentication service.

To enable or disable a user account, you must be a member of the Administrators group or the UserAdministrators group.

Syntax

```
avaki user {<user-grid-path> | <qualified-user-name>}
  --enabled={true | false}>
```

<user-grid-
path> |
<qualified-
user-name>

Specify the name of the user that you are enabling on the authentication service. You can specify the path to the data catalog directory of the user account. For example:

```
/System/Domains/Bedrock/Services/AuthServices/
Grid/DefaultAuthService/Users/wilma
```

Alternatively, you can specify the qualified user name using the following format:

```
<user>@<authservice>.<authservicetype>.<domain>
```

For example:

```
fred@slaterock.nis.bedrock
```

The elements of <qualified-user-name> are as follows:

<user> The name of the user for which you want to obtain the user's name and contact information.

<authservice> The name of the user's authentication service.

<authservicetype> The type of authentication service (Grid, Nis, or Ldap).

<domain> The name of the user's Avaki domain.

--enabled={true | false}> Indicate whether the user account will be enabled (true) or disabled (false) on the authentication service.

avaki user --create

Create a grid user account.

To use this command, you must be a member of the Administrators group or the User-Administrators group, or you must have permission to create user accounts.

Syntax

```
avaki user --create [--real-name=<real-name>] [--email=<email>]  
    <grid-user> <password>
```

--real-name=
<real-name> Optional. The full name of the user.

--email=<email> Optional. The e-mail address of the user.

<grid-user> The name of this user account.

<password> The password for this user account. If you're executing this command on a Unix system and the password contains non-alphanumeric characters (! or &, for instance), enclose the password in single quotation marks to prevent the shell from interpreting the special characters.

avaki user --db-mapping --add

Add a database identity mapping for this user. A database identity mapping overrides the user's default user name and password.

To use this command, you must be a member of the Administrators group or the User-Administrators group.

Syntax

```
avaki user --db-mapping --add {<user-grid-path> |
  <qualified-user-name>} <qualified-dbconn-name> <db-username>
  <db-password>
```

<user-grid-
path> |
<qualified-
user-name>

Specify the name of the user for which you are adding an identity mapping. You can specify the data catalog path to the directory of the user account. For example:

```
/System/Domains/Bedrock/Services/AuthServices/
Grid/DefaultAuthService/Users/wilma
```

Alternatively, you can specify the qualified user name using the following format:

```
<user>@<authservice>.<authservicetype>.<domain>
```

For example:

```
fred@slaterock.nis.bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user for which you want to obtain the user's name and contact information.
<authservice>	The name of the user's authentication service.
<authservice- type>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

<qualified-
dbconn-name>

A database connector on which per-user database identity mappings have been enabled. Before configuring a mapping for a user, enable database identity mappings on the database connector that connects to your database by creating an attribute for the database connector with the following parameters:

Name	dbconn/UseIdentityMapping
Type	string
Value	true

For information about creating new attributes, see [avaki attribute --update \(page 21\)](#).

The format for <qualified-dbconn-name> is:

```
[<domain-name>.]<database-connector-name>
```

For example:

```
Bedrock.sybase
```

The elements of <qualified-dbconn-name> are as follows:

<domain-name> Optional. The grid domain where the database connector is deployed. The default is the current domain.

<database-connector-name> The name of the database connector.

<db-username> The database user name to map to <user-grid-path> or <qualified-user-name>.

<db-password> The database password to map to <user-grid-path> or <qualified-user-name>.

avaki user --db-mapping --delete

Delete a database identity mapping for a user.

To use this command, you must be a member of the Administrators group or the User-Administrators group.

Syntax

```
avaki user --db-mapping --delete {<user-grid-path> |
  <qualified-user-name>} <qualified-dbconn-name>
```

<user-grid-
path> |
<qualified-
user-name>

Specify the name of the user for which you are deleting the identity mapping. You can specify the data catalog path to the directory of the user account. For example:

```
/System/Domains/Bedrock/Services/AuthServices/  
Grid/DefaultAuthService/Users/  
wilma
```

Alternatively, you can specify the qualified user name using the following format:

```
<user>@<authservice>.<authservicetype>.<domain>
```

For example:

```
fred@slaterock.nis.bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user whose database identity mapping you want to delete.
<authservice>	The name of the user's authentication service.
<authservice- type>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

<qualified-
dbconn-name>

A database connector on which per-user identity mappings have been enabled. The format is:

[<domain-name>.]<database-connector-name>

For example:

Bedrock.sybase

The elements of <qualified-dbconn-name> are as follows:

<domain-name> Optional. The grid domain where the database connector is deployed. The default is the current domain.

<database-
connector-
name> The name of the database connector.

avaki user --db-mapping --list

List the database identity mapping for a user.

To use this command, you must be a member of the Administrators group or the User-Administrators group.

Syntax

```
avaki user --db-mapping --list {<user-grid-path> |
    <qualified-user-name>}
```

<user-grid-
path> |
<qualified-
user-name>

Specify the name of the user for which you are listing the identity mappings. You can specify the data catalog path to the directory of the user account. For example:

```
/System/Domains/Bedrock/Services/AuthServices/  
Grid/DefaultAuthService/Users/wilma
```

Alternatively, you can specify the qualified user name by using the following format:

```
<user>@<authservice>.<authservicetype>.<domain>
```

For example:

```
fred@slaterock.nis.bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user whose identity mapping you want to obtain.
<authservice>	The name of the user's authentication service.
<authservice- type>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

avaki user --delete

Delete a user account from a grid domain.

To use this command, you must be a member of the Administrators or UserAdministrators group, or you must have permission to delete user accounts.

Syntax

```
avaki user {--delete | -d} {<users> | {--all | -a}}
```

<code><users> {--all -a}</code>	The name of the user account to delete. To delete all users, specify <code>--all</code> or <code>-a</code> . To specify one or more users, enter a space-delimited list of users.
---	---

avaki user --info

Display a user's name and contact information.

Syntax

```
avaki user --info <qualified-user-name>
```

<code><qualified-user-name></code>	Specify the qualified name of the user account for which you want to obtain a full name and contact information. The format is: <code><user>@<authservice>.<authservicetype>.<domain></code>
--	---

For example:

```
fred@slaterock.nis.bedrock
```

The elements of `<qualified-user-name>` are as follows:

<code><user></code>	The name of the user whose contact information you want to obtain.
<code><authservice></code>	The name of the user's authentication service.

<authservice-type>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

avaki user --list-group

List all the groups to which a user belongs.

Syntax

```
avaki user {--list-group | -l} <qualified-user-name>
```

<user-grid-path> |
<qualified-user-name>

Specify the qualified name of the user for which you want to obtain a group membership list. The format is:

```
<user>@<authservice>.<authservicetype>.<domain>
```

For example:

```
fred@slaterock.nis.bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user for which you want to obtain a group membership list.
<authservice>	The name of the user's authentication service.
<authservice-type>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

avaki view --add-schedule

Add a generation schedule to a view.

To use this command, you must have write permission on the view.

Syntax

```
avaki view --add-schedule <view-path>
  {--periodic-schedule=<period-spec> |
  --cron-schedule=<cron-spec> |
  --one-time-schedule=<time-stamp>}
  [--start-time=<time-stamp>] [[--end-time=<time-stamp>] |
  [--max-iterations=<num-iterations>]]
  [--time-zone=<time-zone>] [--exclusion=<exclusion-name>]+
  [--exclusion-action={ drop | sooner | later }]
  [--reschedule-offset=<offset-spec>]
```

<view-path> The full grid path to the generated view (that is, the output file in which the view data is stored).

--periodic-schedule=<period-spec> Specify a periodic schedule for generating the view. Use the following syntax:

```
--periodic-schedule="<number-of-units> <unit-type>"
```

<number-of-units> An integer from 1 to 2147483648

<unit-type> seconds, minutes, hours, days, weeks, or months

In the following example, the periodic schedule is set to 30 seconds:

```
--periodic-schedule="30 seconds"
```


<pre>--cron-schedule= <cron-spec></pre>	<p>Specify how often a cron-based scheduler for generating views generates the view. Use the following syntax:</p> <pre>--cron-schedule="<seconds> <minutes> <hours> <days-of-month> <months> <days-of-week> [<years>]"</pre> <p>For details about the <cron-spec> syntax, see Appendix C, “Configuring cron schedules”.</p>						
<pre>--one-time- schedule= <time-stamp></pre>	<p>Specify a single time and date for generating the view. Use the following syntax:</p> <pre>--one-time-schedule="{<hh>:<mm>:<ss> {am pm} <mm>--<dd>--<yyyy> now}"</pre>						
<pre>--start-time= <time-stamp></pre>	<p>Specify the time and date that the view generator will start running, in this format:</p> <pre>{<hh>:<mm>:<ss> {am pm} <mm>--<dd>--<yyyy> now}</pre>						
<pre>--end-time= <time-stamp></pre>	<p>Specify the time and date that the view generator will stop running, in this format:</p> <pre>{<hh>:<mm>:<ss> {am pm} <mm>--<dd>--<yyyy> now}</pre>						
<pre>--max-iterations= <num-iterations></pre>	<p>Specify the number of times a schedule will repeat.</p> <pre><num-iterations></pre> <p>An integer from 1 to 2147483648. If no value is specified, the schedule repeats indefinitely.</p>						
<pre>--time-zone= <time-zone></pre>	<p>Specify the time zone for the schedule. The schedule can be specified according to the grid server’s time zone or relative to Greenwich Mean Time (GMT). If no value is specified for the time zone, the server’s local time zone is used. The format for the time zone is as follows:</p> <table border="0" style="margin-left: 40px;"> <tr> <td>GMT+12:00 Fiji</td> <td>GMT+02:00 Cairo</td> </tr> <tr> <td>GMT+11:00 Noumea</td> <td>GMT+01:00 Paris</td> </tr> <tr> <td>GMT+10:00 Sydney</td> <td>GMT London</td> </tr> </table>	GMT+12:00 Fiji	GMT+02:00 Cairo	GMT+11:00 Noumea	GMT+01:00 Paris	GMT+10:00 Sydney	GMT London
GMT+12:00 Fiji	GMT+02:00 Cairo						
GMT+11:00 Noumea	GMT+01:00 Paris						
GMT+10:00 Sydney	GMT London						

GMT+9:30 Adelaide	GMT-01:00 Azores
GMT+09:00 Tokyo	GMT-02:00 Mid-Atlantic
GMT+08:00 Hong Kong	GMT-03:00 Rio de Janeiro
GMT+07:00 Bangkok	GMT-03:30 Newfoundland
GMT+06:30 Rangoon	GMT-04:00 Caracas
GMT+06:00 Dacca	GMT-05:00 New York
GMT+05:45 Katmandu	GMT-06:00 Chicago
GMT+05:30 Calcutta	GMT-07:00 Denver
GMT+05:00 Karachi	GMT-08:00 San Francisco
GMT+04:30 Kabul	GMT-09:00 Juneau
GMT+04:00 Baku	GMT-10:00 Hawaii
GMT+03:30 Tehran	GMT-11:00 Samoa
GMT+03:00 Moscow	

```
--exclusion=  
<exclusion-name>
```

The name of an exclusion to this schedule. For information about creating schedule exclusions, see [avaki scheduleexclusion --create --daily](#) (page 201), [avaki scheduleexclusion --create --weekly](#) (page 205), [avaki scheduleexclusion --create --monthly](#) (page 203), [avaki scheduleexclusion --create --yearly](#) (page 207), and [avaki scheduleexclusion --create --custom](#) (page 199).

```
--exclusion-  
action={drop |  
sooner | later}
```

Specify the action to take if the view generator is scheduled to run during an exclusionary period.

drop	Don't run the data service during the exclusionary period.
sooner	Reschedule the data service to run before the exclusionary period.
later	Reschedule the data service to run after the exclusionary period.

<code>--reschedule- offset= <offset-spec></code>	Specify the amount of time before or after the exclusionary period to run the view generator. Use the following syntax: <code>--reschedule-offset="<number-of-units> <unit-type>"</code> <code><number-of-units></code> An integer from 1 to 2147483648 <code><unit-type></code> Seconds, minutes, hours, days, weeks, or months
--	---

Example

This example adds a schedule in which the view `pebblesmilestones.txt` will be generated every 2 hours. In this example, the time zone is GMT-05:00 New York.

```
$ avaki view --add-schedule /GeneratedViews/  
pebblesmilestonesview.txt --periodic-schedule="2 hours"  
--time-zone="GMT-05:00 New York"
```

avaki view --create --database

Create a view generator that uses a database operation as its input source, then generate the view. To configure view properties such as a description that can be viewed by other users, use [avaki view --set-property](#) (page 275).

To use this command, you must be a member of the DataProviders group.

Syntax

```
avaki view --create --database <view-path>
  <qualified-dbop-name> [[<parameter-value>+]
  [--html | --csv | --stylesheet=<stylesheet-path>]
  [--engine={saxon | xalan | <class-name>}]]
  [--server=<grid-path>] [--no-generate]
  --run-user={<qualified-user-name>}
```

<view-path> The full grid path to the generated view (that is, the output file in which the view data will be stored).

<qualified-dbop-name> The database operation to use as the input source for the view. The format for the qualified name is:

```
[<domain-name>.]<database-connector-name>.<database-operation-name>
```

For example:

```
Bedrock.sybase.listsalaries
```

The elements of **<qualified-dbop-name>** are as follows:

<domain-name> Optional. The grid domain where the database connector is deployed. The default is the current domain.

<database-connector-name> The name of the database connector that **<database-operation-name>** uses.

<dbop-name> The name of a database operation.

<code><parameter-value>+</code>	Optional. One or more parameter values required by the target database operation. The parameter values will be processed in the order specified in the database operation's query. Any parameter value that includes a space must be enclosed in quotation marks. Any string at the end of the command that is enclosed in quotation marks will be interpreted as a single value. (That is, enclose each parameter value in its own pair of quotation marks.)
<code>--html</code>	Optional. Produce a generated view in HTML format. (Use this option only when the input source is Avaki rowset XML, usually.)
<code>--csv</code>	Optional. Produce a generated view in comma-separated values format. (Use this option only when the input source is Avaki rowset XML, usually.)
<code>--stylesheet= <stylesheet-path></code>	<p>The grid path to a style sheet. Required if you specify a custom Java transform engine in the <code>--engine</code> option.</p> <p>If the input source is in XML, you can run an XSLT style sheet to convert the data to another XML schema, HTML, PDF, plain text, etc. In this case, use the <code>--engine</code> option to specify Saxon or Xalan, and use <code>--stylesheet</code> to provide the grid path to an XSLT style sheet.</p> <p>If you are using a custom Java transform engine, use <code>--stylesheet</code> to provide the grid path to a style sheet file that provides any needed information. (You must provide a style sheet file even if your transformation doesn't use it.)</p>
<code>--engine={saxon xalan <class-name>}</code>	<p>Optional. Specify the style sheet engine to use when generating this view. You can choose one of the two engines provided by Avaki, Saxon or Xalan, or you can write a custom transform engine in Java. If you write a custom engine, specify its <code>TransformerFactoryImpl</code> class name. You can find the name in your XSLT processor's JAR files. Here are two common ones:</p> <ul style="list-style-type: none"> • For Saxon 7: <code>net.sf.saxon.TransformerFactoryImpl</code> • For <code>jd.xml</code>: <code>jd.xml.xslt.trax.TransformerFactoryImpl</code> <p>Default: <code>saxon</code></p>

<code>--server= <grid-path></code>	Optional. Specify the name of the Avaki server on which the view will be created. For example, Bedrock. If you omit this option, the view is created on the current server.
<code>--no-generate</code>	Optional. Create a view generator but do not generate a view. (You can use avaki view --generate on page 274 to generate the view later.)
<code>--run-user= {<qualified-user- name>}</code>	Specify the user to create the view as. Specify the qualified user name by using the following format: <code><user>@<authservice>.<authservice- type>.<domain></code> For example: <code>wilma@DefaultAuthService.Grid.Bedrock</code> The elements of <code><qualified-user-name></code> are as follows:
<code><user></code>	The name of the user to create the view as.
<code><authservice></code>	The name of the user's authentication service.
<code><authservice- type></code>	The type of authentication service (Grid, Nis, or Ldap).
<code><domain></code>	The name of the user's Avaki domain.

Example

```
$ avaki view --create --database /gridhome/fred/views/salaries
Bedrock.sybase.listsalaries
```

avaki view --create --data-service

Create a view generator that uses a data service as its input source, then generate the view. To configure view properties such as a description that can be viewed by other users, use [avaki view --set-property](#) (page 275).

To use this command, you must be a member of the DataProviders group.

Syntax

```
avaki view --create --data-service <view-path>
  [<qualified-ds-name> [<parameter-value>+] [--html | --csv |
  --stylesheet=<stylesheet-path>] [--engine={saxon | xalan |
  <class-name>}] [--server=<grid-path>] [--no-generate]
  --run-user={<qualified-user-name>}
```

<code><view-path></code>	The full grid path to the generated view (that is, the output file in which the view data will be stored).
<code><qualified-ds-name></code>	The data service to use as the input source for the view. The format for the qualified name is: [<domain-name>.]<ds-name>
<code><parameter-value>+</code>	Optional. One or more parameter values required by the target data service operation. The parameter values will be processed in the order specified in the database service's query. Any parameter value that includes a space must be enclosed in quotation marks. Any string at the end of the command that is enclosed in quotation marks will be interpreted as a single value. (That is, enclose each parameter value in its own pair of quotation marks.)
<code>--html</code>	Optional. Produce a generated view in HTML format. (Use this option only when the input source is Avaki rowset XML, usually.)
<code>--csv</code>	Optional. Produce a generated view in comma-separated values format. (Use this option only when the input source is Avaki rowset XML, usually.)

<code>--stylesheet= <stylesheet- path></code>	<p>The grid path to a style sheet. Required if you specify a custom Java transform engine in the <code>--engine</code> option.</p> <p>If the input source is in XML, you can run an XSLT style sheet to convert the data to another XML schema, HTML, PDF, plain text, etc. In this case, use the <code>--engine</code> option to specify Saxon or Xalan, and use <code>--stylesheet</code> to provide the grid path to an XSLT style sheet.</p> <p>If you are using a custom Java transform engine, use <code>--stylesheet</code> to provide the grid path to a style sheet file that provides any needed information. (You must provide a style sheet file even if your transformation doesn't use it.)</p>
<code>--engine={saxon xalan <class-name>}</code>	<p>Optional. Specify the style sheet engine to use when generating this view. You can choose one of the two engines provided by Avaki, Saxon or Xalan, or you can write a custom transform engine in Java. If you write a custom engine, specify its <code>TransformerFactoryImpl</code> class name. You can find the name in your XSLT processor's JAR files. Here are two common ones:</p> <ul style="list-style-type: none"> • For Saxon 7: <code>net.sf.saxon.TransformerFactoryImpl</code> • For <code>jd.xslt</code>: <code>jd.xml.xslt.trax.TransformerFactoryImpl</code> <p>Default: <code>saxon</code></p>
<code>--server= <grid-path></code>	<p>Optional. Specify the name of the Avaki server on which the view will be created. For example, <code>Bedrock</code>. If you omit this option, the view is created on the current server.</p>
<code>--no-generate</code>	<p>Optional. Create a view generator but do not generate a view. (You can use avaki view --generate on page 274 to generate the view later.)</p>


```
--run-user=
{<qualified-user-
name>}
```

Specify the user to create the view as. Specify the qualified user name by using the following format:

```
<user>@<authservice>.<authservice-
type>.<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user to create the view as.
<authser- vice>	The name of the user's authentication service.
<authser- vicetype>	The type of the user's authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

avaki view --create --file

Create a view generator that uses a file or a generated view as its input source, then generate the view. To configure view properties such as a description that can be viewed by other users, use [avaki view --set-property](#) (page 275).

To use this command, you must be a member of the DataProviders group.

Syntax

```
avaki view --create --file <view-path> <file-path>
  [--stylesheet=<stylesheet-path>] [--engine={saxon | xalan |
  <class-name>}] [--server=<grid-path>] [--no-generate]
  --run-user={<qualified-user-name>}
```

<code><view-path></code>	The grid path to the output file in which the generated view will be stored.
<code><file-path></code>	The grid path to a file or generated view. This is the input source for the view generator.
<code>--stylesheet= <stylesheet- path></code>	<p>The grid path to a style sheet. Required if you specify a custom Java transform engine in the <code>--engine</code> option.</p> <p>If the input source is in XML, you can run an XSLT style sheet to convert the data to another XML schema, HTML, PDF, plain text, etc. In this case, use the <code>--engine</code> option to specify Saxon or Xalan, and use <code>--stylesheet</code> to provide the grid path to an XSLT style sheet.</p> <p>If you are using a custom Java transform engine, use <code>--stylesheet</code> to provide the grid path to a style sheet file that provides any needed information. (You must provide a style sheet file even if your transformation doesn't use it.)</p>

<pre>--engine={saxon xalan <class-name>}</pre>	<p>Optional. Specify the style sheet engine to use when generating this view. You can choose one of the two engines provided by Avaki, Saxon or Xalan, or you can write a custom transform engine in Java. If you write a custom engine, specify its TransformerFactoryImpl class name. You can find the name in your XSLT processor's JAR files. Here are two common ones:</p> <ul style="list-style-type: none"> • For Saxon 7: net.sf.saxon.TransformerFactoryImpl • For jd.xml: jd.xml.xslt.trax.TransformerFactoryImpl <p>Default: saxon</p>
<pre>--server= <grid-path></pre>	<p>Optional. Specify the name of the server on which the view will be created. For example, Bedrock.</p>
<pre>--no-generate</pre>	<p>Optional. Create a view generator but do not generate a view. (You can use avaki view --generate (page 274) to generate the view later.)</p>
<pre>--run-user= {<qualified-user- name>}</pre>	<p>Specify the user to create the view as. Specify the qualified user name by using the following format:</p> <pre><user>@<authservice>.<authservice- type>.<domain></pre> <p>For example:</p> <pre>wilma@DefaultAuthService.Grid.Bedrock</pre> <p>The elements of <qualified-user-name> are as follows:</p>
<pre><user></pre>	<p>The name of the user to create the view as.</p>
<pre><authser- vice></pre>	<p>The name of the user's authentication service.</p>
<pre><authser- vicetype></pre>	<p>The type of authentication service (Grid, Nis, or Ldap).</p>
<pre><domain></pre>	<p>The name of the user's Avaki domain.</p>

Example

```
$ avaki view --create --file
/GeneratedViews/pebblesmilestonesview.txt
"/Shares/Water Buffalo Lodge/pebblesmilestones.txt"
--run-user=Wilma@DefaultAuthService.Grid.training
```

avaki view --delete

Delete a view.

To use this command, you must be a member of the Administrators group, or you must have delete permission on the view.

Syntax

