IBM Cognos Virtual View Manager
Version 10.2.0

Administration Guide

IBM

Product Information

This document applies to IBM Cognos Business Intelligence Version 10.2.0 and may also apply to subsequent releases. To check for newer versions of this document, visit the IBM Cognos Information Centers (http://publib.boulder.ibm.com/infocenter/cogic/v1r0m0/index.jsp).

Contents

Introduction

This guide is designed for first-time users that are not administrators and are interested in addressing business issues presented by their disparate business systems.

The purpose of this guide is:
- To demonstrate how you can use IBM® Cognos® Virtual View Manager to address your business needs
- To introduce the data modeling aspect of Virtual View Manager

Audience

This documentation is for information technology professionals who want to use IBM Cognos Virtual View Manager to model data resources. Knowledge of relational data sources, hierarchical data sources, and data modeling is recommended.

Finding information

To find IBM Cognos product documentation on the web, including all translated documentation, access one of the IBM Cognos Information Centers (http://publib.boulder.ibm.com/infocenter/cogic/v1r0m0/index.jsp). Release Notes are published directly to Information Centers, and include links to the latest technotes and APARs.

You can also read PDF versions of the product release notes and installation guides directly from IBM Cognos product disks.

Accessibility features

This product does not currently support accessibility features that help users who have a physical disability, such as restricted mobility or limited vision, to use this product. IBM Cognos HTML documentation has accessibility features. PDF documents are supplemental and, as such, include no added accessibility features.

Forward-looking statements

This documentation describes the current functionality of the product. References to items that are not currently available may be included. No implication of any future availability should be inferred. Any such references are not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of features or functionality remain at the sole discretion of IBM.

Samples disclaimer

The Great Outdoors Company, GO Sales, any variation of the Great Outdoors name, and Planning Sample depict fictitious business operations with sample data used to develop sample applications for IBM and IBM customers. These fictitious records include sample data for sales transactions, product distribution, finance, and human resources. Any resemblance to actual names, addresses, contact

numbers, or transaction values is coincidental. Other sample files may contain fictional data manually or machine generated, factual data compiled from academic or public sources, or data used with permission of the copyright holder, for use as sample data to develop sample applications. Product names referenced may be the trademarks of their respective owners. Unauthorized duplication is prohibited.

Chapter 1. Post-installation tasks

This chapter describes some basic tasks that enable secure computing and client connections. These tasks must be performed sometime soon after installation of IBM Cognos Virtual View Manager before serious design and resources may be published.

The following topics are covered in this chapter:

- "Configuring IBM Cognos Virtual View Manager to display Unicode fonts"

  Update Virtual View Manager to support unicode fonts.

- "SSL management" on page 2

  Install a Java Key Store for use with Virtual View Manager.

  Note that this task does not need to be done immediately. However, it is strongly recommended that all server instances are configured with their own certificate prior to deployment with sensitive data.

- "Installing and using special JDBC drivers" on page 2

  Make sure that you have installed the appropriate JDBC drivers for the relational data sources that you plan to use.

  You can use Virtual View Manager's JDBC driver (and edit it if necessary) or install a specific JDBC driver for your installation.

- "Using the IBM Cognos Virtual View Manager ODBC driver" on page 9

  Configure the Virtual View Manager ODBC driver to connect with ODBC data sources and to receive connections from ODBC clients. If you installed the ODBC Client with Virtual View Manager then ODBC Client applications can use a 32-bit driver to connect with the server. The computer on which the ODBC client application resides must be configured to use the Virtual View Manager drivers to properly connect with the server.

- "Configuring IBM Cognos Virtual View Manager for use of a JMS Broker" on page 13

  By default Virtual View Manager supports both Sonic and TIBCO JMS brokers, however some non-distributable JARs must be copied to the installation directory from the JMS installation(s) that will be used.

Configuring IBM Cognos Virtual View Manager to display Unicode fonts

IBM Cognos Virtual View Manager Server returns Unicode characters in all messages carrying data. The server also transforms messages with other UTF encoding formats to Unicode.

To display all Unicode true type fonts, you must manually configure Cognos Virtual View Manager.

Procedure

1. On a Microsoft Windows operating system where Cognos Virtual View Manager is installed, create a folder named **fallback** in the *installation_location*\jre\lib\fonts\ directory.

   For example, *installation_location*\jre\lib\fonts\fallback

2. From `C:/Windows/fonts` or `C:/WINNT/fonts` directory, copy the Arial Unicode MS (TrueType) font file (ARIALUNI.TTF) to the *installation_location*\jre\lib\fonts\fallback directory you just created.

3. Restart Cognos Virtual View Manager.

SSL management

IBM Cognos Virtual View Manager enables specification of the Java Key Store (JKS) used to initiate and establish SSL communications over both HTTPS ports used for secured web services and secured JDBC communications.

Viewing the **SSL Management** page requires a user profile that has the Read All Resources right and change of any of the JKS digital certificate file location, type, or password requires the Modify All Resources right.

A generic JKS file is provided so that development and testing of Web services and JDBC secured over HTTPS ports may proceed without need for immediate installation of a JKS file.

You should configure all Virtual View Manager instances with their own JKS certificate prior to deployment with sensitive data.

Obtain your JKS digital certificate(s) from a Certificate Authority (CA) or generate your own and install it using on the **SSL Management** page in Virtual View Manager Administrator.

Procedure

1. In Virtual View Manager Administrator, click **Configuration** > **SSL**.
2. The **SSL Management** page appears.
3. In the **New Value** column, enter the absolute path to the new JKS file (X.509 compliant certificate file) on the server and click **Apply**.
4. Change the **Java Keystore File Type** and the **Java Keystore Password** values in the same way described above so that the values on server restart match the digital certificate being installed.
5. Restart the server to apply the changes.

Installing and using special JDBC drivers

Some relational data sources require additional JDBC drivers to enable connection, introspection, and use. These data source drivers must be installed separately from the IBM Cognos Virtual View Manager.

The following relational data sources require additional JDBC drivers:
- IBM DB2® (type 2 and type 4)
- IBM DB2 (Mainframe)
- IBM Informix®
- Microsoft SQL Server
- Netezza®
- Teradata
- MySQL
- Sybase
- Neoview

- Oracle (type 2 and type 4)

You must install the necessary driver(s) in the appropriate location(s) so that the Virtual View Manager server and the data source can interact. Each particular data source has a directory within the Virtual View Manager installation directory, which is specified below.

Virtual View Manager provides a JDBC interface and provides methods to connect to relational data sources that may not be formally supported. Custom jars may be written to direct the server to connect using the custom jar. Simply specify the JDBC driver and direct the server to upload it to the system.

One driver is sufficient to connect to any number of the same type of data sources. Once uploaded, the JDBC driver will function like any other JDBC driver, such as Oracle, SQL Server, or MySQL.

