

Management Guide

NetIQ Analysis Center

October 2013



Legal Notice

THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT ARE FURNISHED UNDER AND ARE SUBJECT TO THE TERMS OF A LICENSE AGREEMENT OR A NON-DISCLOSURE AGREEMENT. EXCEPT AS EXPRESSLY SET FORTH IN SUCH LICENSE AGREEMENT OR NON-DISCLOSURE AGREEMENT, NETIQ CORPORATION PROVIDES THIS DOCUMENT AND THE SOFTWARE DESCRIBED IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. SOME STATES DO NOT ALLOW DISCLAIMERS OF EXPRESS OR IMPLIED WARRANTIES IN CERTAIN TRANSACTIONS; THEREFORE, THIS STATEMENT MAY NOT APPLY TO YOU.

For purposes of clarity, any module, adapter or other similar material ("Module") is licensed under the terms and conditions of the End User License Agreement for the applicable version of the NetIQ product or software to which it relates or interoperates with, and by accessing, copying or using a Module you agree to be bound by such terms. If you do not agree to the terms of the End User License Agreement you are not authorized to use, access or copy a Module and you must destroy all copies of the Module and contact NetIQ for further instructions.

This document and the software described in this document may not be lent, sold, or given away without the prior written permission of NetIQ Corporation, except as otherwise permitted by law. Except as expressly set forth in such license agreement or non-disclosure agreement, no part of this document or the software described in this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, or otherwise, without the prior written consent of NetIQ Corporation. Some companies, names, and data in this document are used for illustration purposes and may not represent real companies, individuals, or data.

This document could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein. These changes may be incorporated in new editions of this document. NetIQ Corporation may make improvements in or changes to the software described in this document at any time.

U.S. Government Restricted Rights: If the software and documentation are being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), in accordance with 48 C.F.R. 227.7202-4 (for Department of Defense (DOD) acquisitions) and 48 C.F.R. 2.101 and 12.212 (for non-DOD acquisitions), the government's rights in the software and documentation, including its rights to use, modify, reproduce, release, perform, display or disclose the software or documentation, will be subject in all respects to the commercial license rights and restrictions provided in the license agreement.

© 2013 NetIQ Corporation and its affiliates. All Rights Reserved.

For information about NetIQ trademarks, see <https://www.netiq.com/company/legal/>.

Contents

About This Book and the Library	5
About NetIQ Corporation	7
1 Introducing Analysis Center	9
2 Managing Data Sources	11
2.1 Adding a Data Source	11
2.2 Understanding the Data Sources View	12
2.3 Understanding the ETL and OLAP Processing Jobs	15
2.4 Changing the OLAP Job Schedule	15
2.5 Changing Data Source Properties	16
2.6 Enabling and Disabling a Data Source	16
2.7 Removing a Data Source	16
3 Reviewing AppManager Reports	19
3.1 Application Specific	19
3.2 Exchange	19
3.3 VoIP	20
3.4 SQL Server	20
3.5 Comparison	20
3.6 Dashboards	21
3.7 Events	22
3.8 Performance	22
3.9 Service Levels	24
3.10 Statistical Service Levels	26
3.11 Trend and Prediction	26
3.12 Watch List	27
4 Using Report Templates	29
5 Adding Dynamic Legends to AppManager Knowledge Scripts	31
5.1 What is a Dynamic Legend?	31
5.2 Why Add Dynamic Legends?	31
5.3 Sample Dynamic Legend	32
5.4 Required XML Tags	32
5.5 Adding a Dynamic Legend to a Knowledge Script	33
6 Working with Legend Rules	35
6.1 Identifying a Legend Rule	35
6.2 Removing a Legend Rule	36
6.3 Creating the Removal XML File	36
6.4 Issuing xmlcheckin	37
6.5 Adding a Legend Rule	38

About This Book and the Library

The *NetIQ Analysis Center Management Guide* provides conceptual information about the enterprise-wide reporting of the NetIQ Analysis Center product, based on data from NetIQ AppManager repositories.

Intended Audience

This book provides information for individuals responsible for understanding Analysis Center concepts, and for system administrators and users responsible for installing, configuring, and using Analysis Center for reporting on AppManager data.

Other Information in the Library

The library provides the following information resources:

[*NetIQ Analysis Center User Guide*](#)

Provides conceptual information about Analysis Center. This book also provides an overview of the Analysis Center user interfaces and the Help.

Online Help

Provides context-sensitive information and step-by-step guidance for common tasks, as well as definitions for each field on each window.

Conventions

The library uses consistent conventions to help you identify items throughout the documentation. The following table summarizes these conventions.

Convention	Use
Bold	<ul style="list-style-type: none">♦ Window and menu items♦ Technical terms, when introduced
<i>Italics</i>	<ul style="list-style-type: none">♦ Book and CD-ROM titles♦ Variable names and values♦ Emphasized words
Fixed Font	<ul style="list-style-type: none">♦ File and folder names♦ Commands and code examples♦ Text you must type♦ Text (output) displayed in the command-line interface
Brackets, such as [value]	<ul style="list-style-type: none">♦ Optional parameters of a command
Braces, such as {value}	<ul style="list-style-type: none">♦ Required parameters of a command
Logical OR, such as value1 value2	<ul style="list-style-type: none">♦ Exclusive parameters. Select one parameter.

About NetIQ Corporation

NetIQ Corporation, an Attachmate business, is a leading provider of comprehensive systems and security management solutions that help enterprises maximize IT service delivery and efficiency. With more than 12,000 customers worldwide, NetIQ solutions yield measurable business value and results that dynamic organizations demand. Best-of-breed solutions from NetIQ Corporation help IT organizations deliver critical business services, mitigate operational risk, and document policy compliance. The company's portfolio of award-winning management solutions includes IT Process Automation, Systems Management, Security Management, Configuration Control and Enterprise Administration. For more information, please visit www.netiq.com

Contacting NetIQ Corporation

Please contact us with your questions and comments. We look forward to hearing from you.

Sales Email:	info@netiq.com
---------------------	--

Telephone:	1-713-418-5555 (United States) +353 (0) 91-782-677 (Europe, Middle East, and Africa) For other locations, see our Support Contact Information Web site at www.netiq.com/support
-------------------	---