```
avaki view --delete <view-path>
```

<view-path> The grid path to the generated view file.

avaki view --delete-schedule

Delete a generation schedule from a view.

To use this command, you must have write permission on the view.

Syntax

```
avaki view --delete-schedule <view-path> <schedule-id>
```

<view-path> The grid path to the generated view file.

<schedule-id> The numeric ID of the schedule to delete. To determine a schedule's ID, use [avaki view --list-schedules](#) on page 275.

avaki view --depends

Show the dependencies to and from a view.

To use this command, you must have read permission on the message service on the grid server where the view is located, and you must have read permissions on the view.

Syntax

```
avaki view --depends <view-path>
```

<view-path> The grid path to a generated view or view generator.

avaki view --garbage-collect

Force the garbage collector to run immediately on the specified view. The garbage collector frees up disk space by removing old results for the target view.

To use this command, you must be a member of the DataProviders group.

Syntax

```
avaki view --garbage-collect <view-path>
```

<view-path> The grid path to the generated view or the view generator.

avaki view --generate

Regenerate a view.

To use this command, you must be a member of the DataProviders group.

Syntax

```
avaki view --generate <view-path>
```

<view-path> The grid path to the generated view (if it already exists) or the view generator.

avaki view --info

Display a list of the properties of a generated view and their values. For descriptions of the properties that can be used to configure a view, see [avaki view --set-property on page 275](#).

To use this command, you must be a member of the DataProviders group.

Syntax

```
avaki view --info <view-path>
```

<view-path> The grid path to the generated view or the view generator.

avaki view --list-schedules

List the generation schedules for a view.

Syntax

```
avaki view --list-schedules <view-path>
```

<view-path> The grid path to the generated view file.

avaki view --set-property

Configure the values for the properties that are used when a view is generated. Properties you can set include a description of the view, the style sheet for the view, parameters for the style sheet, and the schedule for the view. See the “[Syntax](#)” section, below, for details. To display current values of the properties, use [avaki view --info](#) (page 274).

To use this command, you must be a member of the DataProviders group.

Syntax

```
avaki view --set-property <view-path> <property-name>
    <property-value>*
```

<view-path> The grid path to the generated view or the view generator.

<property-name> The name of the property to configure.

The property names are as follows:

General Properties

You can configure any of the following parameters for a file view generator, a data service view generator, or a database view generator:

description	Explain what this view is for, who owns it, when it is allowed to run, or other details that will be helpful to those who use it. You can display the description using avaki view --info (page 274).
gc-interval	The time in seconds between garbage collection cycles on expired generated views. Default: 1800 (half an hour) Range: 0 to 2147483647 seconds
gc-timeout	A timer (in seconds) that starts at the last access of an old generated view. When the timeout value is reached, the garbage collector deletes the old generated view. (A generated view is considered old when a newer version of the same view exists.) Set this property to 0 to cause old views to expire immediately upon the appearance of a newer version. Default: 3600 (1 hour) Range: 0 to 2147483647 seconds
gen-view-path	(Read-only) The grid directory in which the generated view will be stored.
notification-list	The grid path to one or more views that this view depends upon. When a specified view changes, the dependent view is regenerated. To specify more than one view, enter a space-delimited list of views.

output.xslt-engine	<p>The class name of the TrAX transformer factory for the style sheet engine you want to use. Avaki includes two style sheet engines:</p> <ul style="list-style-type: none"> • com.icl.saxon.TransformerFactoryImpl • org.apache.xalan.trax.TransformerFactoryImpl <p>Default: com.icl.saxon.TransformerFactoryImpl</p>
output.xslt-path	<p>The full grid path to the style sheet that will be used to format the generated view.</p>
output.xslt-params	<p>One or more colon-separated <name>:<value> pairs, where <name> is a stylesheet property and <value> is its value. The name:value pairs are passed to the stylesheet. If you enter more than one name:value pair, separate the pairs with spaces.</p>
schedule	<p>The schedule or schedules under which this view will automatically run. Syntax:</p> <pre>"every <interval> <units> starting {<timestamp> now}"</pre> <p><interval>: an integer from 0 to 2147483648 <units>: days, hours, or minutes <timestamp>: A date and time, in this format: yyyy-mm-dd hh:mm:ss {am pm}</p> <p>To specify more than one schedule, enter a space-delimited list of schedules.</p> <p>To turn off the schedule, set the schedule to an empty string.</p> <p>See “Examples” on page 279.</p>
server-path	<p>(Read-only) The grid path to the Avaki server on which the generated view will be created.</p>

Database View Generator Properties

You can configure any of the following parameters for a database view generator:

input.dbop	The name of the database operation to use as the input source for a database view generator.
input.input-file	The name of a file that contains data to load into the database. For information on the XML format required for input files, see the <i>Sybase Avaki EII Administration Guide</i> .
input.params	One or more parameters for a SQL query. To specify more than one parameter, enter a space-delimited list of parameters.
input.xml-format	Specify the XML schema for the view generator's input data. Default: default Valid values: default (Avaki's rowset format), oracle_xsu

File View Generator Properties

You can configure the following parameter for a file view generator:

input.file-path	The full grid path to the file to use as the input source for a file view generator.
-----------------	--

<property-value>*

Specify a value for the property you are configuring. Property values are optional only for properties that take lists (multiple values).

Examples

This command makes the view /GeneratedViews/Pebbles dependent on the view /GeneratedViews/Dino so that /GeneratedViews/Pebbles is regenerated when /GeneratedViews/Dino changes.

```
$ avaki view --set-property /GeneratedViews/Pebbles  
notification-list /GeneratedViews/Dino
```

This command changes the SQL parameters used for the database view SalesYTD to 2005.

```
$ avaki view --set-property /GeneratedViews/SalesYTD  
input.params 2005
```

This command sets a view to run every 4 days starting at 1:59:59 PM on Sept. 3, 2005.

```
$ avaki view --set-property /GeneratedViews/Dino schedule  
"every 4 days starting 2005-09-03 01:59:59 pm"
```

This command removes a view's previously configured schedule.

```
$ avaki view --set-property /GeneratedViews/Dino schedule
```

avaki view --update

Configure which user a view should be run as.

To use this command, you must be a member of the Administrators group, or you must have write permission on the view.

Syntax

```
avaki view --update <view-path>
  --run-user={<qualified-user-name> | null}
```

<view-path> The grid path to the generated view file.

--run-user= Specify the user to run the view as. Use this format:
 {<qualified- <user>@<authservice>.<authservicetype>.<domain>
 user-name> | For example:
 null} wilma@DefaultAuthService.Grid.Bedrock

The elements of <qualified-user-name> are as follows:

<user>	The name of the user to create the view as.
<authservice>	The name of the user's authentication service.
<authservice- type>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

If you specify null, the view is run as the current user.

avaki virtualdatabase --allow-dbop-creation

Allow a user to create virtual database operations.

Syntax

```
avaki virtualdatabase --allow-dbop-creation
    <qualified-user-name>
```

<qualified-
user-name>

Specify the user who can create virtual database operations. Use the following format:

```
<user>@<authservice>.<authservicetype>.<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-user-name> are as follows:

<user>	The name of the user who can create virtual database operations.
<authservice>	The name of the user's authentication service.
<authservice-type>	The type of authentication service (Grid, Nis, or Ldap).
<domain>	The name of the user's Avaki domain.

avaki virtualdatabase --disallow-dbop-creation

Prevent a user from creating virtual database operations.

Syntax

```
avaki virtualdatabase --disallow-dbop-creation
    <qualified-user-name>
```

<qualified-user-name> Specify the user who must not create virtual database operations. Use the following format:

```
<user>@<authservice>.<authservicetype>.<domain>
```

For example:

```
wilma@DefaultAuthService.Grid.Bedrock
```

The elements of <qualified-user-name> are as follows:

<user> The name of the user who is not allowed to create virtual database operations.

<authservice> The name of the user's authentication service.

<authservice-type> The type of authentication service (Grid, Nis, or Ldap).

<domain-name> The user's grid domain.

avaki virtualdatabase --execute

Execute an ad-hoc query (a SQL statement) against the Avaki virtual database.

Syntax

```
avaki virtualdatabase --execute [--params={in | out |
  inout}:<parameter-types>] <sql> [<parameter-value>+]
```

--params={in |
out | inout}:
<parameter-
types>

Optional. If the virtual database accepts input parameters or produces output parameters, use this option to specify the data type for each input and output. There is no default.

Use in for input parameters, out for output parameters, and inout for parameters that can be both. If you specify multiple parameters, use semicolons (;) to separate them.

You can use the following <parameter-types>:

BIGINT	LONGVARBINARY
BINARY	LONGVARCHAR
BIT	NUMERIC
BOOLEAN	ORACLE_CURSOR
BLOB	OTHER
CHAR	REAL
CLOB	SMALLINT
DATE	TIME
DECIMAL	TIMESTAMP
DOUBLE	TINYINT
FLOAT	VARBINARY
INTEGER	VARCHAR

<sql>

Enter the SQL statement that the virtual database will execute. Enclose the statement in quotation marks.

<parameter-
value>+

Optional. One or more parameter values required by the target virtual database. If a database operation is associated with the virtual database, the parameter values will be processed in the order specified in the database operation's query. Any parameter value that includes a space must be enclosed in quotation marks. Any string at the end of the command that is enclosed in quotation marks will be interpreted as a single value. (That is, enclose each parameter value in its own pair of quotation marks.)

avaki virtualdatabase --show-tables

Display SQL views in the Avaki virtual database.

Syntax

```
avaki virtualdatabase --show-tables [--schema={DATASERVICE |
  VIRTUALDB | <database-connector-name>}]
  [--table=<table-name>] [--type={TABLE | VIEW}]
```

<pre>--schema= {DATASERVICE VIRTUALDB <database- connector-name>}</pre>	<p>Optional. Display SQL views derived from one of the following sources. If you omit this option, Avaki displays SQL views derived from all sources.</p> <p>DATASERVICE Display SQL views based on data services.</p> <p>VIRTUALDB Display SQL views based on virtual database operations.</p> <p><database-connector-name> Display SQL views based on the specified database connector. Use the database connector name for both provisioned SQL views and SQL views generated from database operations.</p>
<pre>--table= <table-name></pre>	<p>Optional. The name of a SQL view (table) to display.</p>
<pre>--type= {TABLE VIEW}</pre>	<p>Optional. Display SQL views of one of the following types. If you omit this option, Avaki displays SQL views of both types.</p> <p>TABLE Display provisioned SQL views.</p> <p>VIEW Display SQL views generated from database operations and data services.</p>

Example

This command displays all the tables derived from virtual database operations. (There is only one such table in the current domain.)

```
$ avaki virtualdatabase --show-tables --schema=VIRTUALDB
```



```

===== TABLE =====
TABLE_NAME: VDOTESTINGWORKSTABLESTARTABLE
TABLE_CAT: mydomain.sybase.com
TABLE_SCHEM: VIRTUALDB
TABLE_TYPE: VIEW
REMARKS: null
TYPE_CAT: null
TYPE_SCHEM: null
TYPE_NAME: null
SELF_REFERENCING_COL_NAME: null
REF_GENERATION: null
===== COLUMN =====
COLUMN_NAME: title; DATA_TYPE: 12; TYPE_NAME: VARCHAR; COLUMN_SIZE: 128;
DECIMAL_DIGITS: 0; NUM_PREC_RADIX: 0; NULLABLE: 1; REMARKS: null; COLUMN_DEF:
null; CHAR_OCTET_LENGTH: 0; ORDINAL_POSITION: 1; IS_NULLABLE: YES; SCOPE_CATALOG:
null; SCOPE_SCHEMA: null; SCOPE_TABLE: null; SOURCE_DATA_TYPE: 0; PRECISION: 128;
===== COLUMN =====
COLUMN_NAME: guid; DATA_TYPE: 12; TYPE_NAME: VARCHAR; COLUMN_SIZE: 128;
DECIMAL_DIGITS: 0; NUM_PREC_RADIX: 0; NULLABLE: 1; REMARKS: null; COLUMN_DEF:
null; CHAR_OCTET_LENGTH: 0; ORDINAL_POSITION: 2; IS_NULLABLE: YES; SCOPE_CATALOG:
null; SCOPE_SCHEMA: null; SCOPE_TABLE: null; SOURCE_DATA_TYPE: 0; PRECISION: 128;
===== COLUMN =====
COLUMN_NAME: description; DATA_TYPE: 12; TYPE_NAME: VARCHAR; COLUMN_SIZE: 128;
DECIMAL_DIGITS: 0; NUM_PREC_RADIX: 0; NULLABLE: 1; REMARKS: null; COLUMN_DEF:
null; CHAR_OCTET_LENGTH: 0; ORDINAL_POSITION: 3; IS_NULLABLE: YES; SCOPE_CATALOG:
null; SCOPE_SCHEMA: null; SCOPE_TABLE: null; SOURCE_DATA_TYPE: 0; PRECISION: 128;
===== COLUMN =====
COLUMN_NAME: dp; DATA_TYPE: 12; TYPE_NAME: VARCHAR; COLUMN_SIZE: 128;
DECIMAL_DIGITS: 0; NUM_PREC_RADIX: 0; NULLABLE: 1; REMARKS: null; COLUMN_DEF:
null; CHAR_OCTET_LENGTH: 0; ORDINAL_POSITION: 4; IS_NULLABLE: YES; SCOPE_CATALOG:
null; SCOPE_SCHEMA: null; SCOPE_TABLE: null; SOURCE_DATA_TYPE: 0; PRECISION: 128;

```

Number of columns: 4

\$

avaki virtualschema --deploy

Deploy a virtual schema (metadata) model into the data catalog.

Syntax

```
avaki virtualschema --deploy --schemaname=<virtual-schema-name>  
  <local-path>
```

--schemaname= The name of the virtual schema to deploy.
<virtual-
schema-name>

<local-path> The local path to the virtual schema to deploy.

avaki virtualschema --undeploy

Undeploy a virtual schema (metadata) model from the data catalog.

Syntax

```
avaki virtualschema --undeploy  
  --schemaname=<virtual-schema-name>
```

--schemaname= The name of the virtual schema to undeploy.
<virtual-
schema-name>

avaki --whoami

Display your Avaki user name. Similar to the Unix **whoami** command.

Syntax

```
avaki whoami
```

– This command has no options.

Attribute reference

This appendix describes cache attributes that you can create as well as read-only system attributes that Avaki creates as needed. You can use the **avaki attribute** commands to create, modify, and delete attributes. See [page 19](#) for details about **avaki attribute**.

This appendix covers the following types of attributes:

- “Avaki grid server attributes” on [page 290](#)
- “File attributes” on [page 290](#)
- “Cache service attributes” on [page 291](#)

Avaki grid server attributes

This section describes read-only system attributes that provide information about an Avaki grid server.

Attribute	Description
system/ObjectHostName	The DNS name or IP address of the machine on which the object is running.
system/os.name	The name of the operating system on the machine on which the object is running.
system/os.arch	The architecture of the machine on which the object is running.

File attributes

This section describes read-only system attributes that provide information about files in a data grid.

Attribute	Description
system/FileSize	The size of a file (in bytes).
system/ChangeTime	The time a file was last changed.
system/ModificationTime	The time a file was last modified.
system/LastAccessTime	The time a file was last accessed.
system/BlockSize	The size of a file's blocks.
system/Blocks	The number of blocks to a file.

Cache service attributes

By default, the attributes for a cache service are configured on a cache-wide basis (for more information, see the *Sybase Avaki EII Provisioning and Advanced Data Integration Guide*). If you set the following attributes on a per-file basis, the settings override the cache-wide read-only system attributes.

Attribute	Description
cacheable/CoherenceWindow	<p data-bbox="484 432 1368 565">Integer specifying the time in seconds after the content has last been updated during which the proxy cache service assumes the content to be fresh and does not check back with the source for expiration. Set this attribute to -1 for a static file that should never expire.</p> <p data-bbox="484 597 1368 696">This attribute overrides the value of the proxy cache service's default coherence window setting, which configures the coherence window on a cache-wide basis.</p> <p data-bbox="484 727 983 756">Valid values: -1, 0 to 2147483647 seconds</p>
cacheable/OfflineExpiration	<p data-bbox="484 788 1368 921">Integer specifying the time in seconds for which content is allowed to remain valid after its source is determined to be offline. A value of -1 means that this file is always available for offline access when its source is determined to be unreachable.</p> <p data-bbox="484 953 1368 1017">This attribute overrides the value of the proxy cache service's offline expiration setting, which configures the coherence window on a cache-wide basis.</p> <p data-bbox="484 1048 983 1078">Valid values: -1, 0 to 2147483647 seconds</p>

Setting up DGAS properties files

This appendix describes the properties file used by Avaki data grid access servers (DGASes).

About the properties file

If you are starting and configuring a data grid access server using the command line tools, you have the option to use a properties file. The DGAS's properties file lets you specify startup options for the DGAS, including NFS and CIFS behavior, the location and behavior of the caches, and threading options. You can also specify these options in the web UI, but you cannot use a properties file if you set up a DGAS in the web UI.

Store the properties file in a file system where it is accessible to the DGAS. You can store it under any path and name; the default path is the Avaki installation directory. There is no default file name. Our convention is to use the name `dgas.properties`.

The section that follows includes a sample properties file. To display a list that includes the names, descriptions, and default values of all DGAS properties, use a command of this form:

```
avaki dgas startup <dgas-name>
```

Properties file format

Below is a sample of the properties file that sets startup options for a data grid access server. The list of properties changes frequently; for authoritative descriptions, use the following command, in which <dgas-name> is the name of a data grid access server:

```
avaki dgas --get-property-list <dgas-name>
```

Read this sample file as a format guideline. Because this sample would set most properties to their default values, it would have no effect if you used it with your DGAS.

Notes Several property values shown in this sample file are not literal defaults:

- The db-path and cache-path properties require paths as values. The paths in this sample are for a default Windows installation. For a Unix installation, replace the string “C:\Avaki Data Grid 6.0” with the path to the Avaki installation directory on the DGAS machine, and use slashes instead of backslashes in the path.
- The default value of the cifs-server-name property is the name of the machine on which the DGAS is running.

Lines in the sample file that begin with # are comments.

```
#####
# Optional properties file for Avaki Data Grid Access Servers
#
#####
non-idempotent-ops-cache-ttl=10000
lock-manager-enabled=false
cifs-wins-server=WINS_SERVER_NAME
cache-block-size=104857600
inactive-blocks-high-water=-1
precreate-async-write-threads=false
cifs-enabled=false
unix-chown-semantics=true
cifs-max-dirent-per-read=100
cache-frags-per-block=100
inactive-blocks-low-water=31457280
user-cache-size=5
unix-file-mode-semantics=false
cache-enabled=true
cifs-bind-address=*
nfs-preferred-xfer-size=32768
nfs-v3-enabled=true
```

```
non-idempotent-ops-cache-size=20
name=Access1
nfs-max-dirent-per-read=50
async-write-blocks-pool-size=20
max-async-write-bsize=1048576
updates-portmap=true
db-path=C:\Avaki Data Grid 6.0\DGAS\dgas_db
cifs-name-registration-mode=1
max-num-async-write-blocks=5
max-async-write-threads=6
mem-cache-stages=2
cache-path=C:\Avaki Data Grid 6.0\DGAS\dgas_caches
min-async-write-threads=0
file-cache-coherence-window=60000
dir-cache-size=100
nlm-port=0
max-cached-files=50
nfs-port=2049
mount-port=0
min-nfs-udp-threads=0
cifs-server-name=<DGAS host name>
cifs-host-announce-interval=5
netbios-smb-enabled=false
user-mapping-ttl=7200000
tcpip-smb-enabled=false
num-active-blocks=10
mount-server-enabled=true
case-sensitive-searches=false
precreate-nfs-udp-threads=false
mem-fetch-threshold=1048576
dir-cache-coherence-window=120000
file-obj-cache-size=256
cifs-name-registration-frequency=3600
nfs-v2-enabled=true
cifs-server-comment=Avaki CIFS Server
nfs-max-dirplusent-per-read=50
cifs-server-broadcast-address=255.255.255.255
dir-cookie-check-enabled=false
cifs-server-domain=AVAKI
max-nfs-udp-threads=10
```


Configuring cron schedules

You can use a cron specification to indicate the schedule by which the following tasks will be performed:

- Adding a caching schedule for a data service; see **avaki dataservice --add-schedule** (page 48).
- Adding a caching schedule for a database operation; see **avaki dbop --add-schedule** (page 74).
- Adding a pin schedule for a cached directory; **avaki directory --add-schedule** (page 118).
- Tagging a directory or a directory tree for caching within a cache service; see **avaki directory --cache** (page 122).
- Adding a pin schedule for a cached file; see **avaki file --add-schedule** (page 131).
- Tagging a file within a cache service for active caching; see **avaki file --pin** (page 138).
- Adding a schedule by which LDAP users will be automatically imported; see **avaki ldap --add-schedule** (page 150).
- Adding a schedule by which NIS users will be automatically imported; see **avaki nis --add-schedule** (page 178).
- Adding a rehash schedule to a shared directory; see **avaki share --add-rehash-schedule** (page 232).
- Adding a generation schedule to a view; see **avaki view --add-schedule** (page 260).

You can include a *cron schedule option* in any of the commands listed above. The syntax for a cron schedule option is as follows:

```
--cron-schedule="<seconds> <minutes> <hours>
<days-of-month> <months> <days-of-week> [<years>]"
```

Alternatively, you can include a *cron expression* in any advanced schedule in the web user interface. A cron expression's syntax is similar to that of the cron schedule option:

```
<seconds> <minutes> <hours> <days-of-month> <months>
<days-of-week> [<years>]
```

You can use the following values in either type of cron specification:

Cron field	Valid values	Special characters allowed*
<seconds>	An integer from 0 to 59	Comma (,), hyphen (-), asterisk (*), and slash (/)
<minutes>	An integer from 0 to 59	Comma (,), hyphen (-), asterisk (*), and slash (/)
<hours>	An integer from 0 to 23	Comma (,), hyphen (-), asterisk (*), and slash (/)
<days-of-month>	An integer from 1 to 31	Comma (,), hyphen (-), asterisk (*), slash (/), question mark (?), and L
<months>	An integer from 1 to 12 or JAN-DEC	Comma (,), hyphen (-), asterisk (*), and slash (/)
<days-of-week>	An integer from 1 to 7 or SUN-SAT	Comma (,), hyphen (-), asterisk (*), slash (/), question mark (?), L, and pound sign (#)
<years>	Optional. An integer from 1970 to 2099	Comma (,), hyphen (-), asterisk (*), and slash (/)

* See the table on the next page for more information on special characters.