Virtual View Manager server assumes that JDBC drivers conform to the JDBC 2.0 standard. The server does not make any accommodations for JDBC drivers that don't supply correct metadata about the data source. The server does not retrieve result sets that are not consistent with the metadata supplied.

Installing pre-configured JDBC drivers

This section describes how to install pre-configured JDBC drivers from specific locations for connecting to specific data sources.

The IBM Cognos Virtual View Manager location to copy the JAR files into is:

installation_directory\apps\dlm\cis_ds_<datasource_type>\lib

For example, the location for the IBM Informix data source is c:\vvm\apps\dlm\ cis_ds_informix\lib.

Installing pre-configured drivers for DB2 (Type 2 or Type 4)

This section describes how to install pre-configured JDBC drivers for DB2 (Type 2 or Type 4).

Procedure

1. Obtain the appropriate driver file for your version of DB2.
   - Use db2jcc.jar and the accompanying license file, for example, db2jcc_license_cisuz.jar or db2jcc_license_cu.jar
2. Copy the jar file, and license file, to the following Virtual View Manager directory:
   installation_directory\apps\dlm\cis_ds_db2\lib
3. Restart the server.

Installing pre-configured drivers for DB2 z/OS

This section describes how to install pre-configured JDBC drivers for DB2 z/OS®.

Procedure

1. The driver for DB2 z/OS is the same as the driver for DB2 on UNIX. Obtain db2jcc.jar, its accompanying license file, for example, db2jcc_license_cisuz.jar or db2jcc_license_cu.jar, and common.jar.
2. Copy the files to the Virtual View Manager installation directory:
   installation_directory\apps\dlm\cis_ds_db2_mainframe\lib

3. Restart the server.

Installing pre-configured drivers for IBM Informix

This section describes how to install pre-configured JDBC drivers for IBM Informix.

Procedure

1. Obtain the IBM Informix JDBC driver for your version of Infomix.
2. Copy the ifxjdbc.jar driver file to the Virtual View Manager installation directory:

   installation_directory\apps\dlm\cis_ds_informix\lib
3. Restart the server.

Installing pre-configured drivers for Microsoft SQL Server

This section describes how to install pre-configured JDBC drivers for Microsoft SQL Server.

Procedure

1. Obtain the driver file for your version of Microsoft SQL Server.
2. Download the driver version for the appropriate platform.
3. Run the setup.exe installation program, and do the following:
   - For Microsoft SQL 2005, extract the sqljdbc.jar file.
   - For Microsoft SQL 2008, extract the sqljdbc4.jar file.
4. Copy the JAR file(s) to the Virtual View Manager directory:

   `installation_directory\apps\dlm\cis_ds_mssql\lib`
5. Restart the server.

Installing pre-configured drivers for Netezza

This section describes how to install pre-configured JDBC drivers for Netezza.

Procedure

1. Obtain the JDBC driver for Netezza from the CDs that you received with the NPS® system or by contacting the support group at Netezza.
2. Copy the nzjdbc.jar driver file to the Virtual View Manager installation directory:

   installation_directory\apps\dlm\cis_ds_netezza\lib

Installing pre-configured drivers for Teradata

This section describes how to install pre-configured JDBC drivers for Teradata.

Procedure

1. Obtain the JDBC driver for your version of Teradata.
2. Copy the driver files named tdgssconfig.jar, tdgssjava.jar, and terajdbc4.jar to the Virtual View Manager installation directory:

   installation_directory\apps\dlm\cis_ds_teradata\lib
3. Restart the server.

Installing pre-configured drivers for MySQL

This section describes how to install pre-configured JDBC drivers for MySQL.

Procedure

1. Obtain the JDBC driver for your version of MySQL.

2. Copy the mysql-connector-java-3_1_10_1-bin.jar driver file to the Virtual View Manager installation directory:

    installation_directory\apps\dlm\cis_ds_mysql\lib

3. Restart the server.

Installing pre-configured drivers for Sybase

This section describes how to install pre-configured JDBC drivers for Sybase.

Procedure

1. Obtain the JDBC driver for your version of Sybase.
2. Copy the jconn3.jar and jTDS3.jar driver files to the Virtual View Manager installation directory:

    installation_directory\apps\dlm\cis_ds_sybase\lib

3. Restart the server.

Installing pre-configured drivers for Neoview

This section describes how to install pre-configured JDBC drivers for Neoview.

Procedure

1. Obtain the JDBC driver for your version of Neoview.
2. Copy the driver files to the Virtual View Manager installation directory:

    installation_directory\apps\dlm\cis_ds_neoview\lib

3. Restart the server.

Installing pre-configured drivers for Oracle (Type 2 or Type 4)

This section describes how to install pre-configured JDBC drivers for Oracle (Type 2 or Type 4).

Procedure

1. Obtain the JDBC driver for your version of Oracle.
2. Copy the ojdbcXX.jar, xdb.jar, xmlparservX.jar driver files to the Virtual View Manager installation directory:
    *   *For type 4 - installation_directory*\apps\dlm\cis_ds_oracle\lib
    *   *For type 2 - installation_directory*\apps\dlm\cis_ds_oracle_type2\lib
3. Restart the server.

Connecting through JDBC drivers

IBM Cognos Virtual View Manager connects to and introspects underlying data sources so that a virtual, integrated data layer may be created, selectively published, and queried as a single data source.

JDBC connection service requests are created and maintained between the end-user client applications and the Virtual View Manager server and from the server to the underlying data sources.

This administrative section describes the generic JDBC connection URLs between the Virtual View Manager server and the native data sources. The information could also be generally applied to the client-server JDBC connection, though the end-user JDBC client typically connects only to the Virtual View Manager server data source and not directly to the underlying data sources. There are exceptions such as test cases to specify additional attributes for complex transaction handling or to verify proper driver functionality.

Defining new data sources

When the required data source driver is installed and ready for use, developers may create connections with the desired data sources.

You define new data sources in the user home directory or in any appropriate directory. The JDBC driver connection URL is built dynamically using a wizard. The **Add Physical Data Source** wizard enables selection of the data source driver to build the connection URL according to the URL format specifications of that driver class. Custom configurations, URL attribute additions, and implementation specific attributes may be set using the wizard as well.

Each data source driver supports URL service requests with a specific binding format. The general JDBC connection URL format is:

jdbc:<jdbc-subprotocol>:[implementation specific URL attributes]

where <jdbc-sub-protocol> identifies the JDBC implementation and typically identifies the JDBC driver vendor. The implementation specific URL attributes vary widely depending on the vendor data source and implementation specific configurations. A table listing the different data source driver JDBC URL formats and addressable driver class names is provided in "Sample JDBC driver connection URL formats" on page 7.