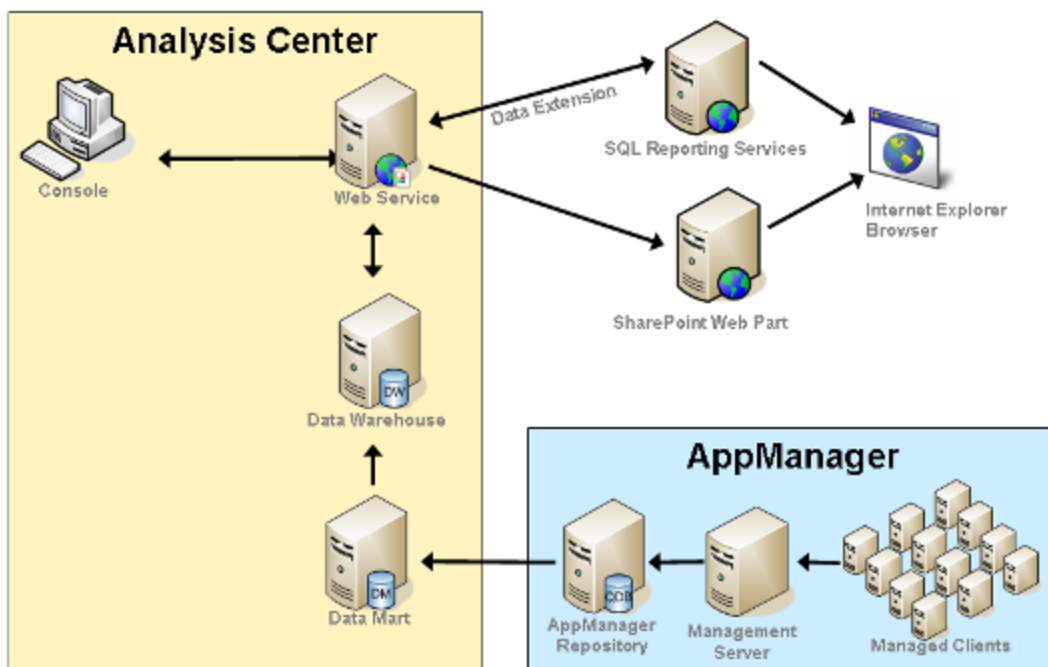
Support Web Site:	www.netiq.com/support
--------------------------	--

1 Introducing Analysis Center

Analysis Center provides enterprise-wide reporting based on data from NetIQ AppManager repositories. Analysis Center for AppManager copies data from AppManager repositories to Data Marts, and from the Data Marts to the Analysis Center Data Warehouse. A Data Mart is a SQL Server database that Analysis Center creates for each AppManager Data Source. The Analysis Center Console sends queries to the Data Warehouse to generate reports.

Analysis Center imports raw data from multiple AppManager repositories, with data from single or multiple repositories in one report. Analysis Center transforms this data into useful information about your computing infrastructure, as opposed to the HTML reports in AppManager, which are based on every data point that you collect. Analysis Center publishes information in the form of reports that you can use as-is or alter to suit your needs. In a single report, you can select multiple Data Sources. You can specify the AppManager repositories by adding a Data Source for each AppManager repository containing data you want to make available in reports. For more information about adding Data Sources, see [Section 2.1, "Adding a Data Source," on page 11](#).

You can generate reports based on the data from multiple repositories or selected views, and view them in various folders in the Navigation pane. The following model illustrates the Analysis Center for AppManager architecture.



For more information about Analysis Center architecture, see the [NetIQ Analysis Center User Guide](#).

2 Managing Data Sources

A Data Source is an AppManager repository that contains data you want to see in reports.

- ♦ [Section 2.1, “Adding a Data Source,” on page 11](#)
- ♦ [Section 2.2, “Understanding the Data Sources View,” on page 12](#)
- ♦ [Section 2.3, “Understanding the ETL and OLAP Processing Jobs,” on page 15](#)
- ♦ [Section 2.4, “Changing the OLAP Job Schedule,” on page 15](#)
- ♦ [Section 2.5, “Changing Data Source Properties,” on page 16](#)
- ♦ [Section 2.6, “Enabling and Disabling a Data Source,” on page 16](#)
- ♦ [Section 2.7, “Removing a Data Source,” on page 16](#)

2.1 Adding a Data Source

Before you can generate reports reflecting AppManager data, you need to add at least one Data Source. The **Add Data Source** Wizard guides you through the process of defining a Data Source.

By default, the wizard:

- ♦ Verifies the Data Source is valid and can be reached.
- ♦ Activates the Data Source, moving it from the Pending Data Sources to the Data Sources folder.
- ♦ After the ETL and OLAP jobs run, you can generate reports against the data. The ETL and OLAP jobs process the data from the AppManager repository to make it available in reports. For more information, see [Section 2.3, “Understanding the ETL and OLAP Processing Jobs,” on page 15](#).

However, you can disable the validity check and perform the Activate operation separately. This is useful in cases where you are adding Data Sources with known problems that you want to resolve at a later time.

You can add a Data Source for an AppManager repository that uses the default or a named instance. For more information about using a named instance instead of a default SQL installation, see the [“Using Named Instances of SQL Server”](#) section in the *NetIQ Analysis Center User Guide*.

To add a Data Source:

- 1 Start the Analysis Center Console. For more information, see the *NetIQ Analysis Center User Guide*.
- 2 From the **Data Sources** menu, select **Add**.
- 3 Analysis Center displays the **Add Data Source Wizard**. Complete the fields by following the instructions on the pages. For more information about fields, see the Help.

- 4 In the **Confirmation** page, select **Set as an enabled Data Source** to enable the Data Source. An enabled Data Source supplies data to the Analysis Center Data Warehouse. If you do not select this option, Analysis Center adds only the Data Source and does not perform other operations such as a pre-check.
- 5 When you are satisfied with the Data Source settings, click **Add** to add the Data Source.

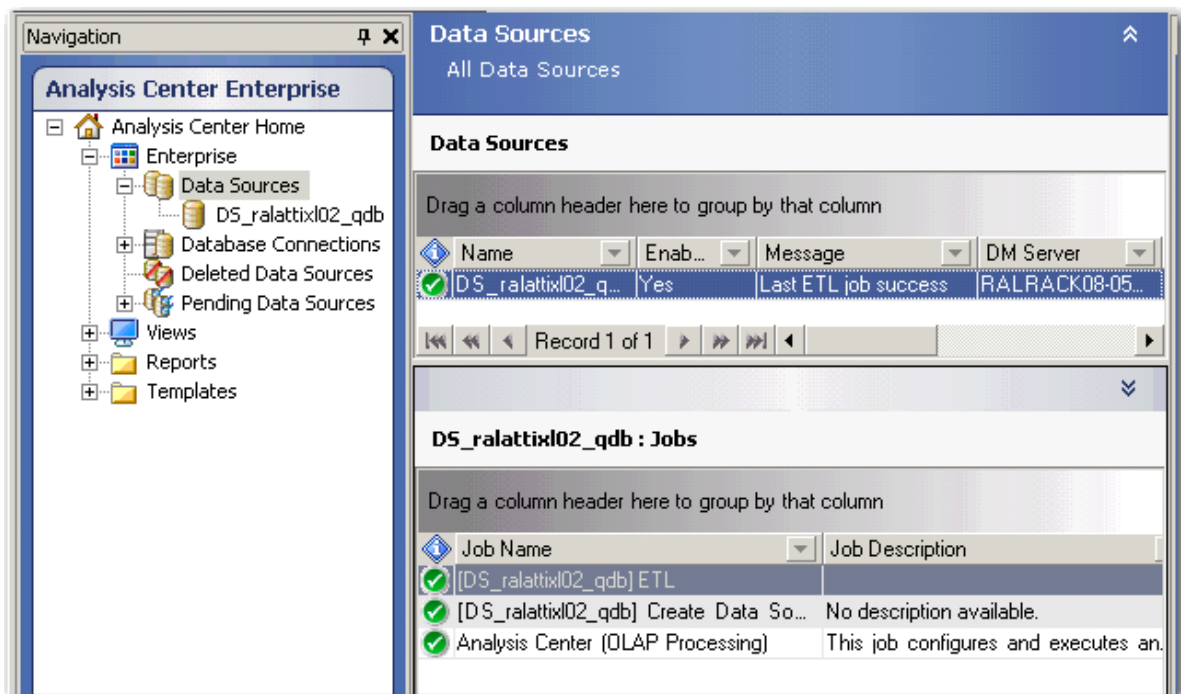
The time it takes to create a Data Source depends on such variables as system resources, network connections, and the amount of data being processed. The process of creating a Data Source involves creating the related Data Mart database, creating the ETL job, and storing the meta data. You use the meta data to manage how the ETL and the OLAP processing jobs process the AppManager data. For more information about ETL and OLAP processing jobs, see the [Section 2.3, “Understanding the ETL and OLAP Processing Jobs,”](#) on page 15.