You can use the following special characters in either type of cron specification:

Character	Description
Comma (,)	A comma may be used to specify additional values. For example, “MON,WED,FRI” in the day-of-week field means “the days Monday, Wednesday, and Friday.”
Hyphen (-)	A hyphen may be used to specify ranges. For example, “10-12” in the hour field means “the hours 10, 11, and 12.”
Slash (/)	A slash may be used to specify increments. For example, “0/15” in the seconds field means “the seconds 0, 15, 30, and 45.” The syntax “5/15” in the seconds field means “the seconds 5, 20, 35, and 50.”
Asterisk (*)	An asterisk can be used to specify all valid values in a given field. For example, “*” in the minute field means “every minute.”
Question mark (?)	A question mark can be used in the day-of-month and day-of-week fields to specify “no specific value.” This character is useful when you need to specify something in one of the two fields, but not in the other.
L	The L character may be used in the day-of-month and day-of-week fields. This character stands for “last” and has a different meaning in each of the two fields. For example, the value “L” in the day-of-month field means “the last day of the month”; day 31 for January, and day 28 for February on non-leap years. If used in the day-of-week field by itself, L means “7” or “SA.”. If used in the day-of-week field after another value, L means “the last xxx day of the month.” For example, “6L” means “the last Friday of the month.”
Pound sign (#)	A pound sign may be used in the day-of-week field to specify “the nth” xxx day of the month. For example, the value of “6#3” in the day-of-week field means the third Friday of the month (day 6 = Friday and “#3” = the 3rd one in the month).

Data type mapping files

When you provision a table into Avaki from a relational DBMS (that is, when you provision a *SQL view*), Avaki uses a default set of mappings between the SQL data types used in the DBMS and the data types used by Avaki's Axion query engine. If the default data type mappings yield unsatisfactory results, you can override them for a particular database connection. To override the default mappings, you must

- Create an XML mapping file as described in this appendix and share it into the Avaki data catalog.
- Use the **--mapping-file** option of the **avaki dbconn --jdbc** command ([page 66](#)) to specify the location of the mapping file.

In this appendix:

- [“Format and contents of type mapping files” on page 302](#)
- [“Setting the source data type: data type inconsistencies” on page 304](#)
- [“Logging of type mapping” on page 305](#)

Format and contents of type mapping files

In your mapping file, specify only the subset of types that need to be mapped differently from the default.

The format of the mapping file is as follows:

```
< mappings >

  < mapping >
    < tableType > all < /tableType >
    < sqlTypeInt > 3 < /sqlTypeInt >
    < sqlTypeName > decimal < /sqlTypeName >
    < outputType > long < /outputType >
    < precisionSource > precision < /precisionSource >
    < scaleSource > decimalDigits < /scaleSource >
    < displaySizeSource > size < /displaySizeSource >
  < /mapping >

  < mapping > . . . < /mappings > . . .

< mappings >
```

Each mapping in the file contains the following information. Elements are required except as noted.

Element	Valid values	Description
tableType	all provisioned generated	To specify different mappings for SQL views (tables) that are provisioned into Avaki and SQL views that are generated from database operations or data services, enter a mapping element for each table type. Otherwise, enter a single mapping element for the designated type and set tableType to all.

Element	Valid values	Description	
sqlTypeInt	ARRAY 2003	JAVA_OBJECT 2000	The integer constant representing the SQL type to be mapped. For example, use 3 for a decimal. See “ Setting the source data type: data type inconsistencies ” on page 304 for more information on setting the value of this element.
	BIGINT -5	LONGVARIABLE -4	
	BINARY -2	LONGVARCHAR -1	
	BIT -7	NULL 0	
	BLOB 2004	NUMERIC 2	
	BOOLEAN 16	OTHER 1111	
	CHAR 1	REAL 7	
	CLOB 2005	REF 006	
	DATALINK 70	SMALLINT 5	
	DATE 91	STRUCT 2002	
	DECIMAL 3	TIME 2	
	DISTINCT 2001	TIMESTAMP 93	
	DOUBLE 8	TINYINT -6	
	FLOAT 6	VARBINARY -3	
	INTEGER 4	VARCHAR 12	
sqlTypeName	String	The name of the SQL type to be mapped. Specify the string the underlying database or database operation returns as the name of the SQL type.	
outputType	boolean	varchar2	The Axion type to which you are mapping the designated SQL type. More information about Axion data types is available at http://axion.tigris.org/datatypes.html
	byte	float	
	unsignedbyte	number	
	short	numeric	
	unsignedshort	clob	
	integer	compressedclob	
	int	blob	
	smallint	compressedblob	
	unsignedinteger	date	
	long	time	
	character	timestamp	
	char	java_object	
	varchar	varbinary	

Element	Valid values	Description
precisionSource	charOctetLen dataType decimalDigits default	precision remarks scopeCatalog scopeSchema
scaleSource	isNullable name nullable numPrecRadix	scopeTable size sourceDataType typeName
displaySizeSource	ordinalPos	
		Optional. The element in the native table's metadata that you are mapping to the <code>precision</code> element in Avaki.
		Optional. The element in the native table's metadata that you are mapping to the <code>scale</code> element in Avaki.
		Optional. The element in the native table's metadata that you are mapping to the <code>displaySize</code> element in Avaki.
		Note: Because Avaki does not allow <code>displaySize</code> to be smaller than the default for its data type, mapping the Avaki <code>displaySize</code> element to a SQL element with a smaller value has no effect.

Setting the source data type: data type inconsistencies

When a back-end database responds to a JDBC request, it sometimes performs type translation. As a result, a column that appears to be of one type when viewed or accessed through a JDBC-based tool or application (the Squirrel SQL client, for example) might appear to be of a different type when viewed with a non-JDBC tool. This is not a common occurrence. However, when you set the value of the `sql-TypeInt` element in a mapping file, be aware that the database's internal notion of a column's data type is not important for mapping purposes—the type you need to capture is the one the database returns through JDBC.

Logging of type mapping

By default, Avaki's log4j logging mechanism records type mapping decisions during the creation of provisioned or generated SQL views (virtual tables). The log entries indicate the JDBC type being mapped to the Avaki/Axion type.

Logging of type mapping is controlled by a log4j category called `TypeMapping`. If you want the type mapping for a particular SQL view to be logged (or not logged), logging of type mapping must be enabled (or disabled) in the log properties file on the grid server that hosts the database operation or data service on which the SQL view is based.

For information on configuration of logging, see the appendix on log properties files in the *Sybase Avaki EII Administration Guide*.

Glossary

Terms printed in *italics* are defined in the glossary.

access control list

(ACL) A list, for a given file, directory, or other Avaki object, of permissions—read, write, execute, delete, and owner—that control which users and groups can view, modify, invoke, and remove the object, and edit the object's ACL.

ACL

See *access control list*.

ad-hoc query

A mechanism that lets you directly query a database in SQL. The query must run through an existing Avaki *database connector*. You can run an ad-hoc query using either the CLI or a *JDBC driver*. Ad-hoc queries can be thought of as single-use *database operations*.

attribute

A property of an *Avaki directory*, file, *service*, or other object. Each attribute has a name, a type (string, integer, float, date, time, or timestamp) and a value. System attributes are read-only; you can change the values of other attributes. You can also create new attributes and add them to objects as needed.

authentication service

A *service* associated with an *Avaki domain* that authenticates an Avaki user's identity and provides security credentials each time the user logs in. Avaki can be configured to use third-party directory services as authentication services for login; for user accounts created directly in the Avaki domain, Avaki uses its own default authentication service.

Avaki directory

Avaki software creates a single, unified namespace that is accessible (subject to Avaki *access control lists*) to all users in the *Avaki domain*. The namespace, called the *data catalog*, is arranged as a hierarchy of Avaki directories (folders). The catalog directory structure is stored by the domain's grid servers and its GDC, while the physical files remain in their original locations in your local file systems. When you work with directories, it's important to distinguish between Avaki directories, which are part of the data catalog, and local directories, which reside in your local file system.

Avaki domain

The basic administrative unit of the Avaki EII system. An Avaki domain consists, at a minimum, of one *grid domain controller* and may also include one or more *grid servers*, *share servers*, *proxy servers*, *data grid access servers*, and *command clients*. See also *domain name*.

Avaki group

A set of users who have the same permissions on one or more Avaki objects. You can use the group name in place of a user name when you set permissions or create *access control lists*.

Avaki installation directory

The directory in your local file system where Avaki software is installed. This is not a *data catalog* directory.

Avaki share

(Also shared directory.) A pointer in the Avaki *data catalog* to a directory or file in the underlying local file system. When you browse the data catalog, Avaki shares look like—and can be accessed like—other Avaki directories. Contrast with *CIFS share*.

Avaki server

A *service* that starts, stops, and monitors other Avaki services on a particular computer. Every server is part of an *Avaki domain*. A server is permanently attached to the computer where it is started. There are several types of server: *data grid access servers*, *grid domain controllers*, *grid servers*, *share servers*, and *proxy servers*.

Avaki Studio

A graphical, metadata-based data integration tool that lets you

- Build data flows by dragging and dropping input sources, operators, and output targets. You can deploy your data flows as *Avaki data services*.
- Import or create *metadata models* and apply them to Avaki objects or use them to build new data services.

In addition, you can use Studio to perform provisioning tasks (creating *database connectors*, *database operations*, *virtual database operations*, and *SQL views*), manipulate *categories*, and edit *ad-hoc queries* and *attributes*.

cache service

(Formerly proxy cache service.) A staging service that stores copies of files, *database operation* results, and *data service* results. Caching improves retrieval performance. To ensure that an object is stored in the cache, you can *pin* a file or directory in the data catalog, or schedule a database operation or data service. A cache service can provide remote caching, local caching, or both. The freshness of cached data is controlled by a data expiration interval that determines how long cached data is considered valid and by a cache coherence window that tells the cache service how often to check whether cached data is still valid. If cached data is too old to satisfy a new request (or is not stored in this cache), the cache service does one of the following:

- If the database operation or data service that produced the data is local to this cache service, the cache service triggers execution of the database operation or data service.
- If the database operation or data service that produced the data is remote from this cache service, this cache service requests the data from the data source's local cache service.

A cache service can be associated with a *data grid access server*, a *grid server*, or a local user in a CLI session. See also *local cache*, *remote cache*, *on-demand caching*, and *scheduled caching*.

category

A mechanism for classifying and organizing the contents of the *data catalog*. Like *Avaki directories*, categories serve as containers for objects in the data catalog. Anything in the data catalog—views, data services, shared files, even Avaki directories themselves—can be assigned to a category. Categories are hierarchical, they have attributes, and Avaki *access control lists* regulate access to them.

CIFS client

A machine that mounts files or directories from the Avaki *data catalog* by connecting to a *CIFS share* through an Avaki *data grid access server*. A CIFS client need not have Avaki software installed. (CIFS—Common Internet File System—is a file-sharing protocol based on the file system implemented by Windows.)

CIFS share

A directory or file that has been exported (shared) from the Avaki *data catalog*. A CIFS share can be mapped into a Windows file system like a network drive. When you browse the Windows file system, CIFS shares look like—and can be accessed like—other files and directories. CIFS shares are created through a *data grid access server*. Contrast with *Avaki share*.

client

Avaki supports several types of client: *Avaki Studio*, *CIFS clients*, *command clients*, *JDBC/ODBC clients*, *NFS clients*, *web clients*, and *WS clients*.

command client

A machine that can issue Avaki commands but does not contribute resources to the *Avaki domain*.

connect port

The connect port on a *grid domain controller*, *grid server*, *data grid access server*, *proxy server*, or *share server* accesses the JNDI naming service or RMI registry for the underlying application server. The connect port is one of many ports that a GDC or server uses to communicate with other Avaki objects. You must supply the connect port number of a target grid server or GDC whenever you connect a new object (another server, a copy of Avaki Studio, or a *command client*, for example) to an *Avaki domain*. When you *interconnect* two Avaki domains, you must supply each domain's connect port number to the other one.

data catalog

A hierarchical structure similar to a file system that encompasses all objects in an *Avaki domain*. The data catalog contains *Avaki directories* and files, *Avaki shares*, *Avaki servers*, *SQL views*, *database operations* and *data services*, and other objects.

data grid access server

(DGAS) An *Avaki server* that makes *Avaki directories* and their contents available to *CIFS clients* and *NFS clients*.

data service

An operation that transforms data obtained from sources in the *data catalog*. Input data can come from any number of sources, including:

- other data services
- data catalog files (which can be *generated views*)
- *Avaki database operations* (which in turn extract the data from relational databases)
- HTTP requests
- Web service invocations

You can generate the code that manipulates the data by creating a *view model* in *Avaki Studio*, or by writing a custom *data service plug-in* using Java, JavaScript, or XSLT. Data service output can be in rowset or XML format. Data services are run by the *execution services* on *grid servers*, they can be scheduled, and their results can be cached.

data service plug-in

The logic for a *data service*, written in Java, JavaScript, or XSLT. Data service plug-ins are modular—you can use the same plug-in for multiple data services. *Avaki Studio* creates data services and plug-ins simultaneously, so if you use Avaki Studio to create data services, you don't have to worry about plug-ins. You can also use the Avaki Plug-in Wizard to create data service plug-ins.

database connector

A mechanism that enables one or more *database operations*, *SQL views*, or *ad-hoc queries* to connect to a relational database.

database operation

(DBOP) A mechanism that can

- extract data from a relational database and deliver it on demand to a *view generator* or a *data service*, or
- modify data in a relational database.

A database operation can be a SQL statement or a stored procedure call.

dependency

A relationship in which an Avaki object requires input from other Avaki objects. A *data service* might require input from one or more *database operations* or from other data services. A *view generator* might depend on a database operation for input. A database operation can serve as an input source for one or more data services or view generators. Generated *SQL views* depend on database operations, virtual database operations, or data services. You can use *Avaki Studio*, the web UI, or the CLI to list input and output dependencies for any data service, database operation, or view.

DGAS

See *data grid access server*.

distributed transaction

A set of related operations (typically SQL operations such as SELECT, INSERT, UPDATE, DELETE, and CALL) that

- involve one or more databases, and
- might lead to unwanted results (such as leaving participating databases in an inconsistent state or producing inconsistent reads) if some of the operations complete and others do not, and therefore
- must all be executed at once, as a single transaction.

The individual operations that make up a distributed transaction are performed by *database operations* that use *database connectors* configured with XA-capable *JDBC drivers*; all the database operations are executed, using the two-phase commit protocol, by a specially configured *data service*. The two-phase commit protocol is designed to ensure that the participating databases will be left in a consistent state—that is, that all the operations in the distributed transaction will be completed, or none of them will.

domain name

A unique alphanumeric identifier for an *Avaki domain*. The domain name is assigned by the Avaki administrator when the Avaki domain is initialized. The domain name has a maximum length of 30 characters.

enterprise information integration

(EII) A software system that

- enables applications and users to access, without replication, both raw and integrated data from multiple heterogeneous distributed data sources while hiding the complexity of the data sources, and
- provides tools enabling users and data owners to further integrate and transform data.

exclusion

See *schedule exclusion*.

execution service

Execution services execute *data services*. There is an execution service on every *grid server*, and you can configure a pool of execution services for load-sharing. When a pool is in place, a data service can be run by any execution service in its grid server's pool.

failover

The transition of control from a failing or unreachable primary *grid domain controller* to a secondary grid domain controller.

federated data access

A scheme that allows independently controlled elements to be shared into a single namespace. Files, user accounts, and other objects maintain their separate identities and remain under the control of their owners, but—subject to access controls—the objects can be accessed, managed, and viewed as if they were part of a single system.

GDC

See *grid domain controller*.

generated view

A file created by a *view generator*; it may contain data obtained from a *database operation*, a *data service*, a file, or an HTTP source. Like other files, generated views exist in a local file system and are shared into the *data catalog*.

grid

A heterogeneous group of networked resources that appears and functions as one operating environment. A data grid like the Avaki Enterprise Information Integration (EII) system provides secure, shared access to data.

grid directory

See *Avaki directory*.

grid domain

See *Avaki domain*.

grid domain controller

(GDC) The first server in an *Avaki domain* is the grid domain controller. The GDC maintains a portion of the Avaki domain's namespace and provides authentication services. It can also run Avaki commands, share data, and monitor other servers. (That is, the GDC functions as a *grid server*.) If the domain is configured for *failover*, it has both a primary GDC and a secondary GDC; the secondary is updated at regular intervals and takes over management of the domain if the primary fails. Any Avaki shares managed by the primary are read-only on the secondary.

grid server

An *Avaki server* that maintains a portion of the *Avaki domain*'s namespace, runs Avaki services such as shares, execution services, caches, and searches, and allows you to run Avaki's web UI and execute Avaki commands.

group

See *Avaki group*.

hard link

Provides an alternate name for an item in the *data catalog*. Changes to the object's other names have no effect on the hard link: you can move or change a file's original name and the hard link will still know where to find the file. To delete a hard-linked object, you must remove its original name. Contrast with *soft link*.

interconnect

To create a unidirectional link from one *Avaki domain* to another. Interconnecting lets an Avaki domain make its *data catalog* visible to users in another domain (subject to Avaki access controls).

JDBC driver

JDBC (Java Database Connectivity) drivers allows application programmers to access database data shared in the *data catalog*. When a JDBC driver accesses data, it returns a JDBC result set that's immediately available to your program. JDBC drivers can:

- Call any *data service* in the data catalog
- Call any *database operation* in the data catalog
- Perform SQL `select` operations against *SQL views* in the data catalog

Sybase offers three JDBC drivers for use with Avaki EII software:

- The included Avaki JDBC driver
- jConnect, Sybase's standard JDBC driver
- An XA-capable driver for use with *database connectors* that support *distributed transactions*

link

See *hard link*, *soft link*.

local cache

A *cache service* that runs on the same *grid server* as a *database operation* or a *data service* that generates cachable data. The local cache stores results produced by local database operations and data services so they don't have to execute for every new request. See also *remote cache*.

metadata model

A construct in *Avaki Studio* that expresses a schema by defining a set of tables. A table in a metadata model can be mapped (linked) to an Avaki object such as a *data service* or a *database operation*, or to a table in a relational database. The mapping lets you address each mapped object by the name of the corresponding table in the metadata model. You can also derive a *view model* schema from a metadata model. When you do this, you ensure that the results of any data service deployed from the view model will conform to the metadata model's schema.

NFS client

A machine that mounts the Avaki *data catalog* (or a portion of it) as a directory by connecting to an Avaki *data grid access server*. An NFS client need not have Avaki software installed. (NFS—Network File System—lets you add file systems located on a remote computer to the directory structure on your own computer.)

ODBC

ODBC (Open DataBase Connectivity) is an API for databases on Windows. An ODBC driver (such as the the Sybase Organic ODBC driver included with Sybase ASE) allows Avaki to communicate with Windows database applications.

on-demand caching

A scheme by which an object is cached only if it's used—for example, results are cached when a *database operation* or a *data service* is executed, or a file is cached when a user or application reads it. On-demand caching uses a fixed expiration interval to determine data freshness. On-demand caching is suitable for objects that are rarely accessed or that change at irregular intervals. Contrast with *scheduled caching*.

pin

To mark an *Avaki directory* or file for *scheduled caching*. See also *cache service*.

plug-in

See *data service plug-in*.

primary GDC

See *grid domain controller*.

proxy server

An *Avaki server* that allows *Avaki domains* on opposite sides of a firewall or a Network Address Translator (NAT) to communicate with one another.

queries

See *ad-hoc query*.

query engine

An *Avaki service* that executes SQL queries against the *SQL views* (tables) that make up the *Avaki virtual database*. A query engine analyzes queries, pushes as much of the work as possible down to the underlying relational database (if there is one), and performs the remaining operations (such as joins across tables from different databases) itself. There is a query engine on each *grid server*.

remote cache

A *cache service* that runs on a grid server that is remote from an *Avaki service* (a *database operation* or a *data service*) that generates cachable data. The remote cache stores results produced by distant services so the results don't have to be fetched over the network to satisfy every new request. Users and applications that access remote data through the cache may have access to cached copies even when the remote data source is unavailable. See also *local cache*.

scheduled caching

A scheme by which an object is cached according to a schedule that you create. The schedule specifies when the object is first cached and how often (or following what trigger event, such as a change to a file) the cache is refreshed. If the object is a *data service* or a *database operation*, the schedule runs it to put fresh results in the cache. Scheduled caching, which overrides other types of caching, is suitable for objects that are updated frequently or on a regular basis. Contrast with *on-demand caching*.

schedule exclusion

A named period of time during which scheduled activities can be prevented from running. You can apply an exclusion to as many schedules as you want. Scheduled activities include refreshing *Avaki shares* and imported user accounts, and caching files, directories, and the results of *database operations*, *data services*, and *generated views*.

secondary GDC

See *grid domain controller*.

service

An Avaki object that performs a function in the domain (stores data or authenticates users, for example). Services provided in Avaki software include *Avaki directories*, *Avaki shares*, *Avaki servers*, *authentication services*, *execution services*, and user accounts.

share

A point of connection between the Avaki *data catalog* and a native file system or file system tool. Avaki supports two kinds of shares: *Avaki shares* and *CIFS shares*.

share server

An *Avaki server* whose only task is to manage *Avaki shares*—local directories that are exported (shared) into the *data catalog*. (Grid servers can also manage shares.)

shared directory

See *Avaki share*.

soft link

A pointer to a particular location (name) in the Avaki *data catalog*. If the object at that location is moved, deleted, or renamed, the soft link leads nowhere. Soft links can be created only in the CLI. Contrast with *hard link*.

SQL view

A virtual table—a *data catalog* entry that represents a table in a relational database, a *database operation*, or a *data service*. SQL views can be created in three ways:

- Provisioned directly from a table in an underlying database
- Generated from a database operation or data service
- Mapped from a database table, a database operation, or a data service, using the *Avaki Studio* metadata model editor

Every SQL view is part of the Avaki *virtual database*. SQL views are treated as relational tables by the Avaki *query engine*. SQL view data can be accessed using standard SQL statements by connecting to Avaki with ODBC or JDBC, or via an Avaki *virtual database operation*.

update notification

A message issued when a *generated view* is updated. A view that receives data from another view can be configured to regenerate itself (using the new data) upon receipt of an update notification.

view generator

A mechanism that does one of the following: extracts data from a file or an HTTP source, obtains data from an Avaki *data service*, or uses an Avaki *database operation* to extract data from a relational database. The view generator can display the data, perform an XSLT transform, save the data as a *generated view* file, and/or update a database. Contrast with *data service*.

view model

The graphical representation of a data flow that you can build in *Avaki Studio*. A view model typically includes one or more input sources (such as *database operations* or *data services*), one or more operations to combine or transform the data, and an output target. When you deploy a view model, it becomes an Avaki data service.

virtual database

The set of all *SQL views* in an *Avaki domain*, including those provisioned from external databases and those generated from *data services* and *database operations*. You can execute SQL queries on the SQL views in the virtual database as if they were tables in a single database.

virtual database operation

A *database operation* whose source database is the Avaki *virtual database* itself. Use virtual database operations if you want to encapsulate and reuse SQL SELECT queries against *SQL views* (provisioned or generated).

web services client

See *WS client*.