JDBC driver connection URL format

Knowing the connection URL format and driver class name enables a direct connection to the underlying data source utilizing the drivers supported by the IBM Cognos Virtual View Manager server. JDBC clients may connect directly with new data sources mediated by the JDBC port.

The driver connection URL format is a template for data source definitions. It contains three literal strings:

<HOST>, <PORT>, and <DATABASE_NAME>

When you add a data source to Virtual View Manager using this driver, the server substitutes the literals <HOST>, <PORT>, and <DATABASE_NAME> with the values you supply for the host name, port number, and database name respectively in the **Add Physical Data Source** window.

The complete syntax of the URL:

jdbc:cognos:dbapi@<HOST>:<PORT>?domain=<domain>&
dataSource=<DATASOURCE>[&NAME=VALUE]*

where
- HOST—valid host name or IP address
- PORT—integer setting on which Virtual View Manager and the host database will communicate
- DOMAIN—user domain (Cognos is the default domain)
- DATASOURCE—data source name

and zero or more optional NAME=<VALUE> pairs may be specified depending on the target data source and driver. Refer to the next table to see the generic URL format used by the Virtual View Manager server to connect to the data source.

Sample JDBC driver connection URL formats

Knowing the connection URL format and driver class name enables a direct connection to the underlying data source utilizing the drivers supported by the IBM Cognos Virtual View Manager server. This section provides examples.

This table lists JDBC driver example URL formats and the corresponding driver class names for supported data sources.

Database / data source	URL format / driver name	Value
Virtual View Manager	URL format:	jdbc:cognos:dbapi@<HOST>:<PORT>?domain= <DOMAIN>&dataSource=<DATABASE_NAME>Normally this is the only JDBC client connection used to establish client-server connections to Virtual View Manager, where cognos is the default <DOMAIN> value unless an LDAP or dynamic domain user is negotiating the connection.
	Driver name:	cs.jdbc.driver.CompositeDriver
DataDirect Mainframe	URL format:	jdbc:neon:; APNA=<APP_NAME>; CPFX=<CATALOG_PREFIX>; DBTY=<DBMS_TYPE>; HOST=<HOST>; PORT=<PORT>; SUBSYS=<DATABASE_NAME>; TRLT=NO; DTFM=ODBC;
	Driver name:	com.neon.jdbc.Driver
	Driver name:	COM.ibm.db2.jdbc.net.DB2Driver
DB2 v8 (type_2)	URL format:	jdbc:db2:<DATABASE_NAME>
	Driver name:	Driver Name:
DB2 v8 (type_4)	URL format:	jdbc:db2://<HOST>:<PORT>/<DATABASE_NAME>
	Driver name:	com.ibm.db2.jcc.DB2Driver
DB2 z/os v8 (type_4) mainframe	URL format:	jdbc:db2://<HOST>:<PORT>/<DATABASE_NAME>
	Driver name:	com.ibm.db2.jcc.DB2Driver
IBM Informix	URL format:	jdbc:informix-sqli://<HOST>:<PORT>/<DATABASE_NAME>: informixserver=<IBM Informix instance name>;user=<user_name>; password=<password> The informixserver variable is the name of the IBM Informix instance, and not the server host name, and <user_name> and <password> are recognized by the IBM Informix server.
	Driver name:	com.informix.jdbc.IfxDriver
MySQL	URL format:	jdbc:mysql://<HOST>:<PORT>/<DATABASE_NAME>

Database / data source	URL format / driver name	Value
	Driver name:	
Netezza	URL format:	jdbc:netezza://<HOST>:<PORT>/<DATABASE_NAME>
	Driver name:	org.netezza.Driver
Oracle (thin), 9i, and 10g	URL format:	jdbc:oracle:thin@<HOST>:<PORT>:<DATABASE_NAME>
	Driver name:	oracle.jdbc.driver.OracleDriver
Oracle type2 (OCI) 9i, and 10g	URL format:	jdbc:oracle:oci:@<DATABASE_NAME>
	Driver name:	oracle.jdbc.OracleDriver
Microsoft SQL Server	URL format:	jdbc:microsoft:sqlserver://<HOST>:<PORT>;databaseName=<DATABASE_NAME>;SelectMethod=<SELECT_MODE>
	Driver name:	com.microsoft.jdbc.sqlserver.SQLServerDriver
Sybase	URL format:	jdbc:sybase:Tds:<HOST>:<PORT>/<DATABASE_NAME>
	Driver name:	com.sybase.jdbc3.jdbc.SybDriver
Teradata	URL format:	jdbc:teradata://<HOST>/DBS_PORT=<PORT>/DATABASE=<DATABASE_NAME>/CHARSET=UTF8,COMPAT_DBS=true
	Driver name:	com.ncr.teradata.TeraDriver

Adding, editing, and removing JDBC drivers

You can install your own JDBC driver on the server. You can also edit or remove the drivers you create, referred to as user-installed drivers.

Adding JDBC drivers

You can install your own JDBC driver on the server.

Procedure

1. The data source wizard lets you add your own JDBC driver to the Virtual View Manager metadata environment. Right-click at an appropriate location in the resource tree, and click **New Data Source**.

2. In the **Add Physical Data Source** window, click **New Driver**.

3. In the **New Driver Information** window, enter the driver information in the fields as follows:

   - **Name**: User-defined name for the driver. Type the name. This is a user-defined name for the driver. After adding it to the server, this name will be listed along with other data source drivers listed on the first window of the **Add Physical Data Source** wizard.

- **Type**: Type of the JDBC driver. Accept **Jdbc**.
- **Jar File**: Click the **Browse** button to locate the JAR file(s) where the driver class is stored. You can also type the name (with path) in the **Jar File** field. You can only upload a JDBC driver JAR file that is visible to Virtual View Manager. For example: *installation_directory*\apps\jdbc\lib\<jar_file>
- **Driver Class Name**: Fully-qualified name of the driver class. For example: oracle.jdbc.driver.OracleDriver
- **Connection URL Pattern**: Format of the URL to connect to the database. For details on the connection URL format, see "Connecting through JDBC drivers" on page 5.

  JDBC drivers are installed in their respective directory, as described in the beginning of the section "Installing and using special JDBC drivers" on page 2. For example, DB2 driver JAR files (db2jcc.jar and the accompanying license file) are installed in the following directory:

  installation_directory\apps\dlm\db2\lib\

4. Click **OK**.

   The driver is displayed with other data source drivers in the first window of the **Add Physical Data Source** window. You can use this driver to add JDBC type data sources to the server.

Editing JDBC drivers

You can edit user-installed JDBC drivers on the server.

Procedure

1. Open the data source wizard.
2. In the **Select Data Source Driver** section, select the user-installed driver that you want to edit.
3. Click **Edit Driver**.
4. In the editing window, make the necessary changes, and click **OK**.