While pre-check background processing is taking place, Analysis Center places the Data Source in the Pending Data Sources folder. If you select the check box to automatically activate the Data Source and it passes the pre-check, then when the pre-check processing is complete, Analysis Center places the Data Source in the Data Sources folder in the Navigation pane. You might need to refresh the window to see the Data Source.

2.2 Understanding the Data Sources View

When you select the Data Sources folder in the Navigation pane, the upper portion of the Results pane displays a list of active Data Sources. The lower portion of the Results pane displays a list of SQL Server Agent jobs related to the selected Data Source. You can right-click a job to display commands to start, stop, enable, or disable the job.

The following picture illustrates the Data Sources view.



NOTE: You cannot generate reports against the Data Source until Analysis Center completes the process of adding the Data Source, creates the cube, and populates the cube with data from the Data Source.

The Data Sources pane displays the following information.

Column Name	Contents
Status	Indicates whether the Data Source is available to provide data for reports. A green check means the Data Source is available. A red X means the Data Source is not available.
Name	Specifies the name of the Data Source.
Enabled	Indicates whether the Data Source is enabled.
Message	Displays a message relevant to the Data Source.
DM Server	Indicates the SQL Server for the Data Mart database.
DM Database	Indicates the name of the Data Mart database.
DS Server	Indicates the SQL Server for the AppManager repository.
DS Database	Indicates the name of the AppManager repository database.
Updated Time	Indicates one of the following: <ul style="list-style-type: none">♦ The time at which the last attempt to create a Data Source either succeeded or failed.♦ The last time the properties of a Data Source were changed.♦ The last time the ETL job ran.
Oldest Data	Indicates the date of the oldest data copied from the AppManager repository database.
Newest Data	Indicates the date of the most-recent data copied from the AppManager repository database.
Days Collected	Indicates the number of days of data copied from the AppManager repository database.

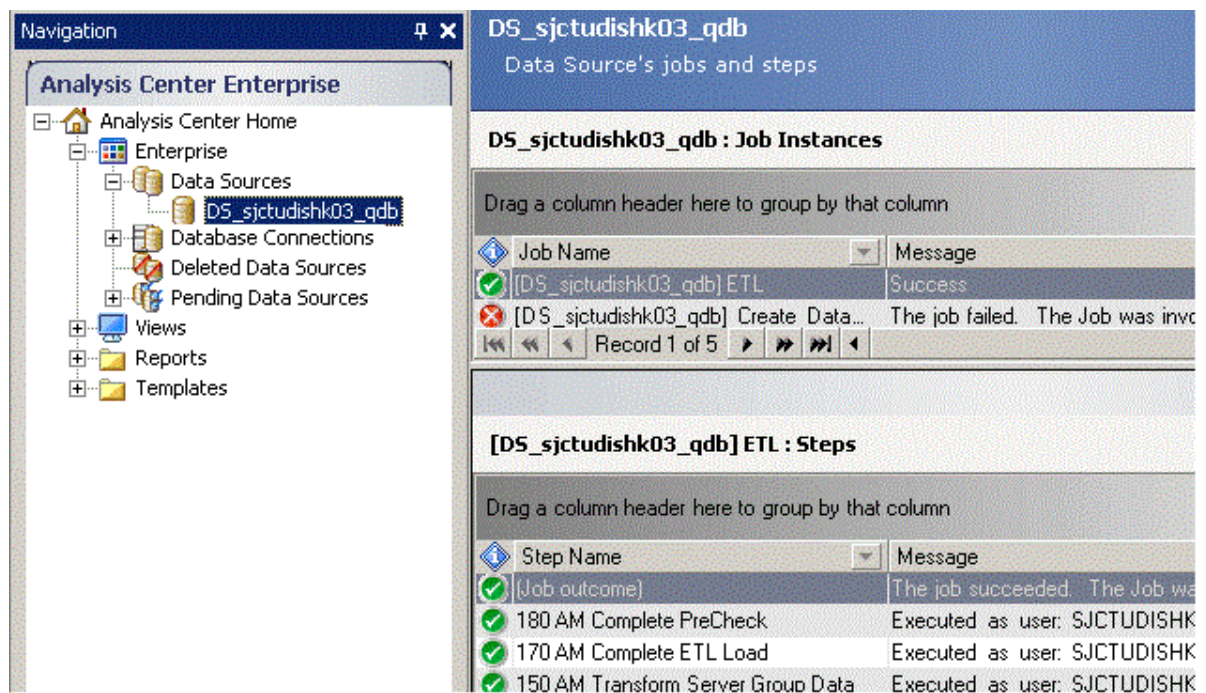
The Jobs pane displays the following information.

Status	Indicates whether the Data Source is available to provide data for reports. A green check means the Data Source is available. A red X means the Data Source is not available.
Job Name	Indicates the name of the SQL Server Agent job.
Job Description	A short description of the SQL Server Agent job.
Message	Displays a message relevant to the Data Source.
Enabled	Indicates whether or not the job is enabled.

Schedule	Indicates the schedule on which the job runs.
Next Run	Indicates the next date and time at which the job is scheduled to run.
Last run	Indicates the last date and time at which the job ran.
Elapsed Time	Indicates the duration of the last run of the job.

When you select an individual Data Source, the upper portion of the Results pane displays a list of SQL Server Agent job instances related to that Data Source. The lower portion of the Results pane displays a list of steps related to the selected job.

The following picture illustrates this view.



You can right-click any job instance to display commands to start, stop, enable, or disable the job. When you select an individual Data Source, the view displays the following information:

Column Name	Contents
Jobs pane	Indicates the jobs that were run and the overall status of the job.
Status	Indicates whether the Data Source is available to provide data for reports. A green check means the Data Source is available. A red X means the Data Source is not available.
Job Name	Indicates the name of the SQL Server Agent job.
Message	Displays a message relevant to the Data Source.
Run Time	Indicates the last date and time at which the job ran.
Elapsed Time	Indicates the duration of the job.

Column Name	Contents
Steps pane	Indicates the steps that comprise the job and the status of each step. Use this pane to discover in which step a failed job failed.
Status	Indicates whether the Data Source is available to provide data for reports. A green check means the Data Source is available. A red X means the Data Source is not available.
Step Name	Indicates the name of the SQL Server Agent job step.
Message	An Analysis Center message relevant to the step.
Run Time	Indicates the last date and time at which the step ran.
Elapsed Time	Indicates the duration of the step.

2.3 Understanding the ETL and OLAP Processing Jobs

An ETL job is a SQL job that extracts data from the Data Source to the Data Mart and the Data Warehouse. An OLAP job is a SQL job that processes the data from the Data Warehouse, and then puts the data into cubes and dimensions for reporting in Analysis Center Console. The ETL job runs according to the schedule you define when you add a Data Source, that is, every 59 minutes. There is a separate ETL job for each Data Source.