WS client

(Also web services client.) A tool or a piece of code that is part of a customer application and that makes SOAP calls to web services on an Avaki grid server. The SOAP calls can request data from the Avaki *data catalog*, from a *database operation*, or from a *data service*.

Master Index

In electronic copies of this book, the index links to other books in the documentation set work only as long as the PDF files are stored in the same directory.

Key

AD: *Administration Guide*

API: *API Guide*

C: *Command Reference*

O: *Overture*

P: *Provisioning & Advanced Data Integration Guide*

S: *Data Integration with Avaki Studio*

Symbols

* asterisks in command syntax AD:xvi, C:xv, P:xi
- hyphens in command syntax AD:xvi, C:xvi, P:xii
+ plus signs in command syntax AD:xvi, C:xv, P:xi
.amm files S:11
.avm files S:11
.js files S:11
.jsi files S:11, S:75
 sample S:115
.NET
 AvakiAPI.disco WSDL discovery file API:3
 sample web services client API:9
 SSL certificates API:9
.project files S:11
<> angle brackets in command syntax AD:xvi, API:vii, C:xv,
 P:xi
= equal signs in command syntax AD:xvi, C:xvi, P:xii
[] square brackets in command syntax AD:xv, C:xv, P:xi
_ (underscore) characters in Avaki names API:81
{ } curly brackets in command syntax AD:xv, C:xv, P:xi
| vertical bars in command syntax AD:xv, C:xv, P:xi

A

About My Domain screen AD:98
AbstractTransformer class P:243
AbstractTransformerFactory class P:244
access control in view models S:74
access control lists, See ACLs
accessibleDBOp SOAP operation API:42
accessibleDS SOAP operation API:36
accessiblePath SOAP operation API:19
accounts for grid users, See users AD:167
ACLs
 about O:45
 adding users and groups AD:243

ACLs (*continued*)

 defined AD:349, API:83, C:307, O:61, P:289, S:175
 deny permissions ineffective for owners, admins O:46
 displaying AD:237, C:186
 for database operations P:22, P:36
 for SQL views P:46
 granting or denying access to everyone O:44
 in grid groups O:43
 interpreting O:48
 modifying AD:239, C:41, S:97
 on cached objects O:50
 on new Avaki shares AD:261
 on new files O:49
 ownership AD:242, O:46
 permissions in AD:242, AD:307
 removing users and groups from AD:242
 sample O:45
 setting for a grid object AD:171
 setting for database operations P:14
 using interconnect IDs to add users and groups to AD:304
Active Directory AD:148
 domain users group AD:157, AD:159, C:155
 See also authentication services, LDAP AD:148
addInputParameter JavaScript method for data service
 plug-ins P:202
addInputStream JavaScript method for data service
 plug-ins P:202
ad-hoc queries
 as web services
 AdHocDBOPExecutionParams complex type API:13
 executeAdHocDBOp SOAP operation API:43
 executeAdHocDBOpWithOutput SOAP operation API:44
 executeAdHocDBOpWithOutputAttach SOAP
 operation API:46
 executeAdHocDBOpWithOutputString SOAP
 operation API:47

ad-hoc queries (*continued*)

- code samples API:74
- defined AD:349, API:83, C:307, O:61, P:289, S:175
- enabling C:66
- enabling on a database connector P:4
- executing C:63
- on virtual database
 - executing C:282
 - parameter types, specifying C:283
 - parameter types, specifying C:65
 - using JDBC driver to run API:69, API:74
- AdHocDBOPExecutionParams complex type API:13
- administrative user accounts, setting up AD:44
- Administrators group O:43
 - about AD:45
 - permissions for AD:240
- admission policies AD:332
 - about AD:85
 - adding C:91
 - creating AD:87
 - deleting AD:88, C:97
 - displaying C:114
 - displaying Windows domains for C:114
 - setting default policies C:112
 - setting Windows domain info C:111
 - unsetting Windows domains for C:116
- aggregate functions S:110
 - in SQL statements, aliasing column names for P:25
- Aggregate operator S:108
- AIX requirements AD:3
- algorithms for join operator S:155
- aliases
 - for column names P:25, S:42
 - aliases for GDC machines AD:10
- Allow permission in ACL AD:243, O:48
- angle brackets in command syntax AD:xvi, API:vii, C:xv, P:xi
- Apache Ant for compiling data service plug-ins P:184
- Apache Axis API:5
 - data catalog example API:18
 - data service example API:35
 - database operations example API:42
- APIs
 - data catalog API:18
 - data services API:34
 - database operations API:40
 - for data service plug-ins
 - about P:185
 - distributed transaction API P:188
 - general data service API P:186
 - TrAX (Transformation API for XML) P:243
 - web services
 - about API:2
 - data service API:35
 - reference API:11
- AROMValue parameters P:212
- As is permission in ACL AD:242
- ASE, see Sybase ASE

- asterisks in command syntax AD:xvi, C:xv, P:xi
- attribute --delete command C:19
- attribute --list command C:19
- attribute --update command C:21
- attributes
 - about AD:245
 - configuring for SQL views P:44
 - creating AD:248
 - and modifying S:100
 - defined AD:349, API:83, C:307, O:61, P:289, S:175
 - deleting AD:254, C:19
 - displaying AD:246
 - displaying details about C:19
 - ldap/importOnDemand AD:158
 - nis/importOnDemand AD:164
 - of cache services C:291
 - of grid servers C:290
 - of patches C:290
 - searching on AD:233
 - setting values AD:252, C:21
 - system attributes AD:248
 - types of AD:250, C:22, S:101
 - user-defined attributes AD:248
 - who can edit S:101
- audit logging
 - about O:14
 - configuring AD:319
 - events captured by AD:322
- audit logs AD:108
- authentication in Avaki O:41
 - using AvakiPrincipal API:13
- authentication services
 - configuring default groups C:218
 - configuring default users C:220
 - configuring GIDs C:217, C:221
 - configuring UIDs C:219, C:222
 - defined AD:349, API:83, C:307, O:61, P:289, S:175
 - deleting AD:166
 - displaying information about AD:166
 - grid, creating groups on C:141
- LDAP
 - adding schedule exclusions for refreshing C:152
 - adding search bases C:160
 - deleting authentication services C:153
 - deleting import schedules C:154
 - deleting search bases C:160
 - displaying information about C:157
 - enabling users C:250
 - importing groups from AD:159, C:155
 - importing users from AD:157, C:155
 - integrating into the grid AD:148, C:157
 - listing import schedules C:158
 - scheduling refreshes AD:185
 - scheduling user imports C:149
 - setting page size for imports AD:145
 - updating C:160

- authentication services (*continued*)
 - NIS
 - adding schedule exclusions for refreshing C:180
 - deleting authentication services C:181
 - deleting import schedules C:181
 - displaying info about C:183
 - enabling users C:250
 - importing groups from AD:165, C:182
 - importing users from AD:164, C:182
 - integrating into the grid AD:162, C:184
 - listing import schedules C:184
 - scheduling user imports C:177
 - updating C:185
 - refreshing imported accounts AD:185
 - specifying for JDBC connections API:69
 - types O:41
 - specifying for JDBC connections API:69
- authentication using AvakiPrincipal API:9
- auto-restart
 - about AD:37, C:5
 - configuring for a DGAS C:5
 - configuring for a GDC AD:38, C:9
 - configuring for a grid server AD:51, C:9
 - configuring for a proxy server C:12
 - configuring for a share server C:15
- avaki attribute --delete C:19
- avaki attribute --list C:19
- avaki attribute --update C:21
- avaki backup C:23
- avaki cache --evict C:24, C:32
- avaki cache --evict --all C:25
- avaki cache --evict --deleted C:26
- avaki cache --get C:27
- avaki cache --invalidate C:27
- avaki cache --invalidate --all C:28
- avaki cache --invalidate-dataservice-results C:29
- avaki cache --invalidate-dbop-results C:30
- avaki cache --list C:31
- avaki cache --set C:33
- avaki cache --unset C:34
- avaki cat C:35
- avaki categories --add-to-category C:35
- avaki categories --create C:36
- avaki categories --delete C:37
- avaki categories --describe C:37
- avaki categories --list C:38
- avaki categories --remove-from-category C:38, C:40
- avaki categories --set-description C:39
- avaki cd C:41
- avaki chmod C:41
- avaki chown C:42
- avaki client C:45
- avaki client --connect command AD:94
- avaki cp C:46
- avaki database operation --list-schedules C:88
- avaki dataservice --add-schedule C:48
- avaki dataservice --create C:52
- avaki dataservice --delete C:52
- avaki dataservice --delete-schedule C:54
- avaki dataservice --depends C:54
- avaki dataservice --execute C:55
- avaki dataservice --generate-sql view C:56
- avaki dataservice --info C:57
- avaki dataservice --list-schedules C:58
- avaki dataservice --update C:58
- avaki dbconn --allow-dbop-creation C:59
- avaki dbconn --delete C:61
- avaki dbconn --disallow-dbop-creation C:62
- avaki dbconn --execute C:63
- avaki dbconn --info C:64
- avaki dbconn --jdbc C:66
- avaki dbconn --provision-tables C:71
- avaki dbconn --show-tables C:73
- avaki dbconn --test C:72
- avaki dbop --add-schedule C:73
- avaki dbop --delete C:78
- avaki dbop --delete-schedule C:78
- avaki dbop --depends C:79
- avaki dbop --execute C:80
- avaki dbop --generate-sql view C:82
- avaki dbop --info C:83
- avaki dbop --jdbc C:83
- avaki dbop --jdbc --create-virtual-dbop C:87
- avaki dgas --add-admission-policy C:91
- avaki dgas --add-group-mapping C:92
- avaki dgas --add-user-mapping C:94
- avaki dgas --cifs-share-info C:95
- avaki dgas --clear-cached-credentials C:95
- avaki dgas --create-cifs-share C:96
- avaki dgas --delete-admission-policy C:97
- avaki dgas --delete-cache C:98
- avaki dgas --delete-cifs-share C:99
- avaki dgas --delete-group-mapping C:99
- avaki dgas --delete-user-mapping C:100
- avaki dgas --disconnect-cifs-client C:101
- avaki dgas --get-cache-size C:101
- avaki dgas --get-cache-statistics C:102
- avaki dgas --get-free-disk-space C:102
- avaki dgas --get-properties C:103
- avaki dgas --get-property C:103
- avaki dgas --get-property-list C:104
- avaki dgas --initialize C:104
- avaki dgas --list-cifs-clients C:105
- avaki dgas --list-cifs-shares C:105
- avaki dgas --list-group-mappings C:106
- avaki dgas --list-user-mappings C:106
- avaki dgas --modify-cifs-share C:107
- avaki dgas --read-log-properties C:107
- avaki dgas --reset-cache-statistics C:108
- avaki dgas --save-cache C:109
- avaki dgas --self-map C:109
- avaki dgas --self-unmap C:111
- avaki dgas --set-admission-policy-domain C:111
- avaki dgas --set-default-admission-policy C:112
- avaki dgas --set-property C:113

- avaki dgas --show-admission-policies C:114
- avaki dgas --show-admission-policy-domain C:114
- avaki dgas --sync-cache C:115
- avaki dgas --unset-admission-policy-domain C:116
- avaki dgas --unset-property C:117
- Avaki directories, See directories, Avaki
- avaki directory --add-schedule C:117
- avaki directory --cache C:122
- avaki directory --delete-schedule C:122
- avaki directory --do-not-cache C:126
- avaki directory --list-schedules C:126
- avaki domain --create C:127
- avaki domain --disconnect C:127
- avaki domain --info C:128
- avaki domain --interconnect C:128
- Avaki domains, See domains, Avaki
- Avaki EII software
 - overview O:1
 - typical deployment O:17
- avaki executionservice --info C:129
- avaki executionservice --set C:129
- avaki file --add-schedule C:130
- avaki file --cache-on-demand C:134
- avaki file --delete-schedule C:135
- avaki file --do-not-cache C:136
- avaki file --list-schedules C:136
- avaki file --pin C:137
- Avaki functions S:73
- Avaki Functions menu S:74
- avaki group --add --user C:138
- avaki group --create C:141
- avaki group --delete C:143
- avaki group --delete --user C:144
- avaki group --info C:145
- avaki group --list-user C:147
- avaki help C:148
- avaki id C:149
- Avaki installation directory AD:350, API:84, C:308, O:62, P:290, S:176
- avaki ldap --add-schedule C:149
- avaki ldap --delete C:153
- avaki ldap --delete-schedule C:154
- avaki ldap --import C:155
- avaki ldap --info C:157
- avaki ldap --integrate C:157
- avaki ldap --list-schedules C:158
- avaki ldap --searchbase C:160
- avaki ldap --update C:160
- avaki ln C:161
- avaki locks --clear C:163
- avaki locks --list C:164
- avaki login C:164
- avaki logout C:165
- avaki ls C:166
- avaki mkdir C:167
- avaki monitor --add C:167
- avaki monitor --clear C:168
- avaki monitor --create C:170
- avaki monitor --delete C:171
- avaki monitor --list C:172
- avaki monitor --result C:172
- avaki monitor --start C:173
- avaki monitor --stop C:174
- avaki mv command C:176
- avaki nis --add-schedule C:177
- avaki nis --delete C:181
- avaki nis --delete-schedule C:181
- avaki nis --import C:182
- avaki nis --info C:183
- avaki nis --integrate C:184
- avaki nis --list-schedules C:184
- avaki nis --update C:185
- avaki passwd C:185
- avaki permissions C:186
- Avaki perspective in Studio S:13
- avaki plugin command P:184
- avaki plugin --generate C:187
- avaki proxy --add C:191
- avaki proxy --delete C:191
- avaki proxy --list C:192
- avaki pwd C:193
- avaki replica --add C:193
- avaki replica --config command C:193
- avaki replica --delete C:194
- avaki replica --info C:194
- avaki replicate --synch C:195
- avaki rm C:195
- Avaki rowset XML
 - class-name element P:279
 - column-display-size element P:279
 - column-index element P:279
 - core schema P:277
 - rowset-specific schema P:279
 - sample schema P:280
 - schema overview P:277
- avaki schedule --delete C:196
- avaki schedule --info C:197
- avaki schedule --list C:197
- avaki schedule --print-iterations C:198
- avaki scheduleexclusion --create --custom C:198
- avaki scheduleexclusion --create --daily C:199
- avaki scheduleexclusion --create --monthly C:201
- avaki scheduleexclusion --create --weekly C:203
- avaki scheduleexclusion --create --yearly C:205
- avaki scheduleexclusion --delete C:207
- avaki scheduleexclusion --info C:208
- avaki scheduleexclusion --list C:209
- avaki search (execute) C:211
- avaki search --create command C:209
- avaki search --delete C:210
- avaki search --get-rehash-level C:212
- avaki search --info C:214
- avaki search --rehash C:215
- avaki search --set-rehash-level C:215
- avaki security --config C:216
- avaki security --default-gid C:217

- avaki security --default-group C:218
- avaki security --default-uid C:219
- avaki security --default-user C:220
- avaki security --gid C:221
- avaki security --info C:222
- avaki security --uid C:222
- avaki server --dgas --connect C:223
- avaki server --dgas --destroy C:224
- avaki server --dgas --stop C:225
- avaki server --grid --connect C:225
- avaki server --grid --destroy C:226
- avaki server --grid --stop C:227
- avaki server --proxy C:228
- avaki server --share --connect C:228
- avaki server --share --disconnect C:229
- avaki server --share --stop C:230
- Avaki servers
 - distribution of data catalog among O:38
 - hardware and operating system requirements for O:16
 - qualified names for O:32
- avaki share --add-rehash-schedule C:231
- avaki share --add-share-servers C:232
- avaki share --create C:235
- avaki share --delete-rehash-schedule C:236
- avaki share --disconnect C:238
- avaki share --get-local-path C:238
- avaki share --get-status C:239
- avaki share --list-rehash-schedules C:239
- avaki share --list-share-servers C:240
- avaki share --rehash C:240
- avaki share --remove-share-servers C:241
- avaki share --set-local-path C:241
- avaki share --set-share-servers C:242
- avaki share --set-status C:243
- avaki share --update-share-servers C:244
- Avaki shares
 - about AD:257
 - adding schedule exclusions for rehashes C:234
 - adding share servers C:232
 - behavior during failover AD:112
 - bringing on line AD:286
 - changing configuration of AD:266
 - changing encryption levels AD:279
 - changing permissions AD:239
 - changing the owner AD:261
 - configuring exclusions for refresh schedules AD:274
 - copying into, out of, and within AD:213
 - creating AD:258, C:235
 - defined AD:350, API:84, C:308, O:62, P:290, S:176
 - deleting C:195
 - disconnecting C:238
 - disconnecting permanently AD:287
 - forcing refresh AD:262
 - icon for O:29
 - linking AD:217
 - local paths for, obtaining C:238
 - modifying load balance factor C:244
 - moving AD:210
- Avaki shares (*continued*)
 - moving source directories AD:283
 - naming of files and directories in AD:207
 - online status, setting C:243
 - organizing O:37
 - permissions on new files in O:49
 - refresh schedules
 - adding C:231
 - deleting C:236
 - listing C:239
 - refresh schedules for AD:266
 - refreshing C:240
 - removing entries from refresh schedules AD:278
 - removing share servers from AD:265
 - renaming AD:212
 - setting load balancing factor AD:280
 - setting local paths C:241
 - setting names AD:260
 - share servers
 - listing C:240
 - removing C:241
 - replacing C:242
 - shutting down AD:287, C:238
 - status, displaying C:239
 - taking off line AD:285
 - uploading files to AD:282
 - with multiple share servers AD:263
 - write access and user accounts AD:12
 - See also share servers
- avaki shell C:245
- avaki sql view --delete C:246
- avaki sql view --get-description C:246
- avaki sql view --set-description C:247
- avaki status C:248
- Avaki Studio
 - about O:9, O:17, S:1
 - Avaki perspective, about S:16
 - defined AD:350, API:84, C:308, O:62, P:290, S:176
 - getting started S:9
 - installing in Windows AD:24
 - limitations of data services created in P:78, S:3
 - log properties file for AD:317
 - metadata models, See metadata models
 - operators S:5, S:107
 - projects, creating S:13
 - requirements for running AD:3
 - setting system properties for AD:129
 - starting S:9
 - time required to upgrade AD:341
 - view models
 - about S:2
 - configuring input sources S:43
 - creating S:29
 - deploying as data services S:50
 - sample workflow for S:29
 - testing S:49
 - workflow S:25

avaki upgrade C:249
 avaki upgrade --info C:250
 avaki user C:250
 avaki user --create C:251
 avaki user --db-mapping --add C:252
 avaki user --db-mapping --delete C:253
 avaki user --db-mapping --list C:255
 avaki user --delete C:257
 avaki user --info C:258
 avaki user --list-group C:258
 avaki view --add-schedule C:259
 avaki view --create --database C:263
 avaki view --create --data-service C:266
 avaki view --create --file C:267
 avaki view --delete C:272
 avaki view --delete-schedule C:272
 avaki view --depends C:272
 avaki view --garbage-collect C:273
 avaki view --info C:274
 avaki view --list-schedules C:274
 avaki view --regenerate C:273
 avaki view --set-property C:275
 avaki view --update C:279
 avaki virtualdatabase --allow-dbop-creation C:280
 avaki virtualdatabase --disallow-dbop-creation C:281
 avaki virtualdatabase --execute C:282
 avaki virtualdatabase --show-tables C:283
 avaki virtualeschema --deploy C:285
 avaki virtualeschema --undeploy C:286
 avaki whoami C:286
 Avaki_JDBCStandAlone.jar file API:66
 Avaki_JDBCStandAlone_Minus3rd.jar file API:66
 AvakiAPI.disco file API:3
 AvakiAPIDocLit.wsdl file API:3
 AvakiAPIRpcEnc.wsdl file API:3
 AvakiAPIWithMIMEDocLit.wsdl file API:3
 AvakiAPIWithMIMERpcEnc.wsdl file API:3
 avakijdbc.properties file API:67
 AvakiPrincipal complex type API:13
 Axis, See Apache Axis

B

backup command C:23
 backups on Avaki servers AD:113
 batch mode, JDBC API:77
 configuring database operations for P:27, P:250
 bindings.xml file
 copying during upgrade AD:344
 on grid servers AD:50
 on proxy servers AD:300
 on share servers AD:59
 block size file attribute C:290
 blocks file attribute C:290
 bootstrapping
 in Unix AD:16
 in Windows AD:23
 brackets, See curly brackets, square brackets, angle brackets

browsers, See web browsers
 build.xml file for data service plug-ins P:184, P:197
 BusinessObjects software unable browse Avaki objects with
 underscores API:81

C

cache --evict --all command C:25
 cache --evict command C:24, C:32
 cache --evict --deleted command C:26
 cache --get command C:27
 cache --invalidate --all command C:28
 cache --invalidate command C:27
 cache --invalidate-dataservice-results command C:29
 cache --invalidate-dbop-results command C:30
 cache --list command C:31
 cache services
 about P:119
 associating with data grid access servers P:113
 associating with grid servers P:111
 coherence windows P:107
 configuring P:116
 configuring per file P:117
 defined AD:351, API:85, C:309, O:63, P:291, S:177
 disassociating from data grid access servers P:114
 disassociating from grid servers P:112
 evicting cached files and directories P:135
 invalidating cached items P:136
 listing P:116
 listing cached data services P:163
 listing cached database operations P:148
 listing cached virtual database operations P:148
 listing pinned files and directories P:134
 on-demand caching P:119
 on-demand caching of database operation and data service
 results P:108
 on-demand caching of files P:107
 overriding default settings P:117
 pinning data services P:152
 pinning database operations P:139
 pinning files and directories P:120
 pinning virtual database operations P:139
 scheduled caching P:119
 scheduled caching of database operation and data service
 results P:109
 scheduled caching of files P:107
 tagging files and directories P:129
 unmarking cached items P:135
 unscheduling cached files and directories P:135
 viewing details about P:116
 See also caches and caching
 cache --set command C:33
 cache --unset command C:34
 caches
 adding schedules for data services C:48
 adding schedules for database operations C:73
 adding schedules for directories C:117
 adding schedules for files C:130

caches (*continued*)