   Use the **Browse** button to locate the JAR file(s) if the location of the file(s) has changed. For example, DB2 driver JAR files are installed in the following directory:

   installation_directory\apps\dlm\ds_db2\lib\

Removing JDBC drivers

You can remove user-installed JDBC drivers on the server.

Procedure

1. Open the data source wizard.
2. In the **Select Data Source Driver** section, select the driver which you want to remove.
3. Click **Delete Driver**.

Using the IBM Cognos Virtual View Manager ODBC driver

If the option to install the ODBC driver is specified during the installation process, the driver is installed.

Using the ODBC driver on Windows operating systems

This section provides the information you need to use the ODBC driver on Windows.

This section covers the following topics:
- "Adding ODBC data sources on Windows operating systems."
- "Overriding the configured settings."
- Code sample for connecting to Cognos Virtual View Manager Server.
- "Supported and non-supported features" on page 11.

Adding ODBC data sources on Windows operating systems

The native Windows driver managers are supported by IBM Cognos Virtual View Manager.

Procedure

1. From the Windows Control Panel, open **Administrative Tools** > **Data Sources (ODBC)**.
2. Click the **User DSN** tab or the **System DSN** tab.

   A **User DSN** is accessible only to the current user. A **System DSN** is accessible to all the users on the system and requires special permission to create and modify.
3. Click the **Add** button.
4. In the **Create New Data Source** screen, select the Virtual View Manager driver, and then click **Finish**.
5. In the **Driver Configuration** window, enter the following information that is required for configuring the driver:
   - **DSN Name**: Name of the data source to which the clients will refer. Once a DSN is created, its name cannot be changed.
   - **Host**: Server name (or IP address) on which Virtual View Manager is running.
   - Port: TCP port used to communicate with Virtual View Manager Server, which must match the port that the server is listening on. With default installation settings, the server listens to port 9401. To confirm the port number, click **Virtual View Manager Server** > **JDBC and ODBC Drivers** > **Communications** > **Port** in Virtual View Manager.
   - **User Name, Password, and Domain**: A Virtual View Manager user name and password. The **Password** is nullable.
   - **Datasource**: Name of the Virtual View Manager data source that the ODBC connection will access This entry sets the default scope of client queries to a particular datasource. Note that querying outside the scope of this data source requires super-qualified tables or stored procedures.
   - **Catalog**: Connects with default data-source catalog
6. Use the **Refresh** button to retrieve the Catalogs available to this user on the server.
7. Use the **Test** button to test the settings in the configuration dialog box.
8. Click **OK**.

   The configured settings you entered are saved, and the data source is added to your computer.

Overriding the configured settings

A client connecting to the server through the ODBC driver can override the configured settings on a data source by adding the appropriate parameters in the connection string.

For example, clients can use a connection string such as

DSN=<value>;UID=<value>;PWD=<value>;DOMAIN=<value>;
HOST=<value>;PORT=<value>;DATASOURCE=<value>; CATALOG=<value>;

Note that for these parameters:
- The original **DSN** value cannot be overridden.
- Upon creation of a **DSN**, you are prompted for a **PWD** entry.
- **HOST** is the host name where the server is running.
- **CATALOG** is optional.

Supported and non-supported features

This section describes the supported and non-supported features of the IBM
Cognos Virtual View Manager ODBC driver.

The IBM Cognos Virtual View Manager ODBC driver supports the following
features:
- Data types: CHAR, VARCHAR, SHORT, LONG, DOUBLE, FLOAT, TIME, DATE,
  TIMESTAMP

  Conversions (indicated by "->" below):

  CHAR, VARCHAR -> VARCHAR BIT, TINYINT, SMALLINT -> SMALLINT
  BIGINT, INT -> INT DECIMAL, REAL, FLOAT, NUMERIC -> FLOAT

  All other types are converted to VARCHAR.

The ODBC driver does not support the following feature:
- Parameters in prepared statements.

Known issue
Fixed width: All CHAR and VARCHAR data types are reported by the driver to be
256 characters. However, the driver supports the retrieval of longer values if the
client provides adequate memory for doing so.

Using the ODBC driver on UNIX and Linux operating systems

You should be familiar with some details if you plan to use the ODBC driver on a
UNIX or Linux operating system.

These details include:
- To install the ODBC driver at anytime after server installation you can log into
  the installation machine as the same user that installed the IBM Cognos Virtual
  View Manager server and the driver.
- To configure the ODBC driver, ensure that you have Read and Write permissions
  on the following files, which are in the C:\Windows directory:

  vvm<version>.xml

  odbc.ini

  odbcinst.ini

  The configuration is done by an interactive utility, driverConfig, which is in the
  <installation_directory>/apps/odbc/<platform>/bin.
- Creating a DSN is done through the configuration utility driverConfig, which
  helps users to reconfigure the driver files (in case, file-location is changed after
  installation), and create, edit, list, or delete DSN entries.

The rest of this section describes the following tasks for using the ODBC driver on UNIX:

- "Setting environment variables"
- "Creating a DSN with driverConfig" on page 13

Setting environment variables

You must set environment variables if you plan to use the ODBC driver on UNIX and Linux.

Procedure

1. Log into the installation machine as the same user that installed Virtual View Manager.
2. Set the following environment variables:
   - VVM_HOME

     The location where the driver is installed. This is the full path to the top-level installation directory for Virtual View Manager server.
   - VVM_DSN_XML

     This is an optional variable. It allows you to specify an alternate location for the DSN (Data Source Name) configuration file, so different users on the same computer can have different sets of DSNs configured. It is the full path to the ODBC DSN configuration file for the Virtual View Manager server. The DSN defaults to the following value, but it can be configured using the driverConfig utility:

     $VVM_HOME/vvm<version>.xml
   - ODBCINI

     odbc.ini defines the DSN entries. This value is the full path to the odbc.ini file. It is generated by the DSN configuration with driverConfig:

     <installation_directory>/odbc.ini
   - ODBCINSTINI

     odbcinst.ini defines the ODBC drivers. This value is the full path to the configuration file odbcinst.ini, which is generated during DSN configuration with driverConfig:

     <installation_directory>/odbcinst.ini
   - LD_LIBRARY_PATH

     This is specific to Solaris-based machines and Linux-based machines. This path refers to the location of the iODBC driver manager files. The default location is:

     <installation_directory>/apps/odbc/lib
   - LIBPATH

     This is specific to AIX-based computers. This path refers to the location of the iODBC driver manager files. The default location is:

     <installation_directory>/apps/odbc/lib
   - SHLIB_PATH

     This is specific to HP-UX-based computers. This path refers to the location of the iODBC driver manager files. The default location is:

     <installation_directory>/apps/odbc/lib

Creating a DSN with driverConfig

This section describes how to create a DSN by running the driverConfig utility program on a UNIX platform.

driverConfig is located in the $VVM_HOME/apps/odbc/<platform>/bin directory.