The ETL job runs only if you select the **Set as an Enabled Data Source** check box when adding the Data Source. For more information about enabling a Data Source, see [Section 2.6, “Enabling and Disabling a Data Source,”](#) on page 16.

The Data Warehouse copies all the data you specify in the initial range, and thereafter, each iteration of the ETL job copies seven days of data. For example, if the initial range includes February, the first iteration of the ETL job copies the most recent week of data, the next iteration copies data from February 21 to February 15, and so on, until it copies all the data from February.

In addition to copying the initial range of data, each iteration of the ETL job also copies any new data. After copying the initial range of data, each iteration of the ETL job copies only data that is new since the previous iteration of the job.

There is one OLAP processing job for the Data Warehouse. This job is owned by the Data Warehouse SQL Server Agent.

The OLAP Processing job takes data copied by each ETL job and inserts it into a cube. This job runs every 15 minutes by default, updating each cube based on the most recent data introduced by the relevant ETL job. If no new data has been introduced, the OLAP processing job does not run.

2.4 Changing the OLAP Job Schedule

To change the OLAP Job Schedule from the Analysis Center Console, click **Tools > Options > Enterprise Options > Data Warehouse**. You can also start the OLAP processing job manually from SQL Server Enterprise Manager.

2.5 Changing Data Source Properties

You can change the Data Source properties to rename the Data Source, change the database type, or modify the schedule of how often the ETL job collects data to process for reports.

To change the properties for a Data Source:

- 1 Start the Analysis Center Console. For more information, see the [NetIQ Analysis Center User Guide](#).
- 2 From the **Data Sources** menu, select **Properties**.
- 3 Click the required tab to change the Data Source properties.
- 4 Click **OK** to save your changes.

2.6 Enabling and Disabling a Data Source

You can disable and enable Data Sources. When you disable the Data Source, Analysis Center disables the ETL jobs associated with the Data Source, and therefore, does not collect data for the disabled Data Source. Analysis Center also does not delete disabled Data Sources, which allows you to enable the Data Source at a later date.

When you disable a Data Source, you can still generate reports based on the data from the Data Source up to the time the last ETL job ran. However, the ETL job stops running, and consequently stops copying data from the original database to the Data Warehouse.

You can manually enable the ETL job for a Data Source that you have disabled in the Console. However, even though the ETL job is running, the OLAP Processing job does not process data from that Data Source. To process the data, enable the Data Source from the Console.

To enable or disable a Data Source:

- 1 Start the Analysis Center Console. For more information, see the [NetIQ Analysis Center User Guide](#).
- 2 In the Navigation pane, click **Data Sources**.
- 3 Select the Data Source.
- 4 Enable or disable the Data Source as follows:
 - ♦ To enable the Data Source, right-click and select **Enable**.
 - ♦ To disable the Data Source, right-click and select **Disable**.

2.7 Removing a Data Source

You can remove a Data Source and add it back, if necessary. When you remove a Data Source, Analysis Center removes the associated ETL job, SSIS packages, and the Data Mart database. Analysis Center marks the Data Source for removal, which means that the next time the OLAP processing job runs, it deletes the data so that it would no longer be available for reports. The OLAP processing job runs at approximately 30 minute intervals. To add a Data Source immediately after you have removed it, you can manually start the OLAP processing job from the SQL Server Enterprise Manager or the Analysis Center Console. The removal process performs the following actions:

- ♦ Displays the Data Source name in the Deleted Data Sources folder.

- ♦ Creates and executes the [Data_Source_Name] Create Data Source Data Marts job to disable both the Data Source and the [Data_Source_Name] ETL job on the Data Mart database.
- ♦ Instructs the Analysis Center OLAP Processing job to perform the following actions for the Data Source.
 - ♦ Removes the physical, regular cubes from the AC_OLAP database.
 - ♦ Deletes the ETL job on the Data Mart SQL Server.
 - ♦ Deletes the ETL job's SSIS packages on the Data Mart SQL Server.
 - ♦ Deletes the Data Mart database, if you chose to remove it.
 - ♦ Redefines the virtual cubes so that the physical, regular cubes for the Data Source are no longer included.
 - ♦ Marks the Data Source as deleted, although it still exists in the AC_Configuration database. However, if no ETL jobs have run against the Data Source, then the Data Source was not fully loaded into the OLAP Data Source dimension. In this scenario, the OLAP processing job completely removes the Data Source.

NOTE: If the SQL Server agent on the Data Mart used by the Data Source is not running at the time that you remove it, the SQL Server agent cannot delete the Data Mart, and the Data Source remains in the console. You can avoid this potential problem by verifying that the SQL Server agent on the Data Mart is running before you remove the Data Source. If the problem has already occurred on the Data Mart computer, restart the SQL Server agent.

To remove a Data Source:

- 1 Start the Analysis Center Console. For more information, see the [NetIQ Analysis Center User Guide](#).
- 2 In the Navigation pane, click **Data Sources** and select **Remove**.
- 3 Click **Yes** to confirm the removal.
- 4 If you want to remove the Data Mart database, click **Yes**.

3 Reviewing AppManager Reports

Analysis Center provides reports specially designed to work with AppManager data and data streams. You can find AppManager reports in the Navigation pane in **Reports > AppManager**.

This chapter briefly describes the purpose of each Analysis Center report within the AppManager category.

- ♦ [Section 3.1, “Application Specific,” on page 19](#)
- ♦ [Section 3.2, “Exchange,” on page 19](#)
- ♦ [Section 3.3, “VoIP,” on page 20](#)
- ♦ [Section 3.4, “SQL Server,” on page 20](#)
- ♦ [Section 3.5, “Comparison,” on page 20](#)
- ♦ [Section 3.6, “Dashboards,” on page 21](#)
- ♦ [Section 3.7, “Events,” on page 22](#)
- ♦ [Section 3.8, “Performance,” on page 22](#)
- ♦ [Section 3.9, “Service Levels,” on page 24](#)
- ♦ [Section 3.10, “Statistical Service Levels,” on page 26](#)
- ♦ [Section 3.11, “Trend and Prediction,” on page 26](#)
- ♦ [Section 3.12, “Watch List,” on page 27](#)

3.1 Application Specific

The Application Specific reports are designed to work specifically with the data and data streams of specific AppManager modules. Analysis Center provides reports for some modules.

You can create other folders within the Application Specific folder to contain reports that you have customized to work with the data and data streams of AppManager modules.

3.2 Exchange

Analysis Center supports Exchange 2000, Exchange 2003, and Exchange 2007 reports. These reports enable you to examine and compare Exchange Server performance data.

For more information about the supported Exchange reports, see the Exchange Server documentation or contact NetIQ Technical Support.

3.3 VoIP

Analysis Center supports VoIP reports. These reports enable you to examine the Cisco VoIP reports, Nortel VoIP reports, and VoIP Quality reports.

For more information about the supported VoIP reports, see the VoIP documentation or contact NetIQ Technical Support.

3.4 SQL Server

The SQL Server reports enable you to examine and compare SQL Server performance data.

Report Name	Description
SQL Server Application Data	Compares and correlates SQL Server performance and event data.
SQL Server Performance	Examines SQL Server performance data.

3.5 Comparison

Comparison reports enable you to compare AppManager data over different time periods and over different computers and computer groups.