- bad port, properties for AD:141
- configuring associated server or user C:33
- configuring threads for AD:131
- data service plug-in, properties for AD:137
- deleting schedules for data services C:54
- deleting schedules for database operations C:78
- deleting schedules for directories C:122
- deleting schedules for files C:135
- DGAS
 - clearing user credentials AD:117
 - clearing user credentials from C:95
 - configuring block size for reads AD:81
 - configuring frags per block for reads AD:81
 - configuring location of AD:73
 - controlling cache size AD:124
 - deleting files and directories AD:119
 - deleting objects from C:98
 - displaying cache statistics C:102
 - displaying current size C:101
 - displaying free disk space on cache machine C:102
 - forcing a refresh AD:121
 - managing AD:117
 - mapping cache AD:336
 - resetting statistics C:108
 - saving a copy C:109
 - saving copies AD:120
 - setting remote caches for AD:90
 - syncing AD:121
 - viewing and resetting statistics AD:123
 - warming and updating C:115
- displaying associated server or user C:27
- displaying tagging information C:31
- for tables in virtual database, property for AD:144
- listing schedules for data services C:58
- listing schedules for database operations C:88
- listing schedules for directories C:126
- listing schedules for files C:136
- local S:182
- local and remote API:63
- marking directories for no caching C:126
- marking files for no caching C:136
- remote object stub, properties for AD:144
- schedule exclusion, properties for AD:144
- scheduled caching S:184
- settable attributes of C:291
- setting invalidate queue for AD:135
- setting local directory for AD:135
- setting remote caches for command clients AD:95
- uncoupling associated server or user C:34

See also cache services and caching

caching

- about O:13
- and JDBC programs O:55
- benefits to performance O:54
- configuring ACLs for O:50

caching (*continued*)

- configuring Avaki clients for O:55
- data service results P:108, S:51
 - tagging for on-demand caching P:159
- database operations P:108
- defined AD:351, API:85, C:309, O:63, P:291, S:177
- files O:56, P:107
- JDBC and caching of database operation results API:62
- local AD:356, API:90, C:314, O:14, O:53, O:68
- local vs. remote O:59
- local, defined P:296
- of data service results O:57
- of database operation results O:57
- on DGAS O:54
- on-demand AD:357, API:91, C:315, O:69, P:296, S:183
- permissions and access control O:59
- remote O:14, O:53
 - defined AD:357, API:91, C:315, O:69, P:297, S:184
- scheduled AD:357, API:91, C:315, O:69
- scheduled, defined P:297
- turning off for specified files and directories P:132
- See also caches and cache services
- callable statements API:72, API:73
- case sensitivity in Avaki naming AD:206
- cat command C:35
- catalog browser S:18
- categories
 - about AD:221
 - adding objects to AD:226, C:35, S:105
 - adding SQL views P:47
 - and permissions AD:222–AD:223
 - browsing AD:222
 - contents of S:18
 - creating AD:224, C:36
 - default, contents of S:18
 - defined AD:351, API:85, C:309, O:63, P:291, S:177
 - deleting AD:230, C:37, S:106
 - displaying S:104
 - for logging AD:318
 - icon for O:29
 - listing categories in domain C:38
 - managing S:103
 - permissions in O:48
 - removing objects from AD:228, C:38, C:40, S:106
 - setting descriptions for C:39
 - showing descriptions C:37
 - using to organize data O:36
 - using to solve access problems O:49
- categories --add-to-category command C:35
- categories --create command C:36
- categories --delete command C:37
- categories --describe command C:37
- Categories directory O:35
- categories --list command C:38
- categories --remove-from-category command C:38, C:40
- categories --set-description command C:39
- cd command C:41
- certificates, SSL, See SSL certificates

- change time file attribute C:290
- characters
 - in column aliases in database operations S:42
 - in command syntax AD:xv, C:xv
 - in cron schedules C:298
 - in domain names, restrictions on AD:41
 - in JavaScript identifiers S:42
 - in metadata model names, restrictions on S:91
 - in names of Avaki objects, restrictions on AD:207
 - wildcards in searches AD:235
- CHARSET JDBC property for ASE and IQ AD:7
- chmod command C:41
- chmod SOAP operation API:19
- chown command C:42
- chown SOAP operation API:20
- chunk size for sorting, controlling AD:139, S:76
- CIFS
 - accessing data grid through AD:338
 - releasing CIFS ports on a DGAS AD:66
- CIFS clients
 - defined AD:351, API:85, C:309, O:63, P:291, S:177
 - disconnecting C:101
 - displaying connected clients C:105
 - requirements for O:17
 - setting up AD:93
- CIFS shares
 - accessing AD:203
 - creating AD:125, C:96
 - defined AD:351, API:85, C:309, O:63, P:291, S:177
 - deleting C:99
 - displaying connected clients C:105
 - displaying information about C:95
 - listing C:105
 - managing AD:125
 - mapping to a network drive AD:204
 - modifying C:107
- class element P:261
- class-name element P:279
- classpath, configuring for JDBC drivers API:67
- client attribute caching AD:336
- client command C:45
- client system properties AD:128
- clients
 - about O:17
 - hardware and operating system requirements for O:16
 - message timeout properties for AD:133
 - setting size of write invalidation queue of cache for AD:136
 - setting system properties for AD:129
 - See also Avaki Studio, CIFS clients, command clients, NFS clients, web clients, WS clients
- code samples
 - ad-hoc queries API:74
 - data catalog API API:18
 - data services API API:35
 - database operations API API:42
 - Java data service plug-ins P:190
 - JDBC batch mode API:77
 - using JDBC drivers API:77
- coherence window cache attribute C:291
- coherence window property, remote AD:141
- coherence windows for caching P:107
- coherenceWindow element P:261
- colors in Studio display, setting S:23
- column-display-size element P:279
- column-index element P:279
- columns
 - aliasing P:25
 - combining with Projection operator S:46
 - from input elements, menus of S:71
 - from input result sets, accessing S:68
 - name property S:60
 - precision property S:61
 - scale property S:61
 - type property S:61
- com.avaki.badPortCacheSize system property AD:141
- com.avaki.badPortExpiration system property AD:141
- com.avaki.cache.cacheDir system property AD:135
- com.avaki.cache.maxReaderThreads system property AD:131
- com.avaki.cache.writeInvalidationQueueSize system property AD:136
- com.avaki.content.encryptionLevel system property AD:139
- com.avaki.dataservice.pluginCacheSize system property AD:137
- com.avaki.dataservice.styleSheetCachePoolSize system property AD:137
- com.avaki.dataservice.styleSheetCacheSize system property AD:137
- com.avaki.DBOProtocolSoTimeout system property AD:134
- com.avaki.generatedXMLIndentSize AD:142
- com.avaki.HttpPort system property AD:140
- com.avaki.HttpsPort system property AD:140
- com.avaki.jobStatusExpiration system property AD:145
- com.avaki.lasInvoker.cacheSize system property AD:144
- com.avaki.lasInvoker.poolSize system property AD:144
- com.avaki.ldap.resultPageSize system property AD:145
- com.avaki.maxActiveCachables system property AD:136
- com.avaki.mux.channelSoTimeout system property AD:135
- com.avaki.mux.connectTimeout system property AD:134
- com.avaki.mux.maxParallelChannels system property AD:142
- com.avaki.mux.maxWriteChunk system property AD:142
- com.avaki.mux.sendBufferSize system property AD:143
- com.avaki.proxy.retryDelay system property AD:133
- com.avaki.proxy.retryTimeout system property AD:133
- com.avaki.proxyIOProtocolSoTimeout system property AD:134
- com.avaki.proxyKeepAliveParams system property AD:140
- com.avaki.queryEngine.sortChunkSize AD:139
- com.avaki.remoteconfig.coherenceWindow system property AD:141
- com.avaki.result.gcInterval system property AD:136
- com.avaki.retryDelay system property AD:133
- com.avaki.retryTimeout system property AD:133
- com.avaki.rmiRegistrySoTimeout system property AD:134
- com.avaki.rpcTimeout system property AD:134
- com.avaki.scheduleExclusionCacheExpiration system property AD:145

- com.avaki.scheduleExclusionCacheSize system property AD:145
- com.avaki.shareIOProtocolSoTimeout system property AD:134
- com.avaki.shareReadBufferSize system property AD:138
- com.avaki.shareReadbufPoolSize system property AD:138
- com.avaki.shareServerCircularLinkChecking system property AD:138
- com.avaki.shareServerThreadPoolSize system property AD:138
- com.avaki.vaultStateCacheSize system property AD:137
- com.avaki.VirtualDbTableCacheSize system property AD:144
- com.sybase.avaki.tdsPort system property AD:50, AD:145, API:71
- command clients
 - connecting C:45
 - defined AD:352, API:86, C:310, O:64, P:292, S:178
 - disconnecting C:45
 - installing in Windows AD:24
 - installing on Unix AD:18
 - obtaining information about C:45
 - setting up AD:94
- commands
 - listing C:148
 - syntax conventions for AD:xv, C:xiv, P:x
 - viewing online usage information C:148
- compatibility properties, setting for Windows 2003 AD:22
- complex types API:12
 - AdHocDBOPExecutionParams API:13
 - AvakiPrincipal API:13
 - DataCatalogAttribute API:14
 - DataCatalogEntry API:15
 - DataCatalogPermission API:15
 - DataServiceExecutionParams API:16
 - DBOPExecutionParams API:16
 - SearchQuery API:17
 - SearchResult API:17
- condition field for Iterator operators S:151
- connect ports
 - default AD:6, AD:9, AD:10
 - defined AD:352, API:86, C:310, O:64, P:292, S:178
 - for DGAS C:224
 - changing AD:74
 - for GDCs C:127
 - changing AD:50
 - for grid servers C:226
 - changing AD:50
 - for proxy servers C:228
 - changing AD:300
 - for share servers C:229
 - changing AD:60
- connectinfo.txt file AD:131
- connection pooling S:36
- connection properties
 - for JDBC drivers API:68
 - for XA drivers C:70, P:7, S:37
- connection strings
 - for databases AD:3
 - for JDBC drivers API:71

- connections in view models, creating S:57
- console view S:22, S:50
- conventions
 - for command syntax C:xiv
 - for commands AD:xv
 - for screen examples AD:xv, API:vi, C:xv, P:xi
- cp command C:46
- cron expressions in schedules AD:185, AD:267, AD:273
- cron schedules
 - configuring C:297
 - values for C:298
- cross-domain messaging
 - disabling AD:313
 - enabling AD:311
- curly brackets in command syntax AD:xv, C:xv, P:xi
- CurrentUser functions S:74
- Custom operator S:111
 - example S:114
- custom types API:12

D

- data access O:11
 - using WS API API:2
- data catalog
 - about O:27
 - defined AD:352, API:86, C:310, O:64, P:292, S:178
 - distribution among Avaki servers O:38
 - names of objects in O:24
 - organizing O:33
 - Avaki shares O:37
 - using categories O:36
 - using links O:36
 - ownership of objects in O:46
 - top-level directories O:32
 - types of entries O:6
- data catalog API API:18
- data catalog SOAP operations API:18
 - accessiblePath API:19
 - chmod API:19
 - chown API:20
 - fileRead API:21
 - fileReadAttach API:21
 - fileReadString API:22
 - fileWrite API:23
 - getAttributes API:23
 - getSystemAttributes API:24
 - getUserAttributes API:24
 - listDomains API:25
 - listSearches API:25
 - ls API:26
 - lsSize API:26
 - mkdir API:27
 - mkdirParents API:27
 - mkdirParentsServer API:28
 - mkdirServer API:29
 - mv API:29
 - permissions API:30

- data catalog SOAP operations (*continued*)
 - removeAttribute API:31
 - rm API:31
 - search API:32
 - setAttribute API:32
 - tester API:33
 - whoami API:33
- data catalog view S:18
- data expiration intervals P:108, S:51
- data grid access servers
 - associating with cache services P:113
 - disabling auto-restart C:8
 - disassociating from cache services P:114
 - enabling auto-restart C:5
 - registering C:5
 - starting C:5, C:6
 - stopping C:7, C:8
 - unregistering C:8
 - See also DGAS
- data grids
 - about O:1
 - defined AD:355, API:89, C:313, O:67, P:294, S:181
 - typical deployment O:17
- data integration O:21, O:23
- data integrity and HTTPS API:8
- data representation O:11
- data security O:10
- data service plug-ins
 - about P:76, P:175
 - addInputParameter JavaScript method P:202
 - addInputStream JavaScript method P:202
 - build.xml file P:184, P:197
 - choice of Java, JavaScript, or XSLT P:176
 - closing streams P:186
 - command for generating C:187
 - configuring P:81
 - creating in Java with the Plug-in Wizard P:183
 - creating in JavaScript P:200
 - creating in XSLT P:180
 - DbopGroupWorkUnit class P:189
 - DbopPipeWorkUnit class P:190
 - defined AD:352, API:86, C:310, O:64, P:292, S:179
 - examples
 - DBOP and CSV merge Java plug-in P:193
 - distributed transaction Java plug-in P:191
 - rowset input and output Java plug-in P:192
 - Execute JavaScript function P:203
 - input sources and output streams P:177
 - InputSource interface P:186
 - JAR files for P:180
 - logging errors P:196
 - manifest files for P:180, P:197
 - modularity and reusability of P:175
 - parameters
 - about P:178
 - specifying for Java plug-ins C:188
 - specifying for XSLT plug-ins P:181
- data service plug-ins (*continued*)
 - ParameterSpec interface P:187
 - Plugin interface P:186
 - prerequisites for writing in Java P:183
 - relationship to .js files in Studio S:11
 - RowSetFactory class P:188
 - setOutputStream JavaScript method P:202
 - StreamingRowSet interface P:187
 - using Java classes and interfaces in JavaScript plug-ins P:200
 - when to use P:78, S:3
 - XAWorkHandler class P:189
 - XAWorkUnit interface P:189
- data service XML schema
 - class element P:261
 - coherenceWindow element P:261
 - dataService element P:262
 - description element P:263
 - initParameter element P:263
 - inputParameter element P:264
 - inputSource element P:265
 - inputStream element P:266
 - isList element P:266
 - jarurl element P:267
 - logicBox element P:268
 - name element P:269
 - outputStream element P:269
 - ref element P:270
 - target element P:270
 - type element P:270
 - urlLogicBox element P:271
 - value element P:272
 - values element P:272
- data services
 - about O:8, O:23, P:49, P:74
 - adding schedule exclusions C:51
 - and distributed transactions P:78
 - caching of results O:57, P:77
 - permissions O:50
 - caching results S:51
 - calling via JDBC API:72
 - components of P:76
 - configuring caching P:108
 - created in Avaki Studio, limitations of P:78, S:3
 - creating C:52, P:80, P:207
 - defined AD:352, API:86, C:310, O:64, P:292, S:178
 - deleted, purging from cache C:26
 - deleting schedules C:54
 - dependencies for S:22
 - deploying from view models in Avaki Studio S:50
 - displaying dependency lists C:54
 - displaying information about C:57, S:20
 - displaying status of C:248
 - evicting from cache P:164
 - execution services, configuring AD:109
 - generating schema for C:55, P:98
 - generating SQL views from C:56, P:100
 - importing descriptors P:92
 - input parameters, configuring P:84

- data services (*continued*)
 - input streams, configuring P:87
 - invalidating all in cache C:28
 - invalidating one in cache C:27
 - invalidating results in cache C:29
 - listing P:93
 - listing caching schedules for C:58
 - listing in cache P:163
 - location in categories S:18
 - marking for scheduled caching P:152
 - modifying C:58
 - modifying permissions AD:239
 - modifying settings P:94
 - names in data catalog O:24
 - nesting operations S:149
 - output streams, configuring P:86
 - provisioning web services as P:205
 - purging all from cache and unspooling C:25
 - purging one from cache and unspinning C:24
 - qualified names for O:31
 - refreshing cached results C:32
 - removing P:103
 - rowsets as input of P:275
 - rowsets as output of P:274
 - running C:55
 - sample workflow for S:29
 - scheduling for caching C:48
 - schema P:257
 - searching for AD:233
 - setting cache sizes for plug-ins AD:137
 - setting up to run distributed transactions P:80
 - specifying grid servers P:213
 - specifying input parameters P:207
 - specifying input streams P:208
 - specifying output streams P:208
 - specifying plug-ins P:207
 - tagging for on-demand caching P:159
 - testing P:102, P:214
 - unscheduling P:164
 - using for distributed transactions O:25
 - viewing P:98
 - viewing dependencies P:97
 - writing your own vs. using Avaki Studio O:24
 - See also data service plug-ins
 - See also view models
- data services API API:34
- data services SOAP operations API:34
 - accessibleDS API:36
 - executeDS API:36
 - getDSOutput API:37
 - getDSOutputAttach API:38
 - getDSOutputString API:38
 - getDSParameters API:39
 - isDSAvakiXML API:40
 - listDSs API:40
- data structures, SOAP complex types API:12
- data type mappings for SQL views P:39
- data types
 - for JDBC API:76
 - mapping
 - about type mapping files C:301
 - command to specify mapping file C:68
 - format of type mapping files C:301
 - inconsistencies C:302
 - logging of mapping decisions C:304
 - setting source data type C:302
 - specifying for ad-hoc query parameters C:65
 - specifying for database operation parameters C:85
 - specifying for parameters for ad-hoc queries on the virtual database C:283
 - specifying for virtual database operation parameters C:88
 - See also type
- database connectors
 - about O:22, P:3
 - adding groups P:16
 - adding users P:15
 - configuring advanced settings P:247
 - configuring JDBC driver JAR file path P:247
 - configuring permissions C:59, C:62
 - creating C:66, P:3, S:31
 - data type mappings for P:39
 - defined AD:353, API:87, C:311, O:65, P:293, S:179
 - deleting C:61
 - displaying information about C:64
 - displaying SQL views provisioned from C:73
 - editing S:38
 - executing ad-hoc queries C:63, C:66
 - finding in catalog S:38
 - getting information about through JDBC API:75
 - JDBC fetch size P:5
 - location in categories S:18
 - managing SQL views P:20
 - modifying C:66, P:8
 - provisioning SQL views from C:71
 - removing P:21
 - removing groups P:18
 - removing users P:18
 - searching for AD:233
 - setting JDBC fetch size S:36
 - testing C:72, P:19
 - viewing P:8
 - viewing associated database operations P:13
- database drivers
 - copying during upgrade AD:341
 - tested with Avaki AD:3
- database identity mappings P:6, S:36
 - about AD:176
 - adding C:252
 - deleting AD:183, C:253
 - displaying AD:180
 - listing C:255
 - modifying AD:182
 - setting up AD:177
- database operation --list-schedules command C:88

database operation SOAP operations API:40

database operations

about O:7, O:22, P:1

access permissions P:22

adding schedule exclusions C:76

allowing groups to create P:16

allowing users to create P:15

caching of results API:62, O:57

permissions O:50

calling with JDBC API:73

calling with ODBC, JDBC, or SOAP P:38

configuring advanced settings P:247

configuring batch mode settings P:250

configuring caching P:108

configuring permissions C:59, C:62

configuring SQL calls P:251

configuring timeouts P:253

creating C:83, P:22, S:38

defined AD:353, API:87, C:311, O:65, P:293, S:179

deleted, purging from cache C:26

deleting C:78

deleting schedules C:78

dependencies for S:22

displaying dependency lists C:79

displaying information about C:83, S:20

displaying status of C:248

evicting from cache P:150

executing P:36

exposing results as SQL view P:34

generating schema for C:80, P:31

generating SQL views from C:82

invalidating all in cache C:28

invalidating one in cache C:27

invalidating results in cache C:30

listing caching schedules for C:88

listing in cache P:148

location in categories S:18

managing P:21

managing metadata P:30

marking for scheduled caching P:139

modifying C:83, P:28

modifying permissions AD:239

names in data catalog O:24

parameter types, specifying C:85, C:88

preventing groups from creating P:18

preventing users from creating P:18

purging all from cache and unscheduling C:25

purging one from cache and unpinning C:24

qualified names for O:31

refreshing cached results C:32

removing P:38

removing SQL views generated from P:35

restricting row output P:248

rowsets as output of P:274

running C:80

sample XML schema P:280

scheduling for caching C:73

searching for AD:233

database operations (*continued*)