Procedure

Run driverConfig using the following command:
driverConfig

Sample interaction

Here is a sample interaction between the system and user.

```
Main Menu
0  Exit this utility
1  Configure ODBC administrator
2  View configuration and DSNs on this system
3  Create/edit a DSN
Enter command> 3
---------
Create/edit a DSN
0  Return to main menu
1  Create a DSN
2  Edit an existing DSN
3  Delete an existing DSN
Enter command> 1
Enter DSN name
> testdsn
Enter host [localhost]
>
Enter port [9401]        (This
is the default port setting)
>
Enter user
> admin
Enter password
> admin
Enter domain
> cognos
Enter datasource
> ds
Enter catalog
> cat
Keep this information?
[testdsn]
host = localhost
port = 9401
uid = admin
password = admin
domain = cognos
datasource = ds
catalog = cat
Enter (y)es or (n)o > y
```

The details for the newly created DSN are saved.

Configuring IBM Cognos Virtual View Manager for use of a JMS Broker

JMS provides a way to publish asynchronous message-based web services.

IBM Cognos Virtual View Manager supports both Sonic and TIBCO JMS brokers, but a few drivers must be copied from the JMS broker installation to the Virtual View Manager installation directory in order to properly connect the two servers.

To enable communications between Virtual View Manager and the JMS broker, several compiled library files, jars, will have to be copied to the server installation directory and the server will have to be restarted.

Enable connection with TIBCO JMS

The following procedure describes how to enable connection with TIBCO JMS.

Procedure

1. Find and copy the file named tibjms.jar from the TIBCO installation to the following Virtual View Manager directory:

   <installation_location>/apps/server/lib
2. Restart the Virtual View Manager server.

Enable connection with Sonic JMS

The following procedure describes how to enable connection with Sonic JMS.

Procedure

1. Find and copy the file named mfcontext.jar and the all jar files that begin with the prefix sonic_ to the following Virtual View Manager directory:

   <installation_location>/apps/server/lib
2. Restart the Virtual View Manager server.

Create and configure a queue and queue connection factory in JMS Broker

The following procedure describes how to create and configure a queue and queue connection factory in JMS Broker.

Procedure

1. For either JMS broker, if you haven't already, configure your JMS broker according to the respective manufacturer's instructions.

   Create a suitable Queue Connection Factory (QCF).
2. Create a suitable Queue.
3. Register the Queue Connection Factory and the Queue with the JNDI.

   Note: Connection to JMS via JNDI is currently supported by Virtual View Manager. Virtual View Manager supports queues, but not JMS topic connection factories.

Add JMS connectors to IBM Cognos Virtual View Manager

Connectors must be configured for use by IBM Cognos Virtual View Manager. Use Virtual View Manager Administrator to create connectors so that the server can publish JMS data services.

For more information on publishing Virtual View Manager Data Services to JMS queues, refer to the "Publishing Resources" of the IBM Cognos Virtual View Manager *User Guide*.

Procedure

1. Start Virtual View Manager Administrator:

2. Click **Configuration** > **Connectors** to open the **Connector Management** page.

3. Add a Connector using the provided button.

   The **Add a JMS via JNDI Connector** window is displayed.

   Almost all fields in the **Add a JMS via JNDI Connector** window must have some value for the Connector to function properly. Many of the fields do have defaults.

4. Enter values in the fields displayed by the **Info** tab.

   - **Connector Name**—convenient, identifier label for the connector

   - **Group Name**—connectors that share an identical group name share a common connection pool. Connector grouping has failover connection pooling because the connection pool is shared. If a connector instance fails other connectors in the group will be able to send and receive messages using the same connection pool.

   - **Annotation**—(optional) adds notes on the JNDI connector. Annotations will be visible on the **Connector Management** page.

5. Enter values in the fields displayed by the **JMS via JNDI** tab.

   - **Initial Context Factory**—Sonic and TIBCO context factories are supported by default, but other context factories may be supported. Other JMS compliant brokers and their respective JNDI context factories may be used to connect with Virtual View Manager. See the documentation for your particular JMS broker for specifics on what is needed to support a JNDI initial context factory. The JNDI initial context factory is usually the class name.

     Type "c" to see the default, suggested string values that may be used with Sonic JMS and TIBCO JMS as values for the initial context factories.

     For Sonic: com.sonicsw.jndi.mfcontext.MFContextFactory

     For TIBCO: com.tibco.tibjms.naming.TibjmsInitialContextFactory

   - **JNDI Provider URL**—URL for connection with the JNDI. TCP protocol is generally used. TIBCO default port is 7222, and Sonic default port is 2506. Make sure the appropriate port in the firewall is opened to allow connections with the JNDI provider.

   - **JNDI User** and **JNDI Password**—JMS JNDI user profile must have sufficient permissions to look up JMS destinations. Passwords are not stored in clear text.

   - **JMS Client ID**—(optional) name the Virtual View Manager connections with the JMS broker.

6. Enter values for additional name-value pairs on the **JNDI Properties** tab.

   The **JNDI Properties** tab enables specification of additional name value pairs.

   Click the plus button to add name value pairs.

   In particular, Sonic requires specification of a domain name, while TIBCO does not require further specification.

7. Enter values for additional name-value pairs on the **Pool** tab.

   The **Pool** tab enables specification of the connection thread timeout and pool size parameters. The default values will be generally adequate for development needs.

   - **Pool Timeout** specifies the maximum waiting time (in seconds) for a new connection. If a connection is not provided within the pool timeout period specified, then a check is made for an available connection using the proper

user and uses that, or if that is not available then the least recently used connection for some other user is dropped and a new connection for the required user is opened.

- **Minimum Pool Size** specifies the number of connections that should remain in the connection pool even when the pool becomes inactive.

  The connection pool is initially empty. When there is a need to connect to JMS via JNDI, the pool creates one connection based on the information provided in the Info panel. Connections remain available even when there is no activity because the time to negotiate new connections can add a significant amount to the query response time.

  If the JMS connection pool has been inactive for a while, the connector pool size will start to shrink based upon the connection inactivity. The **Minimum Pool Size** specifies the minimum number of connections to remain in the pool, to maintain a minimum number of available connections.

- **Maximum Pool Size** specifies the number of connections (both active and idle) allowed to access the data source. When the connection pool limit is reached new incoming requests must wait until the next connection is made available.

  The connection pool is initially empty. When there is a need to connect to the data source, the pool creates a connection based on the information provided in the **Info** panel. As making new connections can take some time and resources, these connections remain available even if they become idle.

  Connectors with identical group names will share the same pool of connections.

Configuring LDAP

You can use LDAP for domain administration. You can also use LDAP as a datasource.

Procedure

1. Obtain the ldapbp.jar file.
2. Copy the ldapbp.jar file to the Virtual View Manager installation directory:
   installation_directory\apps\common\lib
3. Restart the server.