Report Name	Description
Compare Average Performance Data by Time Range	Compares the average values for the same data over two different time periods.
Compare CPU Utilization by Four Time Ranges	Compares the percentage of CPU utilization among computers and computer groups over four different time periods. Each selected computer or computer group represents a row of data, and each time range represents a column of data.
Compare Day of Week By Time Range	Compares activity for days of the week over different time ranges. This report is useful for learning how performance varies for each day of the week during specific time ranges.
Compare Disk Space Utilization by Four Time Ranges	Compares the percentage of disk space utilization among computers and computer groups over four different time periods. Each selected computer or computer group represents a row of data, and each time range represents a column of data.
Compare Memory Utilization by Four Time Ranges	Compares the percentage of memory utilization among computers and computer groups over four different time periods. Each selected computer or computer group represents a row of data, and each time range represents a column of data.
Compare Server Event Occurrence by Time Range	Compares the number of events per server over two different time periods. This report is useful, for example, to see whether hardware upgrades have improved performance as anticipated.

Report Name	Description
Compare Server Event Occurrence Percentage Change by Time Range	Compares the percentage change in events per server from one time period to another. This report is useful, for example, to see whether network upgrades have a positive impact on the transmission of messages.
Compare Server Performance Over Time	Compares the percentage change in events per server from one time period to another. This report is useful, for example, to see whether network upgrades have a positive impact on the transmission of messages.
Compare Server Performance with Multiple Measures	Compares multiple statistics derived from the same data.
Performance Data Percentage Change by Time Range	Compares the percentage change in performance data from one time period to another.

3.6 Dashboards

Dashboard reports enable you to group a number of reports together on the same page, which makes it easy to see related data at a glance. You can use all dashboard reports as the basis for others by adding and deleting the reports in which you are interested, and then saving it under another name.

Report Name	Description
Dashboard	Contains the following reports: <ul style="list-style-type: none"> ♦ Average CPU Utilization by Server. ♦ Average Memory Utilization by Server Over Time. ♦ Disk Space Utilization by Server Over Time. Use this dashboard to create a collection of individual Analysis Center reports.
System Overview Dashboard	Contains the following reports: <ul style="list-style-type: none"> ♦ Average CPU Utilization by Server Over Time. ♦ Average Memory Utilization by Server Over Time. ♦ Disk Space Utilization by Server Over Time. ♦ System Uptime Pie Chart, and System Uptime. Use this dashboard to create a general system overview for the computers in which you are interested.

3.7 Events

The Events reports enable you to examine the event occurrences and event counts by computer, computer group, Knowledge Script, and severity. These numbers might vary dramatically from one another because of differences in event collapsing preference settings for the AppManager repository (with event collapsing, you can have a single parent and child with any number of count increments).

Report Name	Description
Server Event	Examines the number of events per computer or computer group. Displays values for event occurrence and for event count.
Server Event by KS	Examines the overall number of events per Knowledge Script and/or Knowledge Script category raised from the computers you select. Displays aggregate value for event occurrence and/or event count for each selected Knowledge Script or Knowledge Script category.
Server Event by Severity	Examines the overall number of events per severity level raised from the computers you select, for the overall time period you specify. Displays separate values for event occurrence and for event count.
Server Event Over Time	Examines the overall number of events per time interval for the computers you select. Displays separate values for event occurrence and for event count.

3.8 Performance

The Performance reports enable you to examine and compare performance data by computer, computer group, Knowledge Script, server, or Data Source.

Report Name	Description
Average CPU Utilization by Server	<p>Use this report to examine the percentage of CPU utilization by individual computers or groups of computers for the time period you specify.</p> <p>Use the Metric context to specify whether the report contains values for overall utilization (all CPUs on a computer) or distinct values for each CPU. One value is returned for each computer or group for the time period.</p>
Average CPU Utilization by Server Over Time	<p>Use this report to examine the percentage of CPU utilization by individual computers or groups of computers for the time period you specify.</p> <p>Use the Metric context to specify whether the report contains values for overall utilization (all CPUs on a computer) or distinct values for each CPU. One value is returned for each computer or group for each interval of the time period.</p>

Report Name	Description
Average Memory Utilization by Server Over Time	Examines the average use of physical memory, virtual memory, and paging files for individual computers and/or groups of computers over the time period you specify.
Compare Performance Data By Data Source Over Time	Compares performance data from multiple data sources over the time period you specify. Data sources are listed side-by-side (as columns) in the report so you can easily compare values.
Compare Performance Data by KS Over Time	Compares performance data by the Knowledge Scripts from which they were generated. Knowledge Scripts are listed side-by-side (as columns) in the report so you can easily compare values.
CPU and Memory Utilization Over Time	Compares the percentages of CPU and memory utilization by individual computers or groups of computers.
Disk Space Utilization by Machine	Examines the percentage of used disk space and the MBs of available disk space for individual computers or groups of computers.
Disk Space Utilization by Server Over Time	Examines the percentage of used disk space for individual computers or groups of computers, for the time period and interval you specify.
Performance Data Filtered by Data Source Over Time	Filters data by data source. Contains one set of values for each metric you select based on the values from all Data Sources.
Performance Data Filtered by KS Over Time	Filters data by the Knowledge Scripts from which they were generated.
System Downtime Over Time	Examines the downtime for individual computers or groups of computers for the time period you specify. Expresses downtime in seconds. Displays one aggregated value for all selected computers for each time increment.
Total System Downtime by Server	Examines the downtime for individual computers or groups of computers for the time period you specify. Expresses downtime in seconds. Displays one value for each selected computer or computer group.

3.9 Service Levels

The reporting capability of Analysis Center enables organizations to demonstrate the value of IT and how well IT is aligned with business objectives. Run the Service Levels reports to examine server availability and G-A-P performance data.

Report Name	Description
Good-Acceptable-Poor CPU Utilization	Examines relative levels of CPU utilization based on the data you have collected from your computers. Displays overall values for each selected computer.
Good-Acceptable-Poor CPU Utilization Over Time	Examines relative levels of CPU utilization based on the data you have collected from your computers. Displays a set of values per time increment for each selected computer.
Good-Acceptable-Poor Disk Space Utilization	Examines the relative levels of disk space utilization based on the data you have collected from your computers. Displays overall values for each selected computer.
Good-Acceptable-Poor Disk Space Utilization Over Time	Examines the relative levels of disk space utilization based on the data you have collected from your computers. Displays a set of values per time increment for each selected computer.
Good-Acceptable-Poor Memory Utilization	Examines the relative levels of memory utilization based on the data you have collected from your computers. Displays overall values for each selected computer.
Good-Acceptable-Poor Memory Utilization Over Time	Examines the relative levels of memory utilization based on the data you have collected from your computers. Displays a set of values per time increment for each selected computer.
Good-Acceptable-Poor Performance Data	Examines relative levels of performance based on the data you have collected from your computers. Displays overall values for each selected computer.
Good-Acceptable-Poor Performance Data By Machine	Examines relative levels of performance based on the data you have collected from your computers. Displays overall values for each selected computer.
Good-Acceptable-Poor Performance Data Over Time	Examines relative levels of performance based on the data you have collected from your computers. Displays a set of values per time increment for each selected computer.
Performance Relative to Target	Tracks performance relative to a benchmark. Displays values as the percentages by which the target value was exceeded.