setting JDBC fetch size for P:254

setting permissions P:14

SQL statements in C:86

tagging for on-demand caching P:146

transactional behavior of P:79

unscheduling P:150

uses of P:2

viewing P:13, P:28

viewing dependencies P:32

viewing details about P:29

See also virtual database operations

database operations API API:40

database service SOAP operations

accessibleDBOp API:42

executeAdHocDBOp API:43

executeAdHocDBOpWithOutput API:44

executeAdHocDBOpWithOutputAttach API:46

executeAdHocDBOpWithOutputString API:47

executeDBOp API:48

executeDBOpBytesInput API:49

executeDBOpGridFileInput API:50

executeDBOpWithOutput API:50

executeDBOpWithOutputAttach API:52

executeDBOpWithOutputString API:53

getDBOpOutput API:54

getDBOpOutputAttach API:55

getDBOpParameters API:56

getDBOpSchema API:56

getDBOpSchemaAttach API:57

getDBOpSchemaString API:58

getOutputString API:55

getSQL API:58

listDBConns API:59

listDBOps API:59

listDBOpsByDBConn API:60

database, virtual, See virtual database

DatabaseAdministrators group O:44

DatabaseMetaData interface API:75

databases

Avaki tools for working with O:21

connecting to P:3

for Avaki servers, backing up C:23

protecting O:9

schemas, viewing P:9

supported for connecting to Avaki AD:3, AD:5

DataCatalogAttribute complex type API:14

DataCatalogEntry complex type API:15

DataCatalogPermission complex type API:15

DataProviders group O:44

dataservice --add-schedule command C:48

dataservice --create command C:52

dataservice --delete command C:52

dataservice --delete-schedule command C:54

dataservice --depends command C:54

dataService element P:262

dataservice --execute command C:55

dataservice --generate-sql view command C:56

- dataservice --info command C:57
- dataservice --list-schedules command C:58
- dataservice --update command C:58
- DataServiceExecutionParams complex type API:16
- DB2, versions and JDBC drivers for use with Avaki AD:6
- dbconn --allow-dbop-creation command C:59
- dbconn --delete command C:61
- dbconn --disallow-dbop-creation command C:62
- dbconn --execute command C:63
- dbconn --info command C:64
- dbconn --jdbc command C:66
- dbconn --provision-tables command C:71
- dbconn --show-tables command C:73
- dbconn --test command C:72
- DBOPs, See database operations
- dbop --add-schedule command C:73
- dbop --delete command C:78
- dbop --delete-schedule command C:78
- dbop --depends command C:79
- dbop --execute command C:80
- dbop --generate-sql view command C:82
- dbop --info command C:83
- dbop --jdbc command C:83
- dbop --jdbc --create-virtual-dbop command C:87
- DBOPExecutionParams complex type API:16
- DbopGroupWorkUnit class for data services P:189
- DbopPipeWorkUnit class for data services P:190
- db-path option (DGAS) AD:74
- debug mode, enabling in an Avaki shell C:245
- delimiter character for JDBC schema names API:69
- Deny permission in ACL AD:243, O:47
- dependencies S:22
 - defined AD:353, API:87, C:311, O:65, P:293, S:179
 - listing for data services C:54, P:97
 - listing for database operations C:79, P:32
 - listing for view generators C:272, P:228
 - listing for virtual database operations P:59
- description element P:263
- descriptors for data services P:92
- development tools for web services
 - Apache Axis API:5
 - Microsoft Visual Studio API:5
 - SOAP::Lite API:5
 - VB .NET API:5
- DGAS
 - about AD:62
 - adding user self mappings C:109
 - admission policies AD:332
 - about AD:85
 - adding AD:87, C:91
 - deleting AD:88, C:97
 - displaying C:114
 - displaying Windows domains for C:114
 - admission policies
 - setting defaults C:112
 - setting Windows domains for C:111
 - unsetting Windows domains for C:116
- DGAS (continued)
 - changing permissions and ownership AD:334
 - CIFS access to data grid AD:338
 - clearing cached credentials AD:117, C:95
 - configuring associated cache service C:33
 - configuring location of internal caches AD:73
 - configuring to use nondefault ports AD:74, AD:75
 - configuring users and groups AD:67
 - connect port C:224
 - connecting to a domain AD:79, C:223
 - controlling cache size AD:124
 - create CIFS shares C:96
 - default name for AD:73
 - default users, groups, UIDs and GIDs AD:333
 - defined AD:352, API:86, C:310, O:64, P:292, S:178
 - deleting cached objects C:98
 - deleting CIFS shares C:99
 - deleting files and directories from cache AD:119
 - deleting user mappings C:111
 - destroying C:224
 - disconnecting CIFS clients C:101
 - displaying associated cache C:27
 - displaying cache size C:101
 - displaying cache statistics C:102
 - displaying connected CIFS clients C:105
 - displaying free disk space on cache machine C:102
 - displaying information about CIFS shares C:95
 - displaying property descriptions C:104
 - displaying property values C:103
 - dynamic and nondynamic properties AD:83
 - file locking in Unix, interference with AD:64
 - forcing cache to refresh AD:121
 - installing in Unix AD:18
 - installing in Windows AD:24
 - listing CIFS shares C:105
 - listing properties and their values C:103
 - managing cache AD:117
 - mappings
 - default, adding and displaying AD:71
 - for groups, adding AD:70, C:92
 - for groups, deleting C:99
 - for groups, displaying C:106
 - for users, adding AD:70, C:94
 - for users, deleting C:100
 - for users, displaying C:106
 - users and groups, per-DGAS AD:88
 - users, groups, and defaults, domain-wide AD:68
 - modifying CIFS shares C:107
 - NFS clients, not running with AD:64
 - NFS daemons, shutting down before running DGAS AD:66
 - per-DGAS user mappings AD:333
 - ports used by AD:9
 - preparing to start AD:65
 - properties file for AD:76, C:293
 - reading log properties C:107
 - releasing CIFS ports before running AD:66
 - resetting cache statistics C:108

DGAS (continued)

restarting AD:84
 saving a copy of the cache C:109
 saving copies of cache AD:120
 server logs AD:317
 setting a cache service AD:90
 setting location of state database AD:74
 setting properties AD:82, C:113
 setting up NFS clients AD:91
 starting AD:73, C:104
 stopping C:225
 syncing cache AD:121
 time required to upgrade AD:341
 uncoupling associated cache C:34
 unsetting properties C:117
 viewing and resetting cache statistics AD:123
 warming and updating the cache C:115
 See also data grid access servers
 dgas --add-admission-policy command C:91
 dgas --add-group-mapping command C:92
 dgas --add-user-mapping command C:94
 dgas --cifs-share-info command C:95
 dgas --clear-cached-credentials command C:95
 dgas command
 example AD:74
 syntax AD:73
 dgas --create-cifs-share command C:96
 dgas --delete-admission-policy command C:97
 dgas --delete-cache command C:98
 dgas --delete-cifs-share command C:99
 dgas --delete-group-mapping command C:99
 dgas --delete-user-mapping command C:100
 dgas --disconnect-cifs-client command C:101
 dgas --get-cache-size command C:101
 dgas --get-cache-statistics command C:102
 dgas --get-free-disk-space command C:102
 dgas --get-properties command C:103
 dgas --get-property command C:103
 dgas --get-property-list command C:104
 dgas --initialize command C:104
 dgas --list-cifs-clients command C:105
 dgas --list-cifs-shares command C:105
 dgas --list-group-mappings command C:106
 dgas --list-user-mappings command C:106
 dgas --modify-cifs-share command C:107
 dgas --read-log-properties command C:107
 dgas --register command C:5
 dgas --reset-cache-statistics command C:108
 dgas --save-cache command C:109
 dgas --self-map C:109
 dgas --self-unmap C:111
 dgas --set-admission-policy-domain command C:111
 dgas --set-default-admission-policy command C:112
 dgas --set-property command C:113
 dgas --show-admission-policies command C:114
 dgas --show-admission-policy-domain command C:114
 dgas --start command C:6

dgas --stop command C:7
 dgas --sync-cache command C:115
 dgas --unregister command C:8
 dgas --unset-admission-policy-domain command C:116
 dgas --unset-property command C:117
 directories
 adding schedule exclusions for caching C:121, C:125
 adding to cache service C:122
 Avaki directories, defined AD:349, API:83, C:307, O:61,
 P:289, S:175
 Avaki installation API:84, P:290, S:176
 Avaki installation directory AD:350, C:308, O:62
 changing C:41
 changing ownership C:42
 changing permissions for AD:239
 copying AD:213, C:46
 creating AD:208, C:167
 deleted, purging from cache C:26
 deleting AD:219, C:195
 deleting caching schedules C:122
 displaying name of current directory C:193
 evicting from cache P:135
 exporting from the data grid AD:125
 home, creating AD:169
 icon for O:29
 invalidating all in cache C:28
 invalidating from cache P:136
 invalidating in cache C:27
 linking AD:217, C:161
 listing C:166
 listing schedules C:126
 listing those pinned for caching P:134
 marking for no caching C:126, P:132
 marking for scheduled caching P:120
 moving AD:210, C:176
 NFS-mounting AD:92
 permissions in O:48
 purging all from cache and unpinning C:25
 purging from cache and unpinning C:24
 refreshing in cache C:32
 renaming AD:212
 scheduling for caching C:117
 searching for AD:233
 setting ACLs for AD:171
 shared, See Avaki shares
 tagging for on-demand caching P:129
 temp, setting for grid servers AD:135
 top-level, described O:32
 unscheduling from cache P:135
 directory --add-schedule command C:117
 directory --cache command C:122
 directory --delete-schedule command C:122
 directory --do-not-cache command C:126
 directory --list-schedules command C:126
 disk space
 available, displaying for DGAS cache C:102
 requirements for Avaki software AD:4, AD:16

- distributed transactions
 - about O:25, P:78
 - API for executing P:188
 - configuring database connectors for P:7, S:36
 - defined AD:353, API:87, C:311, O:65, P:293, S:180
 - requirements for P:79
 - setting up P:80
 - supported DBMSes P:79
 - two-phase commit protocol P:79
 - DNS aliases for GDC machines AD:10
 - DNS name, setting for a server AD:32
 - document/literal web services API:3, API:5
 - documentation
 - Avaki, list of AD:xii, API:vi, C:xii, O:vi, P:viii, S:viii
 - for Eclipse Workbench S:12
 - domain --create command C:127
 - domain --disconnect command C:127
 - domain --info command C:128
 - domain --interconnect command C:128
 - domain names, defined AD:354, API:88, C:312, O:66, P:294, S:180
 - Domain Users group in Active Directory C:155
 - domains, Avaki
 - creating C:127
 - defined AD:350, API:84, C:308, O:62, P:290, S:176
 - disconnecting C:127
 - displaying information about AD:98
 - getting information about through JDBC API:75
 - interconnecting AD:289, C:128
 - joining together AD:289
 - naming AD:41, AD:354, API:88, C:312, O:66, P:294, S:180
 - obtaining information about C:128
 - planning before install AD:1
 - providers and consumers AD:289
 - remote, logging in to AD:201
 - specifying for JDBC connections API:69
 - DomainUsers group O:44
 - downstream variables menu S:71
 - downstream, defined S:3
 - drivers, See JDBC drivers
 - drivers directory AD:341
 - dynamic and nondynamic properties (DGAS) AD:83
 - dynamic user mappings
 - creating C:109
 - deleting C:111
- E**
- Eclipse Workbench S:12
 - EII, See enterprise information integration
 - elements
 - connecting S:57
 - descriptions of S:59
 - Input Source S:125
 - moving S:56
 - elements (*continued*)
 - names of S:58
 - operators S:5
 - properties dialogs S:58
 - Result S:164
 - selecting S:56
 - with red borders S:60
 - encryption and HTTPS API:8
 - encryption levels for Avaki shares
 - changing AD:279
 - displaying C:222
 - setting at share creation AD:261
 - encryption of grid objects AD:139
 - enterprise information integration, defined AD:354, API:88, C:312, O:66, P:294, S:180
 - equal signs in command syntax AD:xvi, C:xvi, P:xii
 - error handling S:143
 - errors in view models S:60
 - everyone group O:44, O:48
 - examples
 - conventions used in C:xv
 - data catalog web service API:18
 - data services API API:35
 - database operations API API:42
 - web services clients
 - Java API:9
 - Perl API:9
 - VB .NET API:9
 - exclusions, See schedule exclusions
 - execute inputs in parallel field for Iterator operators S:151
 - Execute JavaScript function for data service plug-ins P:203
 - executeAdHocDBOp SOAP operation API:43
 - executeAdHocDBOpWithOutput SOAP operation API:44
 - executeAdHocDBOpWithOutputAttach SOAP operation API:46
 - executeAdHocDBOpWithOutputString SOAP operation API:47
 - executeDBOp SOAP operation API:48
 - executeDBOpBytesInput SOAP operation API:49
 - executeDBOpGridFileInput SOAP operation API:50
 - executeDBOpWithOutput SOAP operation API:50
 - executeDBOpWithOutputAttach SOAP operation API:52
 - executeDBOpWithOutputString SOAP operation API:53
 - executeDS SOAP operation API:36
 - execution services
 - about AD:109, P:77, S:51
 - configuring AD:109, C:129
 - configuring threads for AD:131
 - defined AD:354, API:88, C:312, O:66, P:294, S:180
 - displaying information about C:129
 - executionservice --info command C:129
 - executionservice --set command C:129
 - executionServiceHint JDBC property API:70
 - exiting from an Avaki session C:165
 - expressions in operators S:4
 - expressions menu, using S:71
 - externalCacheService JDBC property API:62, API:70

F

- failover
 - defined AD:354, API:88, C:312, O:66, P:294, S:180
 - managing AD:112
 - setting up a secondary GDC AD:43
- fake_metadata JDBC connection property for ASE API:70
- FAKE_METADATA JDBC property for ASE AD:7
- federated data access AD:354, API:88, O:66, P:294, S:181
- fetch size, See JDBC fetch size
- file --add-schedule command C:130
- file --cache-on-demand command C:134
- file --delete-schedule command C:135
- file --do-not-cache command C:136
- file --list-schedules command C:136
- file locking AD:64
 - suppressing in NFS mount command AD:93
- file --pin command C:137
- file size attribute C:290
- fileRead SOAP operation API:21
- fileReadAttach SOAP operation API:21
- fileReadString SOAP operation API:22
- files
 - .amm files S:11
 - .avm files S:11
 - .js JavaScript files S:11
 - .jsi JavaScript include files S:11, S:75
 - sample S:115
 - adding schedule exclusions for caching C:133, C:140
 - Avaki_JDBCStandAlone.jar API:66
 - Avaki_JDBCStandAlone_Minus3rd.jar API:66
 - avaki_studio.properties AD:129
 - AvakiAPI.disco WSDL discovery file API:3
 - AvakiAPIDocLit.wsdl API:3
 - AvakiAPIRpcEnc.wsdl API:3
 - AvakiAPIWithMIMEDocLit.wsdl API:3
 - AvakiAPIWithMIMERpcEnc.wsdl API:3
 - avakijdbc.properties API:67
- bindings.xml
 - copying during upgrade AD:344
 - on grid servers AD:50
 - on proxy servers AD:300
 - on share servers AD:59
- build.xml for data service plug-ins P:184, P:197
- cached, permissions on O:50
- caching of O:56
- changing ownership C:42
- changing permissions for AD:239
- clearing locks C:163
- configuring caching P:107
- configuring encryption level C:216
- connectinfo.txt AD:131
- copying AD:213, C:46
- copying locally AD:215
- data type mapping
 - about C:301
 - command to specify location C:68
 - format of C:301
 - deleted, purging from cache C:26
 - deleting AD:219, C:195
 - deleting pin schedules for C:135
 - DGAS properties AD:76
 - dgas_log.xml, DGAS log properties file AD:317
 - displaying C:35
 - evicting from cache P:135
 - for data service plug-ins P:180
 - icon for O:29
 - in the data grid O:8
 - install.exe AD:22, AD:24
 - invalidating all in cache C:28
 - invalidating one in cache C:27, P:136
 - JAR files for data service plug-ins P:180
 - jboss-service.xml, request log properties file AD:328
 - jdbc-log4j.properties API:66
 - join.properties file on proxy servers AD:300
 - krb5.conf AD:152
 - linking AD:217, C:161
 - listing C:166
 - listing locks C:164
 - listing schedules C:136
 - listing those pinned for caching P:134
 - log4j.xml,
 - Avaki Studio log properties file AD:317
 - server log properties file AD:317
 - manifest files for data service plug-ins P:180, P:197
 - marking for no caching C:136, P:132
 - marking for scheduled caching P:120
 - moving AD:210, C:176
 - permissions on new files O:49
 - pinning for scheduled caching C:137
 - properties files for DGAS C:293
 - purging all from cache and unpinning C:25
 - purging from cache and unpinning C:24
 - readme AD:12, AD:15, AD:339
 - refreshing in cache C:32
 - renaming AD:212
 - rendering results into O:25
 - scheduling for caching C:130
 - searching for AD:233
 - shareserver.ports
 - on grid servers AD:50
 - on share servers AD:60
 - system.properties AD:129
 - tagging for on-demand caching C:134, P:129
 - temporary, for sorting large result sets S:76
 - unscheduling cached files P:135
 - uploading to the data catalog AD:282
 - Workbench.project S:11
- fileWrite SOAP operation API:23
- Firefox
 - version requirements AD:5
 - setting for selecting run-as users P:27, P:54, P:92, P:224, P:227
- fonts in Studio display, setting S:23

functions
 in expressions S:73
 used with Aggregate operator S:110

G

garbage collection for views C:273
 GDCs, See grid domain controllers
 generated views
 about O:25, P:217, P:240
 defined AD:354, API:88, C:312, O:66, P:294, S:181
 running P:240
 scheduling updates P:231
 transactional consistency of P:217
 GeneratedViews directory O:33
 generating schemas S:19
 Generator operator S:117
 getAttributes SOAP operation API:23
 getCatalogs method API:75
 getCatalogTerm method API:75
 getDBOpOutput SOAP operation API:54
 getDBOpOutputAttach SOAP operation API:55
 getDBOpParameters SOAP operation API:56
 getDBOpSchema SOAP operation API:56
 getDBOpSchemaAttach SOAP operation API:57
 getDBOpSchemaString SOAP operation API:58
 getDSOutput SOAP operation API:37
 getDSOutputAttach SOAP operation API:38
 getDSOutputString SOAP operation API:38
 getDSParameters SOAP operation API:39
 getOutputString SOAP operation API:55
 getSchemas method API:75
 getSchemaTerm method API:75
 getSQL SOAP operation API:58
 getSystemAttributes SOAP operation API:24
 getUserAttributes SOAP operation API:24
 GIDs, configuring AD:68, C:217, C:221
 Global Parameters menu S:71
 grid directories, See directories, Avaki
 grid domain controllers
 backing up and restoring AD:113
 creating C:127
 defined AD:355, API:89, C:313, O:67, P:295, S:181
 DNS aliases for AD:10
 loading AD:14
 ports used by AD:6, AD:8
 primary AD:355, API:89, C:313, O:67, P:295, S:181
 secondary AD:43
 starting AD:14
 stopping AD:38, C:11
 stopping and restarting GDCs registered as services AD:39
 grid domains
 See domains, Avaki
 grid servers
 associating with caches P:111
 backing up and restoring AD:113
 choosing for web services API:5
 configuring AD:48, AD:52

grid servers (continued)
 configuring associated cache service C:33
 configuring cache service threads AD:131
 configuring nondefault ports AD:50
 connecting C:225
 connection info, setting S:23
 defined AD:355, API:89, C:313, O:67, P:295, S:181
 destroying C:226
 disabling auto-restart on C:12
 disassociating from caches P:112
 displaying associated cache C:27
 displaying status of operations on C:248
 enabling auto-restart C:9
 finding connect ports AD:53
 finding server names AD:53
 installing JDBC drivers on AD:49
 monitoring AD:99
 obtaining upgrade information C:250
 ports used by AD:6, AD:8
 registering C:9
 request logs for AD:327
 server logs AD:317
 settable attributes of C:290
 setting location of temp directory for AD:135
 setting plug-in cache size properties AD:137
 setting up command clients on AD:94
 starting AD:50, C:9, C:10
 stopping AD:51, C:11, C:12, C:227
 stopping and restarting AD:52
 time required to upgrade AD:340
 uncoupling associated cache C:34
 unregistering C:12
 upgrading C:249
 grid user accounts, See users
 grid-server --register command C:9
 grid-server --start command C:10
 grid-server --stop command C:11
 grid-server --unregister command C:12
 group --add --user command C:138
 Group By operator S:76, S:121
 group --create command C:141
 group --delete command C:143
 group --delete --user command C:144
 group --info command C:145
 group --list-user command C:147
 group mappings, adding for a particular DGAS C:92
 groups
 about O:43
 activating privileges for newly added users AD:192, AD:243
 adding to ACLs AD:243, S:97
 adding users to AD:191, C:138
 Administrators AD:45, O:43
 configuring default mappings C:218
 creating AD:191, C:141
 DatabaseAdministrators O:44
 DataProviders O:44
 default grid groups O:43
 default groups for DGAS AD:333

groups (continued)

- defined S:176
- deleting AD:198, C:143
- deleting users from C:144
- displaying information about AD:195, C:145
- enabling interconnection access AD:304
- everyone group O:44
- imported groups O:43
 - from LDAP AD:159
 - from NIS AD:165
 - refreshing AD:195
- in Avaki, defined AD:350, API:84, C:308, O:62, P:290
- listing C:144
- listing users in C:147
- making account changes take effect immediately for DGAS access AD:117
- MessagingUsers O:44
- modifying AD:195
- removing from ACLs AD:242
- removing users from AD:193
- setting up for DGAS AD:67
- UserAdministrators AD:45, O:44
- using in ACLs for cached objects O:50

H

hard links

- about O:36
- broken, to generated views P:240
- creating AD:217, C:161
- defined AD:355, API:89, C:313, O:67, P:295, S:182

hardware requirements for Avaki AD:2

help command C:148

help, online, for command line AD:xiii, C:xiii, P:ix

hideCatalogs JDBC property API:70

home directories AD:169

host names

- aliasing for GDCs AD:10
- setting for servers AD:32

HTTP and HTTPS ports

- default AD:8, AD:9, AD:10
- properties for AD:140

HTTP and web services API:7, API:8

HTTP POST problem in web browsers AD:5

HTTP request logs, See request logs

HTTPS and web services API:7, API:8

hyphens in command syntax AD:xvi, C:xvi, P:xii

I

IATEMPDIR environment variable AD:16

IBM AIX O:16

IBM AIX requirements AD:3

IBM DB2, versions and JDBC drivers for use with Avaki AD:6

icons for grid objects in the data catalog O:29

id command C:149

identity mapping P:6, S:36

imported groups, See groups

imported user accounts AD:167

increment field for Iterator operators S:151

indent size property for XML files AD:142

initialize field for Iterator operators S:151

initParameter element P:263

inner join S:156

input parameters

- creating P:88
- for data services
 - configuring P:84
 - deleting P:86

Input Source element S:125

input sources

- accessing columns from S:68
- browsing for in data catalog view S:18
- configuring for view models in Avaki Studio S:43
- creating S:55
- error handling S:143
- finding S:18

input streams, for data services, configuring P:87

inputParameter element P:264

inputSource element P:265

InputSource interface for data services P:186

inputStream element P:266

installation directory AD:350, API:84, C:308, O:62, P:290, S:176

installing Avaki

- about AD:13
- in Unix AD:16
- in Windows AD:23
- preparation and planning AD:1
- system requirements AD:2

integration, See data integration

interconnection IDs

- creating AD:295, C:149
- using in permissions AD:304, C:43
- using to provide cross-domain data access AD:295

interconnections between grid domains

- about AD:289
- breaking C:127
- creating AD:291
- defined AD:355, API:89, C:313, O:67, P:295, S:182
- disconnecting domains AD:314
- enabling access AD:299
- prerequisites for AD:94
- setting up C:128

two-way, exposing users AD:308

user access methods AD:294

viewing interconnected domains AD:310

Interconnects directory O:33

Internet Explorer

Avaki version requirements AD:5

setting for selecting run-as users P:27, P:54, P:92, P:224, P:227

Intersection operator S:148

performance S:76

IP address, setting for a server AD:32

isDSAvakiXML SOAP operation API:40

isList element P:266
 Iterator operator S:149
 example S:152

J

JAR files

for Avaki JDBC driver API:66
 configuring path for second JAR on one grid server P:247
 for data service plug-ins P:180
 for jConnect API:67

jarurl element P:267

Java

data service plug-in code samples P:190
 sample web services client API:9
 writing data service plug-ins in P:183