Chapter 2. Virtual View Manager configuration

IBM Cognos Virtual View Manager provides a large number of configuration parameters that display configuration settings, usage data, and allow administrator modification of settings and behavior. Usually none of the configuration settings require modification for development environments, but certain settings enable optimal implementation for various production settings.

This chapter introduces some of the configuration tasks you can perform to track information and control Virtual View Manager behavior.

The following topics are covered in this chapter:
- "The Configuration window"
- "Fine tuning memory" on page 19
- "Case sensitivity" on page 19
- "Case sensitivity and trailing spaces mismatches" on page 19
- "Dealing with settings mismatches" on page 21
- "Trailing spaces" on page 22

Configuration rights and privileges

All IBM Cognos Virtual View Manager users (users with the Access Tools right) may display the Configuration panel from the Studio > Administration menu.

Display of all parameters in the Configuration panel requires both the Access Tools right and the Read All Config right. If a different set of rights are held by the user, then only an appropriate subset of the available parameters will be visible as read-only settings.

Modification of configuration parameters requires the Modify All Config right in addition to the Access Tools right.

The Configuration window

The Configuration window provides access to all IBM Cognos Virtual View Manager configuration parameters. You access the Configuration window by selecting Administration > Configuration.

The configuration parameters are grouped into three main categories: Virtual View Manager Server, Data Sources, and Studio. When you open the subfolders in each category and select a parameter indicated by the configuration parameter icon, studio displays the current configuration setting for that parameter along with a description in the right pane.

All parameters, whether read-only or configurable, have a type and a description that make their function and purpose clear. Read-only parameter values are grayed out and labeled (Current). If the value may be changed with the associated parameter is appropriately labeled (On Server Restart). Those parameters that require the server to be restarted are clearly marked in the parameter description.

Virtual View Manager is optimized for a typical development environment where an individual or a moderately sized team work together to prepare prototypes for test and later for a production deployment. Some configuration settings should be assessed prior to test and production deployment so that performance for a particular implementation environment is optimal.

Configuration parameters - server

Several configuration parameters are available. Configuration parameters are grouped by type.

The configuration parameters, grouped by type, are listed in the following table.

Folder	Configuration Parameters
API > Protocol	Current and supported API protocol versions
Communications	DN, Java Keystore file and settings
Configuration	Repository, Debug, E-mail, Files, Hooks, General Info (host name, IP, version,...), License, Metadata (change log, cache sizes, purge thresholds), Monitor settings, Network (HTTPS hostname verification, FTP, HTTP, and HTTPS proxy settings), Security (Anonymous and Dynamic logins), and transaction logging
Events and Logging	Storage conditions, Event Generation (Cache, Data Source, Request, Resource, Session, Storage, System Overview, Transaction and Trigger events), and Logging (custom logger, database logger, file logger, memory, and SNMP settings)
JDBC and ODBC Drivers	Client communications settings, data fetch default, requests and session time-outs
Memory	Java heap (current and total available memory), managed memory settings, maximum memory for a request
Runtime Processing Information	I/O samples, Repository (Privilege, Resource, and User cache usage), Requests tracking, Sessions, Storage, Transactions, Triggers, and Wait Queue settings
SQL Engine	Cardinality, Query Plan, Prepared Statement, and Runtime Stats caches, Case Sensitivity, Ignore Trailing Spaces, Logging for resource usage, and query statistics, SQL optimization settings, and overrides

All parameters, whether read-only or configurable, have a type and a description that make their function and purpose clear. Read-only parameter values are grayed out and labeled (Current). If the value may be changed with the associated parameter is appropriately labeled (On Server Restart). Those parameters that require the server to be restarted are clearly marked in the parameter description.

Virtual View Manager is optimized for a typical development environment where an individual or a moderately sized team work together to prepare prototypes for test and later for a production deployment. Some configuration settings should be assessed prior to test and production deployment so that performance for a particular implementation environment is optimal.

Fine tuning memory

If you want to change the default memory setting, (512 MB), use the Configuration window (Administration > Configuration menu option), and navigate to Memory (Configuration > Virtual View Manager Components > Virtual View Manager Server > Membory). You can modify the settings for Java heap.

Consider the following to determine the configuration for optimal memory:

- Queries run faster with more memory. So, giving the server as much memory as possible is highly desirable.
- However, giving the server too much memory can cause excessive paging (see "Paging"), which can degrade performance significantly.

Paging

Paging occurs if the total amount of memory of all the running applications exceeds the amount of physical memory.

In this situation, the operating system temporarily moves parts of the running applications onto the disk so that the applications won't crash when memory is exhausted. When a paged-out memory location is accessed, the operating system will restore that area of memory from disk and then, to make room, move some other part of memory to disk. Consequently, what should be a simple memory access becomes two disk operations and performance suffers. Some amount of paging is fine on the client side. But on the server, where things are a lot more controlled and optimal performance is desired, paging must be minimized.

Case sensitivity

By default, IBM Cognos Virtual View Manager is set to be not case sensitive. While the SQL specification encourages the use of case-sensitive string comparison, many databases default to a non-case sensitive comparison. Getting the correct results from a query requires knowledge of which type of comparison is used. For example, the test ('abc' = 'ABC') returns FALSE for a case sensitive comparison and TRUE for a non-case sensitive comparison.

Changing the case sensitivity setting might impact existing queries, in that the results they return may be changed or performance may be affected.

Case sensitivity and trailing spaces mismatches

Case sensitivity and trailing space mismatches are often encountered in enterprise environments with many different database systems.

With IBM Cognos Virtual View Manager, case sensitivity and trailing spaces mismatches only occur under the following conditions:

- There is a mismatch between Virtual View Manager and the underlying data source's case sensitivity and/or trailing spaces settings.
- There is a WHERE clause with a CHAR or VARCHAR in the clause.

Virtual View Manager handles case sensitivity and trailing space mismatches by following the conventions defined under **Administration** > **Configuration**.

These settings may be overridden on a query by query basis; however, this practice should be considered very carefully to avoid providing queries to clients that could produce unexpected results. Consider the following example.

A client submits a simple SQL statement such as:

```
SELECT v1.balance FROM accounts v1 WHERE v1.account_name
= 'bob'
```

The client is aware of what case sensitivity it wants to use. If it submits this to a case sensitive database, then it expects to only get accounts with exactly 'bob' as the name. If it submits this to a case insensitive database, it expects to get accounts with 'bob', 'BOB', and 'Bob'. If the client knows the database is case sensitive and it wants an insensitive compare then it would submit:

```
WHERE UPPER(v1.account_name) = UPPER('bob')
```

The same is true of Virtual View Manager. However, in the case where Virtual View Manager is not case sensitive and the underlying database is case sensitive, Virtual View Manager will add the UPPER function to the SQL sent to the underlying database. Unfortunately, doing this will invalidate an existing index—in the previous example, the index on account_name would be invalidated, causing a table scan.