Report Name	Description
System Uptime	<p>Lists the total system uptime and downtime for computers and/or computer groups over the time period you specify. Expresses uptime and downtime in hours and percentages. Returns a value for:</p> <ul style="list-style-type: none"> ♦ System Total Time ♦ System Downtime ♦ System Uptime ♦ System Downtime (%) ♦ System Uptime (%)
System Uptime by Hour	<p>Examines the percentage of uptime and downtime by time interval for computers and/or computer groups over the time period you specify. Displays two values (percentage of uptime and downtime) per time interval for all selected computers.</p>
System Uptime By Server	<p>Examines the percentage of uptime and downtime per computer and/or computer group for the time period you specify. Displays two values per computer or computer group for the time period you specify.</p>
System Uptime Over Time	<p>Examines the percentage of uptime and downtime by time interval for computers and/or computer groups over the time period you specify. Displays two values (percentage of uptime and downtime) per time interval for all selected computers.</p>
System Uptime Pie Chart	<p>Examines the percentage of uptime and downtime for an individual computer or computer group for the time period you specify. Displays two overall values (percentage of uptime and downtime) for the time period you specify.</p> <p>NOTE: : You can include multiple computers or groups, but a single pie chart is not suited to illustrating multiple 100% totals. If you want to use this report for multiple computers/groups, use a different chart type.</p>

3.10 Statistical Service Levels

These reports are similar to the Service Level reports. The Statistical reports use statistical data instead of raw count data to determine service levels. These reports take considerably less time to generate. Consider using Statistical Service Level reports rather than Service Level reports where possible to improve performance.

Use the Parameters on the Properties tab to set thresholds for what you consider good and poor performance:

Report Name	Description
Statistical Good-Acceptable-Poor CPU Utilization	Examines relative levels of CPU utilization based on the data you have collected from your computers. Displays overall values for each selected computer.
Statistical Good-Acceptable-Poor CPU Utilization Over Time	Examines relative levels of CPU utilization based on the data you have collected from your computers. Displays a set of values per time increment for each selected computer.
Statistical Good-Acceptable-Poor Disk Space Utilization	Examines the relative levels of disk space utilization based on the data you have collected from your computers. Displays overall values for each selected computer.
Statistical Good-Acceptable-Poor Disk Space utilization Over Time	Examines the relative levels of disk space utilization based on the data you have collected from your computers. Displays a set of values per time increment for each selected computer.
Statistical Good-Acceptable-Poor Memory Utilization	Examines the relative levels of memory utilization based on the data you have collected from your computers. Displays overall values for each selected computer.
Statistical Good-Acceptable-Poor Memory Utilization Over Time	Examines the relative levels of memory utilization based on the data you have collected from your computers. Displays a set of values per time increment for each selected computer.
Statistical Good-Acceptable-Poor Performance Data By Machine	Examines relative levels of performance based on the data you have collected from your computers. Displays overall values for each selected computer.
Statistical Good-Acceptable-Poor Performance Data Over Time	Examines relative levels of performance based on the data you have collected from your computers. Displays a set of values per time increment for each selected computer.

3.11 Trend and Prediction

The Trend and Prediction reports enable you to plot a trend line and a normalized trend line.

Report Name	Description
Performance Data Trend and Prediction	Plots lot a trend line based on a range of all available data.

Report Name	Description
Performance Data Trend and Prediction (Normalized)	Plots a trend line based on a range of existing, normalized data, that is, data that excludes extreme, isolated values.

3.12 Watch List

The Watch List reports enable you to find Top N Computers and Bottom N computers based on Disk Space Available, Disk Utilization, Event performance, and so on.

Report Name	Description
Bottom N Computers by Disk Space Available	Lists the computers that have the lowest amount of available disk space (in MB).
Bottom N Computers by Disk Utilization	Lists the computers that have the lowest percentage of disk utilization.
Bottom N Downtime Server	Lists the bottom N computers with the lowest amount of downtime. Expresses downtime in seconds.
Bottom N Performance Server	Lists the bottom N computers with the lowest values for the metrics you select.
Bottom N Server with Least Event Occurrence	Lists the computers that are generating the least number of events. Displays values for event occurrence and for event count.
Top N Computers by Disk Space Available	Lists the computers that have the greatest available disk space (in MB).
Top N Computers by Disk Utilization	Lists the computers that have the highest percentage of disk utilization.
Top N Downtime Server	Lists the computers that have the greatest amount of downtime for the period you specify. Expresses downtime in seconds.
Top N KS with Most Event Occurrence	Lists the Knowledge Scripts that are raising the most events from the computers you select. Displays values for event occurrence and for event count.
Top N Performance Server	Lists the top N computers with the highest values for the metrics you select.
Top N Server with Most Event Occurrence	Lists the computers that are generating the greatest number of events. Displays values for event occurrence and for event count.

4 Using Report Templates

Analysis Center includes a number of templates that you can use as the basis for your own reports. The templates are located in the **Templates** folder in the Navigation pane.

- ♦ Although each template includes all the context controls, the data in each report is determined primarily by one context for the columns, and one for the rows. The template names indicate which contexts are implemented. For example, the *Metric by Day of Week* template has metric data in each column, aggregated by weekday in each row.
- ♦ Although the contexts referenced in the template name determine the column and row data, you can filter that data using the other context controls.

For example, with *Metric by Day of Week*, you can include any number of metrics as columns, and include or exclude any of the week days. You can also set any time range for the report, select any measure, and include data from any available Data Sources.

All the templates follow this same logic. The short description of the template under the title tells you which context is represented by column and which context is represented by row. After you have configured a template to report on the data in which you are interested, use the *Save Report As* command to save it to an appropriate name and location under the Reports folder. For more information, see the [NetIQ Analysis Center User Guide](#). You must save the template before you can deploy it as a report.

Analysis Center supports drill-down reports for AppManager. For more information, see the [NetIQ Analysis Center User Guide](#).

5 Adding Dynamic Legends to AppManager Knowledge Scripts

This chapter provides a high-level explanation of how to modify AppManager Knowledge Scripts to generate dynamic legends. To implement this solution, you should be familiar with the process of making modifications to existing Knowledge Scripts or creating new Knowledge Scripts.

For more information about customizing Knowledge Scripts, see the [AppManager documentation](#).

- ♦ [Section 5.1, “What is a Dynamic Legend?,” on page 31](#)
- ♦ [Section 5.2, “Why Add Dynamic Legends?,” on page 31](#)
- ♦ [Section 5.3, “Sample Dynamic Legend,” on page 32](#)
- ♦ [Section 5.4, “Required XML Tags,” on page 32](#)
- ♦ [Section 5.5, “Adding a Dynamic Legend to a Knowledge Script,” on page 33](#)

5.1 What is a Dynamic Legend?

A dynamic legend is a block of XML text that accompanies each data stream collected by a Knowledge Script. The dynamic legend describes the data stream so that Analysis Center can accurately and uniquely identify the data to process it for inclusion in the Data Warehouse.

A dynamic legend includes:

- ♦ A static description, which appears as a member name in the second level of the Metric dimension hierarchy in the Data Warehouse.
- ♦ A dynamic description, which appears as a member name in the third level of the Metric dimension hierarchy in the Data Warehouse.

5.2 Why Add Dynamic Legends?

You want to add this capability to a Knowledge Script when:

- ♦ You have changed the legend format of an existing Knowledge Script.
- ♦ You have created a new Knowledge Script with a unique legend format.