Java transformers

error logging P:246
 implementing P:243
 installing P:245
 referring to other documents P:245
 using P:245

java.io.tmpdir system property AD:135

java.protocol.handler.pkgs system property AD:141

java.rmi.server.hostname system property AD:32

java.security.krb5.conf system property AD:143, AD:153

java.security.krb5.kdc system property AD:153

java.security.krb5.realm system property AD:153

java.sql.DatabaseMetaData interface API:75

Javadoc, Avaki, accessing P:185

JavaScript

files S:11
 include files S:11, S:75
 sample S:115
 methods on data service plug-in objects P:201
 resources for learning about S:67
 using Java classes and interfaces in data service plug-ins P:200
 writing data service plug-ins in P:200

JavaScript expressions

about S:4, S:66
 menu for constructing S:71
 uses of in Avaki Studio S:vii

jConnect, See JDBC drivers

JDBC

accessing data catalog through O:24
 and caching of database operation results API:62
 connection properties API:68
 data types API:76
 directing queries to a particular grid server API:70
 result set types API:75

JDBC drivers

about API:61
 Avaki
 choosing version of API:66
 connection properties API:68
 JAR files for API:66

JDBC drivers (continued)

Avaki

setting classpath for API:67
 when to use API:65

batch mode API:77

choosing API:65

configuring for a database connector P:5

configuring two versions on one grid server P:247

connection strings API:71

defined AD:356, API:90, C:314, O:68, P:295, S:182

for supported DBMSes AD:3

installing AD:49

jConnect

changing default port API:71

setting classpath for API:67

using with Sybase databases AD:6

when to use API:65

loading API:68

prerequisites for using API:64

sample code API:77

setting system properties for API:67

supported features API:74

JDBC fetch size

setting for database connectors P:5, S:36

setting for database operations P:254

JDBC schema names API:69

jdbc-log4j.properties file API:66

join algorithms S:155

Join operator S:154

in tutorial S:44

performance S:76

join types S:156

join.properties file on proxy servers AD:300

JRE versions supported by Avaki AD:5

K

keepalive properties for HTTP ports AD:140

Kerberos

configuring with LDAP authentication services AD:152

system properties for AD:143

krb5.conf Kerberos configuration file AD:152

L

last access time file attribute C:290

LBF AD:280

LDAP

authentication services AD:148

See also authentication services, LDAP

authentication through DGAS AD:86

configuring for Kerberos access AD:152

disabling import on login AD:157

host port, default and nondefault AD:149

importing users on login AD:157

specifying a nondefault host port C:158

ldap --add-schedule command C:149

ldap --delete command C:153
 ldap --delete-schedule command C:154
 ldap --import command C:155
 ldap --info command C:157
 ldap --integrate command C:157
 ldap --list-schedules command C:158
 ldap --searchbase command C:160
 ldap --update command C:160
 links
 command for creating C:161
 uses of in data catalog O:36
 See also hard links, soft links
 Linux requirements AD:3
 listDBConns SOAP operation API:59
 listDBOps SOAP operation API:59
 listDBOpsByDBConn SOAP operation API:60
 listDomains SOAP operation API:25
 listDSs SOAP operation API:40
 listSearches SOAP operation API:25
 ln command C:161
 load balancing factor for share servers AD:280, C:244
 local caches AD:356, API:90, C:314, O:14, O:68, P:296, S:182
 locks command C:163
 locks on files
 clearing C:163
 obtaining a list C:164
 log properties file, sample AD:323
 log4j AD:318, P:196
 logging
 audit logs AD:108
 categories of loggable events AD:318
 configuring audit logging AD:319
 for data service plug-ins P:196
 for JDBC API:66
 for TrAX transformers P:246
 for type mapping C:304
 HTTP request logs AD:108
 log4j properties files for servers and Studio AD:317
 properties files for request logs AD:327
 viewing the server log AD:107
 logging in AD:199, C:164
 logical operators S:72
 logical operators in searches AD:234
 logicBox element P:268
 login command C:164
 login info, setting S:23
 logout command C:165
 ls command C:166
 ls SOAP operation API:26
 lsSize SOAP operation API:26

M

manifest files for data service plug-ins P:180, P:197
 mappings
 between Avaki and local users/groups for DGAS AD:68
 between Avaki users and database users, See database identity mappings

mappings (continued)
 database identity
 adding C:252
 deleting C:253
 listing C:255
 default, setting up AD:71
 DGAS default AD:69, C:220
 DGAS domain-wide
 for groups, setting up AD:70
 for users, setting up AD:70
 users, groups, and defaults AD:68
 DGAS dynamic C:109
 for data types in SQL views C:68, C:301
 for users and groups for DGAS C:94, C:109
 per-DGAS
 adding for groups C:92
 adding for users C:94
 deleting C:100
 deleting for groups C:99
 per-DGAS, users and groups AD:88
 self mappings for users C:109
 See also data type mappings
 maximum concurrent data services setting for execution services AD:109
 memory requirements for Avaki software AD:3
 message tests in monitor services AD:101
 message timeout properties for Avaki servers and clients AD:133
 MessagingUsers group and user accounts O:44
 metadata O:13, S:3
 Metadata directory O:33
 metadata models
 about S:77
 creating S:84
 defined AD:356, API:90, C:314, O:68, P:296, S:182
 deleting S:94
 deploying C:285, S:91
 deriving S:92
 editing S:84
 files associated with S:11
 importing S:79
 mapping to Avaki objects S:88
 naming scheme S:91
 undeploying C:286, S:94
 Microsoft SQL Server, versions and JDBC drivers for use with Avaki AD:6
 Microsoft Visual Studio API:5
 MicroSoft Windows O:16
 MIME in Avaki web services API:3
 minus signs in command syntax AD:xvi, C:xvi, P:xii
 mkdir command C:167
 mkdir SOAP operation API:27
 mkdirParents SOAP operation API:27
 mkdirServer SOAP operation API:29
 mkdorParentsServer SOAP operation API:28
 models, See metadata models and view models
 modification time file attribute C:290

monitor --add command C:167
 monitor --clear command C:168
 monitor --create command C:170
 monitor --delete command C:171
 monitor --list command C:172
 monitor --result command C:172
 monitor services
 monitor --start command C:173
 monitor --stop command C:174
 monitoring
 about AD:99
 adding tests C:167
 configuring AD:101
 creating monitor services C:170
 deleting monitor services AD:106, C:171
 deleting tests AD:105
 disabling and enabling tests AD:104
 listing active tests C:172
 logging AD:107
 message tests AD:101
 ping tests AD:100
 removing tests C:168
 restarting tests AD:105, C:173
 stopping tests AD:105, C:174
 viewing results AD:103, C:172
 mount port for DGAS AD:81
 mount protocol port, default AD:9
 Mozilla, Avaki version requirements AD:5
 Multiplexer operator S:157
 multiplexing socket properties AD:142
 mv command C:176
 mv SOAP operation API:29
 MySQL
 configuring XA driver for P:7, S:36
 versions and JDBC drivers for use with Avaki AD:6

N

name element P:269
 name property for columns S:61
 names of Avaki objects
 about O:29
 avoiding underscores in when using BusinessObjects API:81
 case sensitivity and restrictions AD:206
 changing AD:212
 metadata models and mapped tables S:91
 of elements S:58
 qualified names O:30
 restrictions on AD:41
 three-part O:24
 navigator in Studio S:17
 nesting operations in data services S:149
 .NET, See .NET under Symbols at the beginning of the index
 Netscape requirements AD:5
 NFS
 and permissions AD:12
 configuring NFS port for DGAS AD:81

NFS (*continued*)
 port, default AD:9
 shutting down before starting a DGAS AD:65
 NFS clients
 attribute caching for AD:336
 defined AD:356, API:90, C:314, O:68, P:296, S:183
 older, accessing data grid through AD:337
 requirements for O:16
 setting up AD:91
 NFS URLs AD:93
 NIS
 disabling import on login AD:163
 importing users on login AD:163
 See also authentication services, NIS
 nis --add-schedule command C:177
 nis --delete command C:181
 nis --delete-schedule command C:181
 nis --import command C:182
 nis --info command C:183
 nis --integrate command C:184
 nis --list-schedules command C:184
 nis --update command C:185
 NLM AD:64
 NLM protocol port, default AD:9
 notifications, See update notifications

O

object host name grid server attribute C:290
 octothorpe AD:33, AD:129
 ODBC
 accessing data catalog through O:24
 defined AD:356, API:90, C:314, O:68, P:296, S:183
 support for API:80
 ODBC drivers, using with Avaki API:80
 offline expiration cache attribute C:291
 on-demand caching
 about P:119
 defined AD:357, API:91, C:315, O:69, P:296, S:183
 of database operation and data service results O:57, P:108
 of files O:56, P:107
 online help for command line AD:xiii, C:xiii, P:ix
 operating systems supported by Avaki O:16
 operations, monitoring AD:99
 operators
 about S:5
 adding to a view model S:55
 Aggregate S:108
 connecting S:57
 Custom S:111
 descriptions of S:59
 Generator S:117
 Group By S:121
 in searches AD:234
 Input Source S:125
 Intersection S:148
 Iterator S:149
 Join S:154

- operators (*continued*)
 - logical, in expressions S:72
 - moving S:56
 - Multiplexer S:157
 - names of S:58
 - Order By S:159
 - performance considerations S:5
 - Projection S:161
 - properties dialogs S:58
 - Result S:164
 - Select S:165
 - selecting S:56
 - sort-based, performance of S:76
 - Splitter S:166
 - Union S:168
 - Update S:169
 - with red borders S:60
 - Oracle 10g
 - configuring XA driver for P:7, S:36
 - versions and JDBC drivers for use with Avaki AD:6
 - Order By operator S:159
 - performance S:76
 - os arch grid server attribute C:290
 - os name grid server attribute C:290
 - outer-full join S:156
 - outer-left join S:156
 - outer-right join S:156
 - output streams, for data services, configuring P:86
 - outputStream element P:269
 - ownership of objects in the data catalog
 - about O:46
 - setting AD:242
- P**
- palette in Avaki Studio view model editor S:20, S:54
 - parameters
 - accessing in expressions S:71
 - adding S:64
 - deleting S:65
 - displaying S:63
 - for data service plug-ins
 - about P:178
 - specifying for Java plug-ins C:188
 - specifying for XSLT plug-ins P:181
 - for testing view models S:50
 - in Avaki Studio, about S:4
 - input, for data services, configuring P:84
 - mapping input parameters S:144
 - modifying S:65
 - reordering S:65
 - validating S:64
 - ParameterSpec interface for data services P:187
 - passwords
 - changing AD:175, C:185
 - specifying for JDBC connections API:69
 - patches, settable attributes of C:290
 - performance S:5, S:76
 - benefits of caching O:54
 - tracking, enabling in an Avaki shell C:245
 - Perl, sample web services client API:9
 - permissions
 - about O:45
 - changing AD:239, C:41
 - granted by grid groups O:43
 - hiding objects with O:49
 - on new Avaki shares AD:261
 - on shared data AD:12
 - setting in ACLs S:97
 - values for AD:242, O:47
 - viewing for Avaki services C:186
 - See also ACLs
 - permissions command C:186
 - permissions SOAP operation API:30
 - perspectives in Avaki Studio
 - Avaki S:13
 - defined S:12
 - Resource S:11
 - pin for caching, defined AD:357, API:91, C:315, O:69, P:297, S:183
 - ping tests in monitor services AD:100
 - planning an Avaki deployment AD:1
 - platforms supported by Avaki O:16
 - plugin command P:184
 - plugin --generate command C:187
 - Plugin interface for data services P:186
 - Plug-in Wizard and creating data service plug-ins P:183
 - plug-ins, See data service plug-ins
 - plus signs in command syntax AD:xvi, C:xv, P:xi
 - ports
 - bad port cache AD:141
 - changing, See ports, nondefault
 - CIFS, releasing before running a DGAS AD:66
 - default, for Avaki servers AD:6
 - HTTP and HTTPS, See HTTP and HTTPS ports
 - LDAP host
 - default and nondefault AD:149
 - specifying C:158
 - NFS, default AD:9
 - nondefault
 - configuring for DGAS AD:74, AD:75
 - configuring for grid servers AD:50
 - configuring for proxy servers AD:300
 - configuring for share servers AD:59
 - RMI, See RMI ports
 - SMB, default AD:9
 - specifying in WSDL API:6
 - SSL, See SSL ports
 - TDS AD:8, AD:50, AD:145, API:71
 - ports, connect, See connect ports
 - POST problem in web browsers AD:5
 - precision property for columns S:61
 - preferences for Avaki Studio, setting S:23
 - primary GDCs AD:357, C:315, O:69

privacy and HTTPS API:8
 Projection operator S:161
 in tutorial S:46
 projects
 creating S:13
 defined S:12
 properties
 cache sizes for data service plug-ins AD:137
 DGAS
 configuring AD:82
 controlling cache size AD:124
 displaying AD:124, C:103
 displaying descriptions C:104
 listing C:103
 properties file C:293
 setting C:113
 unsetting C:117
 for cache services AD:135
 for encryption of grid objects AD:139
 for HTTP and HTTPS ports AD:140
 for HTTP keepalives on proxy servers AD:140
 for JDBC clients AD:128
 for JDBC connections API:68
 for Kerberos AD:143
 for multiplexing sockets AD:142
 for remote object stub cache AD:144
 for schedule exclusion cache AD:144
 for server request logs AD:327
 for servers acting as clients AD:128
 for share servers AD:138
 for virtual database table cache size AD:144
 for XA connections C:70, P:7, S:37
 Java system properties, providing to JDBC driver API:67
 JDBC, specifying for a database connector C:67
 Kerberos default realm AD:153
 Kerberos key distribution center AD:153
 location of Kerberos configuration file AD:153
 message timeouts for Avaki servers and clients AD:133
 remote coherence window for configurations AD:141
 setting server's host name or IP address AD:33
 system. See system properties.
 views
 displaying C:274
 setting C:275
 XML indent size AD:142
 properties dialog boxes S:58
 provisioning data O:21
 proxy --add command C:191
 proxy --delete command C:191
 proxy --list command C:192
 proxy routing tables
 about AD:289
 configuring AD:292, C:191
 displaying C:192

proxy servers
 about AD:289
 configuring AD:299
 configuring nondefault ports AD:300
 connecting C:228
 defined AD:357, API:91, C:315, O:69, P:297, S:183
 deleting from the routing table C:191
 destroying C:228
 disabling auto-restart C:14
 enabling auto-restart AD:300, C:12
 finding connect port numbers AD:303
 finding server names AD:303
 installing in Windows AD:24
 installing in Unix AD:18
 ports used by AD:10
 registering for auto-restart AD:302, C:12
 request logs for AD:327
 server logs for AD:317
 setting HTTP keepalive properties for AD:140
 setting up C:191
 starting AD:301, C:12, C:13
 stopping C:14, C:228
 stopping and restarting AD:301
 time required to upgrade AD:341
 unregistering C:14
 when to deploy AD:290
 proxy-server --register command AD:302, C:12
 proxy-server --start command C:13
 proxy-server --stop command C:14
 proxy-server --unregister command C:14
 pwd command C:193

Q

qualified names
 about O:30
 for data services, specifying API:35, C:51, C:267
 for database connectors, specifying C:254
 for database operations, specifying C:25, C:264
 for groups, specifying C:43, C:142, C:145, C:147, C:148
 for users, specifying C:42, C:53, C:60, C:86, C:89, C:141,
 C:145, C:165, C:251, C:266, C:280
 queries, See ad-hoc queries
 query engine
 defined AD:357, API:91, C:315, O:69, P:297, S:184
 mapping data types for C:301
 sort chunk size property AD:139
 queryCacheTTL JDBC property API:71

R

RAM requirements for Avaki software AD:3
 range input for Iterator operators S:151
 readme file AD:12, AD:15, AD:339
 red borders on operators in Studio S:60
 Red Hat Linux requirements AD:3
 ref element P:270

refresh schedules
 for Avaki shares AD:266, C:231
 advanced AD:272
 calendared AD:271
 exclusions for AD:274
 listing C:239
 one-time AD:269
 periodic AD:270
 removing AD:278, C:236
 for data services P:152
 advanced P:157
 calendared P:155
 one-time P:155
 periodic P:157
 for database operations P:139
 advanced P:144
 calendared P:142
 one-time P:142
 periodic P:144
 for files or directories P:120
 advanced P:127
 calendared P:125
 one-time P:125
 periodic P:127
 for generated views P:231
 advanced P:236
 calendared P:234
 one-time P:233
 periodic P:236
 for LDAP authentication services AD:185
 for virtual database operations P:139
 advanced P:144
 calendared P:142
 one-time P:142
 periodic P:144
 refreshing users on login AD:149
 reindex interval for search services AD:232
 remote caches AD:357, API:91, C:315, O:14, O:69, P:297,
 S:184
 removeAttribute SOAP operation API:31
 replica --add command C:193
 replica --config command C:193
 replica --delete command C:194
 replica --info command C:194
 replica --synch command C:195
 request logs
 configuring AD:327
 viewing AD:108
 requirements, pre-installation AD:2
 Result element S:164
 result sets
 accessor functions S:73
 combining S:44
 large, providing space for sorting S:76
 types in JDBC API:75

rm command C:195
 rm SOAP operation API:31
 RMI ports
 default AD:8, AD:9, AD:10
 linked to grid server connect ports AD:50
 linked to share server connect ports AD:60
 routing tables, configuring AD:292, C:191
 row-level access control S:74
 RowSetFactory class for data services P:188
 rowsets O:11, P:273
 rpc/encoded web services API:3, API:5
 rpcinfo command AD:66
 Rudi port AD:352, C:310, O:64
 run-as users, See users and user accounts

S

Saxon C:265, C:268, C:271, P:220, P:224, P:227
 scale property for columns S:61
 schedule --delete command C:196
 schedule exclusions
 about AD:274, P:166
 adding for Avaki share rehashes C:234
 adding for data services C:51
 adding for directories C:125
 adding for files C:133, C:140
 adding for LDAP authentication services C:152
 adding for NIS authentication services C:180
 adding for views C:262
 adding to Avaki directories C:121
 adding to database operations C:76
 applying to schedule entries P:168
 caching properties for AD:144
 configuring AD:274, P:166
 creating custom C:198
 creating daily C:199
 creating monthly C:201
 creating weekly C:203
 creating yearly C:205
 defined AD:358, API:92, C:316, O:70, P:298, S:184
 deleting C:207
 displaying information about C:208, P:171
 listing names C:209
 schedule --info command C:197
 schedule --list command C:197
 schedule --print-iterations command C:198
 scheduled caching
 about P:119
 defined AD:357, API:91, C:315, O:69, P:297, S:184
 of database operation and data service results O:59, P:109
 of files O:57, P:107
 scheduleexclusion --create --custom command C:198
 scheduleexclusion --create --daily command C:199
 scheduleexclusion --create --monthly command C:201
 scheduleexclusion --create --weekly command C:203
 scheduleexclusion --create --yearly command C:205
 scheduleexclusion --delete command C:207
 scheduleexclusion --info command C:208

- scheduleexclusion --list command C:209
- schedules
 - adding for data services C:48, P:152
 - adding for database operations C:73, P:139
 - adding for directories C:117
 - adding for files C:130
 - adding for LDAP user importation C:149
 - adding for NIS user importation C:177
 - adding for views C:259
 - adding for virtual database operations P:139
 - creating cron specifications C:297
 - creating custom exclusions C:198
 - creating daily exclusions C:199
 - creating monthly exclusions C:201
 - creating weekly exclusions C:203
 - creating yearly exclusions C:205
 - cron expressions in AD:185, AD:273
 - deleting C:196
 - deleting exclusions C:207
 - deleting for data services C:54
 - deleting for database operations C:78
 - deleting for directories C:122
 - deleting for files C:135
 - deleting for LDAP user importation C:154
 - deleting for NIS user importation C:181
 - deleting for views C:272
 - displaying exclusion information C:208
 - displaying information about C:197
 - exclusions, see also schedule exclusions
 - execute permissions required O:59
 - for refreshing LDAP authentication services AD:185
 - for view generators and generated views P:231
 - listing C:197
 - listing execution times C:198
 - listing for data services C:58
 - listing for database operations C:88
 - listing for directories C:126
 - listing for files C:136
 - listing for LDAP authentication services C:158
 - listing for NIS authentication services C:184
 - listing for views C:274
 - listing names of exclusions C:209
 - refresh, See refresh schedules
 - types of AD:185, AD:267
- schemas
 - about S:3, S:60
 - enabling browsing on a database connector P:4
 - expressed in metadata models S:77
 - for Avaki data services API:75
 - for operators, column properties of S:60
 - for virtual database operations and their SQL views API:75
 - for virtual database operations, generating P:57
 - generating S:19
 - generating for data services C:55, P:98
 - generating for database operations C:80, P:31
 - getting information about through JDBC API:75
 - getting via JDBC API:75
- schemas (*continued*)
 - JDBC schema names API:69
 - modifying S:161
 - types in Avaki domains P:51, P:64
 - viewing for databases P:9
 - See also metadata models
 - search (execute) command C:211
 - search --create command C:209
 - search --delete command C:210
 - search --get-rehash-level command C:212
 - search --info command C:214
 - search --rehash command C:215
 - search --set-rehash-level command C:215
 - search SOAP operation API:32
 - searches O:13
 - configuring rehash intervals C:215
 - creating search services AD:231, C:209
 - deleting AD:236
 - deleting search services C:210
 - displaying search service information C:214
 - performing AD:233, C:211
 - rehashing search services C:215
 - reindex interval, setting AD:232
 - viewing rehash intervals C:212
 - SearchQuery complex type API:17
 - SearchResult complex type API:17
 - secondary GDCs AD:43
 - adding C:193
 - deleting C:194
 - forcing updates C:195
 - setting refresh intervals C:193
 - setting update interval C:195
 - viewing C:194
 - See also grid domain controllers
 - security
 - .NET API:9
 - about O:10
 - ACLs O:45
 - authentication O:41
 - configuring encryption levels C:216
 - displaying encryption levels C:222
 - for web clients API:8
 - HTTPS API:8
 - setting permissions C:41
 - SSL certificates API:8
 - user accounts and groups O:43
 - viewing permissions C:186
 - security --config command C:216
 - security --default-gid command C:217
 - security --default-group command C:218
 - security --default-uid command C:219
 - security --default-user command C:220
 - security --gid command C:221
 - security --info command C:222
 - security --uid command C:222
 - Select operator S:165
 - server connect ports, See connect ports
 - server --dgas --connect command C:223