Settings affecting query performance

To determine if your configuration settings are affecting query performance, you can evaluate any filter nodes or the SQL underlying each FETCH node in the Execution Plan to determine if case sensitivity or trailing spaces settings are impacting the query. Focus primarily on the WHERE clause or any filter nodes.

One of the two major issues that can come up is that some string comparisons in the WHERE clause have RTRIM or UPPER functions applied to them which is manifested in the FETCH node. Wrapping a column with a function such as UPPER or RTRIM will prevent the underlying system from using an index on that column. This is necessary to provide correct results, but it can affect performance.

If a filter is applied at the server level, rather than the database level, all rows must be returned from the underlying table which may also affect performance on large tables.

Review the following matrix to determine the possible impact of different case sensitivity and trailing spaces settings:

Virtual View Manager Setting	Underlying Data Source Setting	Virtual View Manager Query Behavior
case_sensitivity=true	case_sensitivity=true	None
case_sensitivity=true	case_sensitivity=false	Performs WHERE clause string comparison in Virtual View Manager instead of pushing down to database.
case_sensitivity=false	case_sensitivity=true	Adds UPPER to both sides
case_sensitivity=false	case_sensitivity=false	None
ignore_trailing_spaces=true	ignore_trailing_spaces=true	None
ignore_trailing_spaces=false	ignore_trailing_spaces=true	Performs WHERE clause string comparison in Virtual View Manager instead of pushing down to database.
ignore_trailing_spaces=false	ignore_trailing_spaces=false	None

The Virtual View Manager query engine is designed to get the correct and consistent answer regardless of the configurations of the underlying data sources. If you find an RTRIM in the WHERE clause, it is because Virtual View Manager is configured to ignore trailing spaces while the underlying data source does not ignore them. Likewise if you find an UPPER, it means that Virtual View Manager is configured to ignore case while the underlying database is sensitive to case.

Dealing with settings mismatches

There are two ways to deal with settings mismatches. First, the system wide configuration values for case sensitivity and trailing spaces can be modified via the Administration > Configuration menu. This is only useful if the data sources are fairly homogeneous in regard to this behavior. Changes to this setting should be well-considered or avoided as they will cause all other query plans to be re-evaluated to accommodate the new setting.

Second, if the data sources have varying policies for case sensitivity and/or trailing spaces, these values can be modified on a per-query basis by using SQL query options. This facility is useful when numerous types of data sources are used with varying case-sensitivity and/or trailing space settings.

These query hints should be used with an understanding that the global contract provided by IBM Cognos Virtual View Manager is overridden. It must be communicated to clients querying this published resource that the contractual behavior has been overridden.

In the previous example, the developer would use this syntax immediately after the SELECT keyword:

```
{option ignore_trailing_spaces="false", case_sensitive="true"}
```

Impact on string comparison

The case-sensitive policy affects all forms of string comparison.

The impacted functions and operators include:
- Comparison operators in WHERE and JOIN ON: = < <= >= > <>
- REPLACE(src,pattern,escape)

  The pattern is matched according to the policy.
- MIN(column)

  The strings ABC and abc are considered the same, so either may be chosen by this function.
- MAX(column)

  The strings ABC and abc are considered the same, so either may be chosen by this function.
- GROUP BY

  The strings ABC and abc are considered the same, so the group will include both sets of values.
- ORDER BY

  The strings ABC and abc are considered the same, so they will sort together and may be intermixed.

Case sensitivity does not affect the actual value of strings. Case is preserved in all cases. It only affects the comparison between strings.

Trailing spaces

Most databases perform string comparisons while ignoring any spaces at the end of the string values. For example, the test ('abc' = 'abc') is TRUE. Some databases do make use of trailing spaces, so this test would return FALSE.

IBM Cognos Virtual View Manager policy

IBM Cognos Virtual View Manager's default policy is to ignore trailing spaces. Changing the policy may affect existing queries, in that the results they return may be changed or performance may be affected.

Impact on string comparison

The trailing spaces policy affects all forms of string comparison. The affected functions and operators include:

- Comparison operators in WHERE and JOIN ON: = < <= >= > <>
- LENGTH(column)

  The string length returned does not count trailing spaces.

- MIN(column)

  he strings 'abc   ' and 'abc' considered the same, so either may be chosen by this function.

- MAX(column)

  The strings 'abc   ' and 'abc' are considered the same, so either may be chosen by this function.

- GROUP BY

  The strings 'abc   ' and 'abc' are considered the same, so the group will include both sets of values.

- ORDER BY

  The strings 'abc   ' and 'abc' are considered the same, so they will sort together and may be intermixed.

The trailing spaces policy does not affect the actual value of strings. Trailing spaces are preserved in all cases. It only affects the comparison between strings.

Impact on server performance

Performance may be impacted by the choice of policy. IBM Cognos Virtual View Manager makes every attempt to run as much of the query as possible in the underlying database. This is always possible when the Virtual View Manager policy is set to match the database's policy, but when the policies are different, some portions of the query may be executed in Virtual View Manager instead of in the database.

Whenever possible, you should match the policy of the server to that of the underlying database. If this is not desirable because you want to present a different policy or because there are multiple underlying databases with different policies, the following performance guidelines apply:

- If Virtual View Manager ignores trailing spaces and the database does not, comparisons can only be pushed to the database if the data is known to be without trailing spaces for both values. You can force this condition by wrapping values in TRIM() or RTRIM() functions.
- If Virtual View Manager does not ignore trailing spaces and the database ignores trailing spaces, and the data is not known to be without trailing spaces for both

values, the server will add an RTRIM() function to both values to ensure the underlying database performs a ignore trailing spaces compare. This should not affect server performance.

Studio locking

A locking configuration setting can force IBM Cognos Virtual View Manager users to acquire a lock prior to changing a resource. The server web services API does not honor this configuration setting.

This setting is in effect when **Studio** > **Locking** > **Enabled** is set to true.

Changing the configuration requires the Modify All Config and of course the Access Tools rights.

Procedure

1. From Studio, click **Administration** > **Configuration**.
2. Click **Studio** > **Locking** at the bottom of the configuration window and set **Enabled** to **True**.
3. **Apply** the changes, and click **OK**.

Results

This configuration change is not immediately propagated to other instances connected to the server, but any attempt to save resources will force a check to see whether the lock is enabled.

The requirement for resource locking prior to changing and saving resources may be disabled for the entire server by any administrator with the Modify All Resources right. Disable locking by toggling of the value of **Enabled** to **False**.

Existing locks will persist regardless of whether Virtual View Manager is requiring locks prior to modification. When locking is disabled users may still optionally use resource locks so that simultaneous changes are not made to a resource being revised by more than one person at a time.