Analysis Center can process dynamic legend data without having to rely on a legend-processing rule for each data stream collected by a Knowledge Script.

To see an example of how dynamic legends are used in an existing Knowledge Script, examine the code for the AppManager version 6.0 of NT_LogicalDiskSpace.

5.3 Sample Dynamic Legend

Below is a sample dynamic legend from the NT_LogicalDiskSpace Knowledge Script.

```
<?xml version="1.0" encoding="utf-8"?>

<APPMANAGER>

  <LEGENDDETAIL>

    <ORIGINKSNAME>NT_LogicalDiskSpace</ORIGINKSNAME>

    <APPLICATION>NT</APPLICATION>

    <ACTYPE>Logical Disk</ACTYPE>

    <DESCRIPTION>Logical Disk Space Available {MB}

  </DESCRIPTION>

  <OBJECTINFO>Logical Disk C:</OBJECTINFO>

    <OBJECTNAME>NT_LogicalDiskObj = C:</OBJECTNAME>

    <FULLPATHOBJECTNAME>

  </FULLPATHOBJECTNAME>

    <FULLPATHDELIM>

  </FULLPATHDELIM>

    <UNIT>MB</UNIT>

  </LEGENDDETAIL>

</APPMANAGER>
```

5.4 Required XML Tags

The following XML tags are required in a dynamic legend:

XML Tag	Description
<ORIGINKSNAME>	The full name of the Knowledge Script.
<APPLICATION>	A valid AppManager application name such as NT, IIS, or SQL. You can query the Component table in the AppManager repository to get a list of valid application names.
<DESCRIPTION>	The static description of the data stream. This should describe the data stream without using dynamic data such as server name, drive letter, database name, or variable information. For example, Logical Disk Space Available (MB) .

XML Tag	Description
<OBJECTINFO>	The dynamic description of the data stream. This should describe the data stream using dynamic data such as server name, drive letter, database name, or other variable information. For example, Logical Disk C:
<UNIT>	The unit applied to the data. For example: <ul style="list-style-type: none"> ♦ % ♦ # ♦ MB ♦ KB ♦ Msec ♦ Sec

5.5 Adding a Dynamic Legend to a Knowledge Script

To provide a dynamic legend for a data stream collected by a Knowledge Script, use the `CreateData` callback. The `CreateData` callback replaces the `DataHeader`, `DataLog`, and `DynaCreateData` callbacks in the following format:

```
CreateData streamId, legend, dynalegend, objlist, val, agentmsg, msgtype [,schema]
[,loglimit] [,lowWM] [,hiWM] [,deletefile]
```

6 Working with Legend Rules

Analysis Center relies on legend rules to process the AppManager data that define Knowledge Script data streams from older Knowledge Scripts that do not support the newer dynamic XML legends, or that are not properly formed dynamic XML. These rules configure Analysis Center to map each non-dynamic legend to a unique Metric dimension member in the Data Warehouse.

You can remove any legend rule that maps incorrectly or is no longer valid. You can also add a legend rule that you previously removed.

For more information about dynamic legends, see [Chapter 5, “Adding Dynamic Legends to AppManager Knowledge Scripts,”](#) on page 31.

- [Section 6.1, “Identifying a Legend Rule,”](#) on page 35
- [Section 6.2, “Removing a Legend Rule,”](#) on page 36
- [Section 6.3, “Creating the Removal XML File,”](#) on page 36
- [Section 6.4, “Issuing xmlcheckin,”](#) on page 37
- [Section 6.5, “Adding a Legend Rule,”](#) on page 38

6.1 Identifying a Legend Rule

You can easily identify any legend rule. The legend rules are contained in XML files sorted by AppManager module name in C:\Program Files\NetIQ\Analysis Center\Config\Legend on the Data Warehouse computer.

For example, in CIM.LEGEND.XML, you can see the mapping of the ASR Reboot Count data stream legend of the CIM_ASRHealth Knowledge Script.

```
<RULE ID="CIM_ASRHealth.1">
  <EXPRESSION>^ASR Reboot Count$</EXPRESSION>
  <ITEM IDX="0" LABEL="Number of Reboots" MATCH="TRUE" />
  <UNITS>#</UNITS>
  <KSName>CIM_ASRHealth</KSName>
  <ACApplication>CIM</ACApplication>
  <ACObject NAME="" />
  <ACMetricCategory NAME="" />
  <SampleHiddenLegend>ASR Reboot Count</SampleHiddenLegend>
  <StaticDataDescription>Automatic Server Recovery (ASR) Health</StaticDataDescription>
```

```

<DynamicDataFormat>Number of Reboots</DynamicDataFormat>

<Type NAME="CIMT_ASR" />

</RULE>

```

Knowledge Scripts that collect more than one data stream will have more than one rule. The additional rules for the data stream will have a unique number assigned to the rule. So in the above example, if there were a second data stream rule for the CIM_ASRHealth Knowledge Script, the "ID" for the rule would be CIM_ASRHealth.2.

6.2 Removing a Legend Rule

You remove a legend rule by defining it in XML, and thereafter, pass the rule to the `XmlCheckIn` utility. The first step in the legend rule removal process is to mark the rule for removal. When a rule is marked as removed it will not be included in future legend processing unless it is re-added. Removing a legend rule can have one of the two consequences, depending on the status of the rule.

- ♦ **Processed legend.** The rule for a legend that has already been mapped, is *marked* as deleted in the Data Warehouse database. In this scenario, the rule cannot be physically removed because it has been used, but because it is marked for deletion, it will never be used again.
- ♦ **Unprocessed legend.** The rule for a legend that has never been mapped, is physically removed from the list of rules stored in the Data Warehouse database.

To remove a legend rule:

- 1 Create an XML file that identifies the rule that you want to remove. For more information, see [Section 6.3, "Creating the Removal XML File," on page 36](#)
- 2 Run the `xmlcheckin` command. For more information, see [Section 6.4, "Issuing xmlcheckin," on page 37](#)

6.3 Creating the Removal XML File

Your XML file can identify one or more rules that you want to remove. If you are removing several rules that have the same prefix or suffix, you can use a percent(%) wildcard character.

To remove one or more rules without using a wildcard character, create an XML file that has the following format:

```

<?xml version="1.0" encoding="UTF-8" ?>

- <xml>

- <LEGENDPARSERS>

- <LEGENDPARSER>

  <DELETERULE ID="ruleID" />

</LEGENDPARSER>

</LEGENDPARSERS>

</xml>

```

where `ruleID` is the name of the rule taken from the rules XML file. For more information, see [Section 6.1, “Identifying a Legend Rule,” on page 35](#). Use the `<DELETERULE ID>` tag once for each rule that you want to delete.

To remove multiple rules using a wildcard character, create an XML file that has the following format:

```
<?xml version="1.0" encoding="UTF-8" ?>
- <xml>
- <LEGENDPARSERS>
- <LEGENDPARSER>
  <DELETERULE WildCard="prefixofrule.%" />
  </LEGENDPARSER>
</LEGENDPARSERS>
</xml>
```

where `prefixofrule` is the portion of the rule name common to all of the rules that you want to delete. Be sure that there are not any rules that you want to keep that have the same prefix. If there are, then you will want to provide a more specific prefix, or supply `<DELETERULE ID>` tags to list the specific rules that you want to delete.