- server --dgas --destroy command C:224
- server --dgas --stop command C:225
- server --grid --connect command C:225
- server --grid --destroy command C:226
- server --grid --stop command C:227
- server logs
 - configuring AD:317
 - viewing AD:107
- server --proxy command C:228
- server --share --connect command C:228
- server --share --disconnect command C:229
- server --share --stop command C:230
- servers, backing up databases for C:23
- servers, Avaki
 - defined AD:350, API:84, C:308, O:62, P:290, S:176
 - displaying software version of AD:99
 - finding names of AD:131
 - in a typical deployment O:18
 - monitoring AD:99
 - ports used by AD:6
 - problems communicating with AD:32
 - setting cache service properties for AD:135
 - setting host name or IP address to advertise AD:32
 - setting message timeout properties for AD:133
 - See also DGAS, grid domain controllers, grid servers, share servers, and proxy servers
- servers, proxy, See proxy servers
- services, Avaki
 - copying C:46
 - defined AD:358, API:92, C:316, O:70, P:298, S:184
 - icon for O:29
- setAttribute SOAP operation API:32
- setOutputStream JavaScript method for data service plug-ins P:202
- share --add-rehash-schedule command C:231
- share --add-share-servers command C:232
- share --create command C:235
- share --delete-rehash-schedule command C:236
- share --disconnect command C:238
- share --get-local-path command C:238
- share --get-status command C:239
- share --list-rehash-schedules command C:239
- share --list-share-servers command C:240
- share --rehash command C:240
- share --remove-share-servers command C:241
- share servers
 - about AD:54
 - adding to Avaki shares AD:263
 - backing up and restoring AD:115
 - before disconnecting AD:54, C:230
 - configuring a machine with one share server AD:55
 - configuring multiple share servers on one machine AD:59
 - configuring to use nondefault ports AD:59
 - connecting to grid servers C:228
 - defined AD:358, API:92, C:316, O:70, P:298, S:184
 - disabling auto-restart C:17
 - disconnecting from grid servers C:229
 - enabling auto-restart C:15
 - finding connect ports AD:59
 - finding server names AD:59
 - installing in Windows AD:24
 - installing on Unix AD:18
 - local path AD:261
 - modifying load balance factor C:244
 - multiple AD:55, AD:263
 - obtaining upgrade information C:250
 - ports used by AD:9
 - registering C:15
 - registering for auto-restart AD:57
 - removing from Avaki shares C:241
 - removing from shares AD:265
 - replacing for Avaki shares C:242
 - server logs for AD:317
 - setting for Avaki shares AD:260
 - setting load balancing factor AD:280
 - setting system properties for AD:138
 - starting AD:56, C:15, C:16
 - stopping C:17, C:230
 - stopping and restarting AD:57
 - time required to upgrade AD:341
 - unregistering C:17
 - upgrading C:249
 - write access and user accounts AD:12
- share --set-local-path command C:241
- share --set-share-servers command C:242
- share --set-status command C:243
- share --update-share-servers command C:244
- shared directories, See Avaki shares
- shares, See Avaki shares and CIFS shares
- Shares directory O:34
- shares, CIFS AD:125
- share-server --register command AD:57, C:15
- share-server --start command C:16
- share-server --stop command C:17
- share-server --unregister command C:17
- shareserver.ports file
 - on grid servers AD:50
 - on share servers AD:60
- shell command C:245
- shells, Avaki, accessing C:245
- shortcuts created in Windows installations AD:27
- SMB ports, default AD:9
- SOAP
 - formal definition API:1
 - learning about API:1
 - over HTTP API:8
 - over HTTPS API:8
 - standards compliance API:1
- SOAP clients, See web services clients
- SOAP complex types API:12
 - AdHocDBOPExecutionParams API:13
 - AvakiPrincipal API:13
 - DataCatalogAttribute API:14

SOAP complex types (*continued*)

DataCatalogEntry API:15
 DataCatalogPermission API:15
 DataServiceExecutionParams API:16
 DBOPExecutionParams API:16
 SearchQuery API:17
 SearchResult API:17

SOAP operations

accessibleDBOp API:42
 accessibleDS API:36
 accessiblePath API:19
 chmod API:19
 chown API:20
 components of web service API:4
 data catalog API:18
 data services API:34
 database operations API:40
 executeAdHocDBOp API:43
 executeAdHocDBOpWithoutOutput API:44
 executeAdHocDBOpWithoutOutputAttach API:46
 executeAdHocDBOpWithoutOutputString API:47
 executeDBOp API:48
 executeDBOpBytesInput API:49
 executeDBOpGridFileInput API:50
 executeDBOpWithOutput API:50
 executeDBOpWithOutputAttach API:52
 executeDBOpWithOutputString API:53
 executeDS API:36
 fileRead API:21
 fileReadAttach API:21
 fileReadString API:22
 fileWrite API:23
 getAttributes API:23
 getDBOpOutput API:54
 getDBOpOutputAttach API:55
 getDBOpParameters API:56
 getDBOpSchema API:56
 getDBOpSchemaAttach API:57
 getDBOpSchemaString API:58
 getDSOutput API:37
 getDSOutputAttach API:38
 getDSOutputString API:38
 getDSParameters API:39
 getOutputString API:55
 getSQL API:58
 getSystemAttributes API:24
 getUserAttributes API:24
 isDSAvakiXML API:40
 listDBConns API:59
 listDBOps API:59
 listDBOpsByDBConn API:60
 listDomains API:25
 listDSs API:40
 listSearches API:25
 ls API:26
 lsSize API:26
 mkdir API:27
 mkdirParents API:27

SOAP operations (*continued*)

mkdirParentsServer API:28
 mkdirServer API:29
 mv API:29
 permissions API:30
 removeAttribute API:31
 rm API:31
 search API:32
 setAttribute API:32
 tester API:33
 whoami API:33
 SOAP::Lite API:5
 sockets, multiplexing AD:142
 setting maximum channels AD:142
 setting maximum write AD:142
 setting send buffer size for AD:143
 soft links
 about O:36
 creating AD:217, C:161
 defined AD:358, API:92, C:316, O:70, P:298, S:185
 not used in searches AD:233
 software requirements for Avaki AD:2
 Solaris requirements AD:3
 sort chunk size
 controlling S:76
 for query engine AD:139
 spaces
 in Windows install pathnames, avoiding AD:26
 to separate arguments in Avaki commands AD:xvi, C:xvi,
 P:xii
 SPARC/Solaris requirements AD:3
 Splitter operator S:166
 SQL
 aggregate functions and aliasing columns S:42
 as prerequisite for Avaki Studio users S:vii
 statements
 in database operations C:86, P:251
 in virtual database operations C:90
 SQL Server, versions and JDBC drivers for use with
 Avaki AD:6
 sql view --delete command C:246
 sql view --get-description command C:246
 sql view --set-description command C:247
 SQL views
 about O:8, O:22, P:38
 adding to categories P:47
 configuring attributes P:44
 data type mappings for C:68, C:301, P:39
 defined AD:358, API:92, C:316, O:70, P:298, S:185
 deleting C:246
 displaying descriptions C:246
 displaying tables provisioned from database connectors C:73
 enabling provisioning on a database connector P:4
 from data service results, generating P:60, P:100
 generated from database operations, removing P:35
 generating from data services C:56
 generating from database operations C:82, P:34
 location in categories S:18

SQL views (*continued*)

- managing P:20
- modifying P:43
- modifying descriptions C:247
- names in data catalog O:24
- provisioning P:39
- provisioning from database connectors C:71
- qualified names for O:31
- removing P:44
- schema types for P:51, P:64
- schemas for S:22
- searching for AD:233
- table types for API:75
- viewing P:42
- viewing and modifying ACLs P:46
- square brackets in command syntax AD:xv, C:xv, P:xi
- SSL certificates
 - about API:8
 - generating AD:33
 - installing AD:39
 - planning for AD:8
- SSL ports, default AD:8, AD:9, AD:10
- status command C:248
- status of grid operations, monitoring AD:99
- stored procedures API:73
- StreamingRowSet interface for data services P:187
- streams, closing in data service plug-ins P:186
- Studio, See Avaki Studio
- style sheet engines
 - for data service view generators P:227
 - for database operation view generators P:224
 - for file view generators P:220
- style sheet engines for database view generators C:265, C:268
- style sheet engines for file view generators C:271
- Sun JDK for compiling data service plug-ins P:184
- Sun ONE Directory Server AD:148
 - See also authentication services, LDAP AD:148
- Sun Solaris requirements AD:3
- SuSE Linux requirements AD:3
- Sybase ASA, versions and JDBC drivers for use with Avaki AD:6
- Sybase ASE
 - configuring XA driver for P:7, S:36
 - connection property required for API:70
 - versions and JDBC drivers for use with Avaki AD:7
- Sybase contact information AD:xvi, API:vii, C:xvi, O:vi, P:xii, S:ix
- Sybase IQ, versions and JDBC drivers for use with Avaki AD:7
- syntax conventions for commands AD:xv, C:xiv, P:x
- system attributes AD:248, S:101
- System directory O:34
- system properties
 - about AD:128
 - descriptions of AD:131
 - setting on Avaki Studio AD:129
 - setting on clients AD:129
 - setting on servers AD:129
 - setting with JDBC driver API:67

- system requirements for Avaki software AD:2, O:16
- system.properties file AD:33, AD:129

T

- table schema view S:22
- table types for SQL views API:75
- tables
 - deleting SQL views C:246
 - displaying descriptions C:246
 - generating from data services C:56
 - generating from database operations C:82
 - in metadata models S:77
 - arranging in editor S:84
 - making accessible via JDBC S:91
 - mapping to Avaki objects S:88
 - naming scheme for S:91
 - in virtual database, displaying C:283
 - mapping data types for C:301
 - modifying descriptions C:247
 - provisioned from database connectors, displaying C:73
 - provisioning as SQL views O:22
 - provisioning from database connectors C:71
 - qualified names for O:32
 - schemas for, displaying S:22
 - See also SQL views
- tabs for view models S:21
- target element P:270
- TCP channel sockets AD:142
- TDS port
 - changing AD:50, AD:145, API:71
 - default AD:8
- technical support contact information AD:xvi, API:vii, C:xvi, O:vi, P:xii, S:ix
- temp directories for grid servers, setting AD:135
- Templates class P:244
- tester SOAP operation API:33
- testing an upgraded grid domain AD:346
- testing function for WS clients API:33
- text conventions C:xiv
- timeout properties for Avaki server communications AD:133
- timeouts, configuring for database operations P:253
- transactions, See distributed transactions P:78
- TrAX API P:243
- two-phase commit protocol P:79
- type element P:270
- type mapping, See data types, mapping
- type property for columns S:61
- TypeMapping log4j category C:304
- types for variables S:70

U

- UID attribute in LDAP authentication services AD:150
- UIDs, configuring AD:68, C:219, C:222
- underscore characters in Avaki names API:81
- Union operator S:168
- Unix file mode semantics, setting AD:337

- Unset permission in ACL AD:242, O:47
- update intervals for GDCs, setting C:195
- update notifications
 - configuring P:238
 - defined AD:359, API:93, C:317, O:71, P:299, S:185
 - enabling AD:311
- Update operator S:169
- upgrade command C:249
- upgrade --info command C:250
- upgrading Avaki software
 - interoperability of different versions AD:339
 - preparation steps AD:341
 - upgrade planning AD:340
 - upgrade steps AD:342
- upstream, defined S:3
- urlLogicBox element P:271
- user attributes AD:248, S:101
- user command C:250
- user --create command C:251
- user --db-mapping --add command C:252
- user --db-mapping --delete command C:253
- user --db-mapping --list command C:255
- user --delete command C:257
- user groups, See groups
- user --info command C:258
- user --list-group command C:258
- UserAdministrators group AD:45, O:44
- users and user accounts
 - about O:43
 - adding database identity mappings C:252
 - adding to ACLs AD:243, S:97
 - adding to groups AD:191, C:138
 - administrative accounts, setting up AD:44
 - changing passwords AD:175, C:185
 - clearing credentials from DGAS cache AD:117
 - configuring associated cache service C:33
 - configuring default mappings C:220
 - configuring dynamic mappings C:109
 - configuring self mappings C:109
 - creating accounts AD:168, C:251
 - creating home directories AD:169
 - default users for DGAS AD:333
 - deleting accounts AD:189, C:257
 - deleting database identity mappings C:253
 - deleting from groups C:144
 - disabling import on login (LDAP) AD:157
 - disabling import on login (NIS) AD:163
 - displaying and changing account information AD:187
 - displaying associated cache C:27
 - displaying full names and contact information C:258
 - displaying names C:286
 - enabling interconnection access AD:304
 - enabling on authentication services C:250
 - exposing in a two-way interconnect AD:308
 - giving access to other domains AD:294
 - grid accounts AD:167
 - imported accounts AD:167
 - users and user accounts (*continued*)
 - importing from LDAP services AD:157
 - importing from NIS services AD:164
 - importing on login (LDAP) AD:157
 - importing on login (NIS) AD:163
 - listing database identity mappings C:255
 - listing group membership for C:258
 - logging in AD:199
 - logging out when newly added to groups AD:192, AD:243
 - making account changes take effect immediately for DGAS access AD:117
 - managing AD:167
 - mapping Avaki users to database users, See database identity mappings
 - MessagingUser O:44
 - qualified names for O:32
 - querying whether enabled in LDAP AD:151
 - refreshing imported accounts AD:185
 - refreshing on login (LDAP) AD:149
 - removing from ACLs AD:242
 - removing from groups AD:193
 - roles for O:43
 - run-as users
 - browser setting for selecting P:27, P:54, P:92, P:224, P:227
 - for data service views P:227
 - for data services P:92
 - for database operation views P:223
 - for database operations P:26
 - for virtual database operations P:54
 - setting run-as user for views C:279
 - setting up for DGAS AD:67
 - setting up local accounts for Avaki AD:11
 - specifying for JDBC connections API:69
 - uncoupling associated cache C:34

V

- validation error expressions S:65
- validation expressions S:64
- value element P:272
- values element P:272
- variables
 - about S:69
 - allowed types for S:70
 - downstream, menu of S:71
 - in Avaki Studio, about S:4
 - updating S:69
- VB .NET, See .NET
- versions
 - of Avaki software, displaying AD:99, C:148, C:250
 - SOAP API:1
 - TrAX P:243
 - WSDL API:1
- vertical bars in command syntax AD:xv, C:xv, P:xi
- view --add-schedule command C:259
- view --create --database command C:263
- view --create --data-service command C:266

- view --create --file command C:267
- view --delete command C:272
- view --delete-schedule command C:272
- view --depends command C:272
- view --garbage-collect command C:273
- view generators
 - about O:8, O:25, P:217
 - caching of input files P:241
 - configuring update notifications for P:238
 - defined AD:359, API:93, C:317, O:71, P:299, S:185
 - for data services
 - setting up P:225
 - specifying a style sheet engine P:227
 - for database operations
 - setting up P:221
 - specifying a style sheet engine P:224
 - for files
 - setting up P:218
 - specifying a style sheet engine P:220
 - for large data sets and unsupported formats P:242
 - listing dependent operations P:228
 - modifying P:229
 - non-XSLT-based P:242
 - removing P:239
 - rowsets as inputs of P:275
 - running P:240
 - scheduling updates P:231
 - troubleshooting P:240
 - using TrAX transformers P:242
- view --info command C:274
- view --list-schedules command C:274
- view models
 - about O:23, S:2
 - configuring input sources S:43
 - creating S:42
 - defined AD:359, API:93, C:317, O:71, P:299, S:185
 - deploying as data services S:50
 - error handling S:143
 - files associated with S:11
 - opening, saving and closing S:17
 - red borders showing errors S:60
 - sample workflow for S:29
 - schemas S:3
 - tabs for, in a project S:21
 - testing S:49
 - view model editor S:20
- view --regenerate command C:273
- view --set-property command C:275
- view --update command C:279
- ViewLibrary category, contents of S:18
- views
 - adding generation schedules C:259
 - adding schedule exclusions C:262
 - configuring values for SQL parameters C:275
 - configuring with database sources C:263
 - configuring with data-service sources C:266
 - configuring with file sources C:267
- views (*continued*)
 - deleting C:272
 - deleting generation schedules C:272
 - listing generation schedules for C:274
 - obtaining information about C:274
 - regenerating C:273
 - removing old results C:273
 - setting run-as user C:279
 - showing dependencies C:272
 - specifying a style sheet engine C:265, C:268, C:271
 - updating C:279
 - See also generated views, SQL views, view generators, view models
- virtual database
 - about O:22
 - configuring attributes P:70
 - defined AD:359, API:93, C:317, O:71, P:299, S:185
 - displaying SQL views C:283
 - executing ad-hoc queries on C:282
 - schema
 - browsing catalogs P:64
 - browsing schemas P:64
 - browsing tables P:64
 - table cache size system property AD:144
 - types of schemas P:51, P:64
 - viewing and modifying ACLs P:72
- virtual database operations
 - about O:23, P:49
 - access permissions P:50
 - allowing creation of C:280
 - allowing groups to create P:67
 - allowing users to create P:65
 - creating C:87, P:50
 - defined AD:359, API:93, C:317, O:71, P:299, S:186
 - evicting from cache P:150
 - executing P:61
 - generating SQL views from P:60
 - listing in cache P:148
 - location in categories S:18
 - managing P:50
 - marking for scheduled caching P:139
 - modifying P:55
 - preventing creation of C:281
 - preventing groups from creating P:69
 - preventing users from creating P:68
 - qualified names for O:31
 - removing P:63
 - schemas for, generating P:57
 - SQL statements in C:90
 - tagging for on-demand caching P:146
 - unscheduling P:150
 - viewing P:55, P:57, P:59
 - viewing dependencies P:59
 - viewing details P:56
- virtual database service, configuring access permissions P:65
- virtualdatabase --allow-dbp-creation command C:280
- virtualdatabase --disallow-dbp-creation command C:281
- virtualdatabase --execute command C:282

virtualdatabase --show-tables command C:283
 virtualeschema --deploy command C:285
 virtualeschema --undeploy command C:286
 virtual schema models, See metadata models

W

web browsers

requirements for Avaki software AD:5
 setting for selecting run-as users P:27, P:54, P:92, P:224,
 P:227

web services API:2, API:8

about API:2
 access permissions API:9
 client examples API:9
 data catalog API:18
 data services API:34
 database operations API:40
 development framework API:5
 document/literal API:3, API:5
 provisioning, about P:205
 rpc/encoded API:3, API:5
 security API:8
 with MIME API:3

web services clients

defined AD:359, API:93, C:317, O:71, P:299, S:186
 requirements for API:4

web services description language, See WSDLs

whoami command C:286

whoami SOAP operation API:33

wildcard characters in searches AD:235

Windows

avoiding install pathnames with spaces AD:26
 installing Avaki in AD:23
 installing on Windows 2003 AD:22
 requirements for AD:3
 services, running under avaki local user account AD:11
 shortcuts for AD:27
 update for HTTP POST problem in web browsers AD:5
 versions supported by Avaki O:16

Windows domains

displaying for DGAS admission policies C:114
 setting for DGAS admission policies C:111
 unsetting for DGAS admission policies C:116

Windows Services list C:5

Workbench S:12

workspace directory for Avaki Studio

described S:12
 setting S:10

WS API

accessibleDBOp API:42
 accessibleDS API:36
 accessiblePath API:19
 authentication API:9
 authorization API:9
 chmod API:19
 chown API:20
 data access API:2

WS API (continued)

data services SOAP operations API:34
 executeAdHocDBOp API:43
 executeAdHocDBOpWithOutput API:44
 executeAdHocDBOpWithOutputAttach API:46
 executeAdHocDBOpWithOutputString API:47
 executeDBOp API:48
 executeDBOpBytesInput API:49
 executeDBOpGridFileInput API:50
 executeDBOpWithOutput API:50
 executeDBOpWithOutputAttach API:52
 executeDBOpWithOutputString API:53
 executeDS API:36
 fileRead API:21
 fileReadAttach API:21
 fileReadString API:22
 fileWrite API:23
 getAttributes API:23
 getDBOpOutput API:54
 getDBOpOutputAttach API:55
 getDBOpParameters API:56
 getDBOpSchema API:56
 getDBOpSchemaAttach API:57
 getDBOpSchemaString API:58
 getDSOutput API:37
 getDSOutputAttach API:38
 getDSOutputString API:38
 getDSParameters API:39
 getOutputString API:55
 getSQL API:58
 getSystemAttributes API:24
 getUserAttributes API:24
 grid server API:6
 HTTP API:7, API:8
 HTTPS API:7, API:8
 isDSAvakiXML API:40
 listDBConns API:59
 listDBOps API:59
 listDBOpsByDBConn API:60
 listDomains API:25
 listDSs API:40
 listSearches API:25
 ls API:26
 lsSize API:26
 mkdir API:27
 mkdirParents API:27
 mkdirParentsServer API:28
 mkdirServer API:29
 mv API:29
 permissions API:30
 ports API:6
 removeAttribute API:31
 rm API:31
 search API:32
 setAttribute API:32
 tester API:33
 whoami API:33

WS clients, See web services clients

WSDLs
 about API:2
 as SOAP contracts API:3
 AvakiAPI.disco discovery file for .NET clients API:3
 AvakiAPIDocLit.wsdl API:3
 AvakiAPIRpcEnc.wsdl API:3
 AvakiAPIWithMIMEDocLit.wsdl API:3
 AvakiAPIWithMIMERpcEnc.wsdl API:3
 choosing API:5
 document/literal API:3
 editing API:6
 locations of API:6
 provided by Avaki API:11
 rpc/encoded API:3
 standards compliance API:1
WSDLs directory O:35

X

X Window System libraries required for Avaki install on
 Unix AD:16
XA drivers, configuring for database connectors C:69, P:7, S:36
Xalan C:265, C:268, C:271, P:220, P:224, P:227
XWorkHandler class for data services P:189
XWorkUnit interface for data services P:189
XML data in Avaki O:11, P:273
XML indent size property AD:142
XML schema
 Avaki rowset
 class-name element P:279
 column-display-size element P:279
 column-index element P:279

XML schema (*continued*)
 core schema P:277
 overview P:277
 rowset-specific schema P:279
 sample schema P:280
data service
 class element P:261
 coherenceWindow element P:261
 dataService element P:262
 description element P:263
 initParameter element P:263
 inputParameter element P:264
 inputSource element P:265
 inputStream element P:266
 isList element P:266
 jarurl element P:267
 logicBox element P:268
 name element P:269
 outputStream element P:269
 ref element P:270
 target element P:270
 type element P:270
 urlLogicBox element P:271
 value element P:272
 values element P:272

XSLT
 in view generators, when not to use P:242
 using in data service plug-ins P:180
 See also Xalan, Saxon