You can use a combination of `<DELETERULE ID>` and `<DELETERULE WildCard>` tags in your file. For example:

```
<?xml version="1.0" encoding="UTF-8" ?>
- <xml>
- <LEGENDPARSERS>
- <LEGENDPARSER>
  <DELETERULE WildCard="NT_PhysicalDiskB%" />
  <DELETERULE ID="Commerce_AuthCheckSuccess.1" />
  </LEGENDPARSER>
</LEGENDPARSERS>
</xml>
```

Save your removal XML file as `<nameoffile>.xml`.

6.4 Issuing `xmlcheckin`

From a command prompt at the root of the Analysis Center\bin directory, enter the following command: `xmlcheckin <nameoffile>.xml`

If your `<nameoffile>.xml` is not located in the \bin directory, include the full path in the command, for example:

```
xmlcheckin C:\temp\<nameoffile>.xml
```

6.5 Adding a Legend Rule

You can add a legend rule that has been deleted or marked for deletion.

To add a legend rule:

- 1 Create an XML file that identifies the rule you want to add, in the following format:

```
<?xml version="1.0" encoding="UTF-8" ?>
- <xml>
- <LEGENDPARSERS>
- <LEGENDPARSER>
  <ADDRULE ID="ruleID" />
</LEGENDPARSER>
</LEGENDPARSERS>
</xml>
```

where `ruleid` is the name of the rule that you want to add.

- 2 Save the file as `<nameoffile>.xml`.
- 3 Issue the following command from a command prompt at the root of the `\bin` directory:
`xmlcheckin <nameoffile>.xml`.

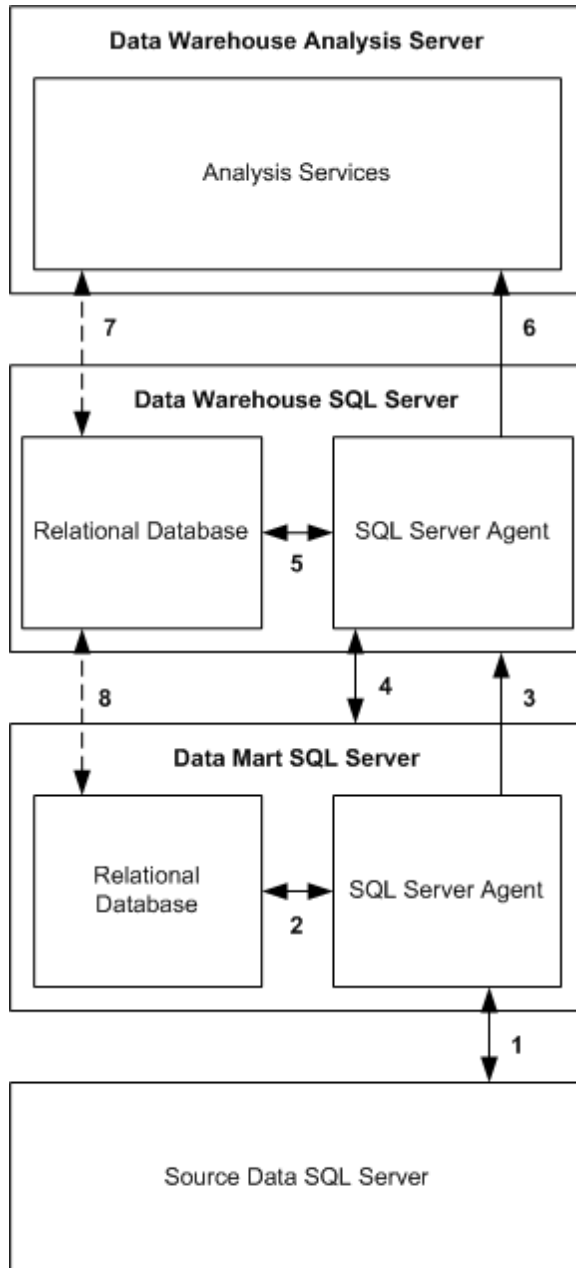
7 Understanding Data Connections and Flow

The various SQL Server and Analysis Services components make data connections. The SQL Server Agent on the Data Mart server reads from the source database, and writes to the Data Mart and Data Warehouse relational databases. The SQL Server agent owns the **ETL** job. The ETL job is the process by which data is extracted from an AppManager repository, transformed to a format appropriate to the Data Warehouse, and then loaded into the Data Warehouse.

The SQL Server Agent on the Data Warehouse server creates and reads from the Data Mart database, reads from and writes to the Data Warehouse relational database, and invokes the SSIS packages that process the cubes in the Data Warehouse multidimensional database (the OLAP processing job).

The Analysis Services on the multidimensional database server read data in the Data Warehouse relational database, as well as views in that database that are based on fact data in the Data Marts.

A legend describing the connections, and how Analysis Center gets its data follows the illustration.



Connection 1: Data Mart SQL Server Agent to source SQL Server. This connection is defined when you create a Data Source from the Analysis Center Console. This connection can use either Windows or SQL Server Authentication, depending on the security configuration of the source SQL Server and the choices you make during configuration of the Data Source.

This connection is used to copy data from the source database.

Connection 2: Data Mart SQL Server Agent to Data Mart SQL Server database. This connection is defined when you install the Analysis Center Data Warehouse. This connection can use either Windows or SQL Server Authentication, depending on the security configuration of the Data Mart SQL Server and the choices you make during the Data Warehouse installation.

This connection is used to write source data to the Data Mart database and to prepare that data for the Data Warehouse.

Connection 3: Data Mart SQL Server Agent to Data Warehouse SQL Server. This connection is defined when you install the Analysis Center Data Warehouse. This connection can use either Windows or SQL Server Authentication, depending on the security configuration of the Data Mart SQL Server and the choices you make during the Data Warehouse installation.

This connection is used to copy data from the Data Mart to the Data Warehouse relational database.

Connection 4: Data Warehouse SQL Server Agent to Data Mart SQL Server. This connection is defined when you install the Analysis Center Data Warehouse. This connection can use either Windows or SQL Server Authentication, depending on the security configuration of the Data Mart SQL Server and the choices you make during the Data Warehouse installation.

This connection is used to create the Data Mart databases.

Connection 5: Data Warehouse SQL Server Agent to Data Warehouse relational database. This connection is defined when you install the Analysis Center Data Warehouse. This connection can use either Windows or SQL Server Authentication, depending on the security configuration of the Data Warehouse SQL Server and the choices you make during the Data Warehouse installation.

This connection is used to further process data.

Connection 6: Data Warehouse SQL Server Agent to Analysis Services. This connection is defined when you install the Analysis Center Data Warehouse. This connection must use Windows authentication.

This connection is used to initiate cube and dimension processing.

Connection 7/8: Analysis Services to Data Warehouse SQL Server to Data Mart SQL Server. This connection is defined when you install the Analysis Center Data Warehouse. This connection can use either Windows or SQL Server Authentication, depending on the security configuration of the Data Warehouse and Data Mart SQL Servers and the choices you make during the Data Warehouse installation.

This connection is used to get data for cube and dimension processing from the Data Warehouse relational database and from the Data Mart database. The Data Warehouse relational database contains dimensional data and views to fact data in the Data Mart database; the Data Mart database contains fact data.

For more information about Analysis Center components, and how Analysis Center works, see the [NetIQ Analysis Center User Guide](#).