Chapter 3. IBM Cognos Virtual View Manager domain administration

IBM Cognos Virtual View Manager supports the Cognos, dynamic, and LDAP domains, each of which controls a particular set of users and groups that can access Virtual View Manager. This chapter describes the Cognos domain and how to create and manage its users and groups.

The following topics are covered in this chapter:

Configuration and management of the LDAP and dynamic domains are documented in the following chapters:

The Cognos domain

The Cognos domain comprises users and groups defined within IBM Cognos Virtual View Manager. Virtual View Manager has predefined specific users and groups in the Cognos domain which you can use and modify as appropriate. You can create additional users and groups within the Cognos domain to meet your specific needs.

Administration of the Cognos domain involves creating new users and groups, changing user passwords, and granting privileges to users and groups to access the resources.

The main tool used to manage domain users and groups is Virtual View Manager Administrator. You can access Virtual View Manager Administrator in two ways:

- From Virtual View Manager, click **Administration** > **Launch Virtual View Manager Administrator**.
- From a web browser using the following URL:

  http://*server_name:port_number*/manager

  The default port number is 9400.

The **Users** page in Virtual View Manager Administrator is where you manage users and groups in the Cognos domain as well as in the LDAP and dynamic domains.

Domain management

Domain management entails adding and removing domains and the users and groups assigned to a domain. Two domains—cognos and dynamic—are already defined for use when IBM Cognos Virtual View Manager is installed.

The **Domain Management** page in Virtual View Manager Administrator lists the defined domains and provides links to view the groups and users within those respective domains.

The **Domain Management** page is used primarily for the specification of LDAP domains and for the selection and deselection of those external groups that will have rights and privileges to view and use defined resources.

Domain management for LDAP domain configurations are described in Chapter 4, "LDAP domain administration," on page 33. For more information on dynamic domain administration refer to Chapter 5, "Dynamic domain administration," on page 49.

Group management

You can create groups of users who need similar rights to perform administrative tasks on the server, and groups who need access to create, view, access, and change objects defined with IBM Cognos Virtual View Manager. Developers, operations personnel, and administrators should each have their own groups to access Virtual View Manager Administrator and other tools and options.

Group rights templates enable quick assignment of rights based on an expected level of interaction with Virtual View Manager. Group rights templates exist for: Administrators, Developers, Operations, Backup, Restore, Backup & Restore, and End Users. For more information, see "Group and User Rights Templates" in the Security chapter of the Virtual View Manager *User Guide*.

As an example, end users should belong to groups with no group rights. Typically end users are not allowed to change data source definitions, change server configuration settings, or back up servers. End users simply use JDBC, ODBC, or web service-enabled applications to trigger data requests and procedure calls that get executed in the background without further user interaction or need for additional rights.

Built-in groups

There are three built-in groups—admin, all (cognos), and all (dynamic)—which are created by the system and cannot be deleted.

- admin (cognos)—this group has administrative privileges. The admin user is a system-provided member of this group. Other users can be added to or removed from this group by anyone with administrative privileges.
- all (cognos)—this group contains all users except for the following: anonymous, nobody, system, and users of the dynamic domain. User membership is automatically maintained by the system.
- all (dynamic)—This group contains all users in the dynamic domain.

Adding groups to the Cognos domain

You can add any number of groups to the Cognos domain. When you add a group, you define the rights for that group. After you've created a group, you can add users to it.

Procedure

1. Start Virtual View Manager Administrator.
2. Click the **Users** tab, and click **Group Management**.
3. Click the **Add Group** button.
4. In the **Add a Group** window, enter the name for the new group.
5. Select the group rights template that is most appropriate for the new group. Customize the rights as required. Refer to the IBM Cognos Virtual View Manager *User Guide* description of "Group and User Rights Templates" for more information on the rights and what they will allow.
6. Add notes in the **Annotation** field to help developers identify the users, usage, and rights associated with the group. This will help with the setting of permissions on new resources and other future administration.
7. Click **OK**.

   The group is added to the **Group Management** page.

Removing groups from the Cognos domain

Administrators with the Modify All Users and the Access Tools rights may remove groups from the Cognos domain. Removing a group deletes any associated rights and privileges from group members.

Deletion of a Cognos domain group does not remove its member users from the Virtual View Manager.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **Group Management**.
2. Select one or more groups using the check box and click **Remove Groups**.

Results

Removing LDAP groups does nothing to LDAP configurations and definitions, but it does remove LDAP users and any group associated rights and privileges from the Virtual View Manager system.

Removing an externally defined LDAP group

IBM Cognos users who were members of a group deleted from the Cognos domain might still have the rights and privileges that were associated with that group. If this is the case, the rights and privileges are present because of membership in other groups or the rights and privileges were explicitly assigned directly to the user.

Deletion of a Cognos domain group does not remove its member users from the Virtual View Manager.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **Domain Management**.

2. Select the LDAP domain using the left most column radio button, and click **Edit External Groups**.

User management

User administration involves adding a user to the domain, removing a user from the domain, adding users to a group, removing users from a group, and changing passwords.

Built-in users and their privileges

The Cognos domain has the following users that are automatically created: admin, anonymous, nobody, and system. These users are permanent in the system and cannot be removed.

- admin—This user has privileges to access and use any resource in the system. admin can also grant/revoke privileges to other users. The admin user cannot be removed from the system. The admin user has a home folder (/users/admin).
- anonymous—This user is provided for anonymous login for JDBC clients and Web service clients. By default, anonymous logins are disabled. anonymous users must be explicitly given privileges to access IBM Cognos Virtual View Manager resources.
- nobody—This user cannot log in or be removed. Abandoned resources owned previously by a user that no longer exists in the system are given to nobody.
- system—This user cannot log in or be removed. It owns items that even the users with administrative privileges cannot modify.

Members of the all group, meaning all Cognos users and all dynamic users, have Read privileges for all folders created with the installation. Newly created folders and resources do not include privileges for members of the all group. Privileges must be assigned by the creator/owner of the resource, or by an administrator or user explicitly given the Grant right on that object.

All semi-editable folders, such as /shared, /services/databases and /services/webservices, have no privileges but they can be edited.

All pre-created tables/procedures have Select and Execute privileges for the all groups (in the cognos and dynamic domains) and the anonymous user in cognos; for example, /services/databases/system, /services/webservices/system, and /lib.

- By default anonymous users cannot invoke any Web services. To make Web services available to anonymous users, grant the Read privilege to /services/webservices, then grant Read to the both the data service and the port that you want the anonymous user to access and use.
- anonymous users cannot connect to the server using JDBC because no Virtual View Manager data service of the type database is automatically available. To enable them, you should grant Read privileges to services/databases, the data service, and any catalogs or schemas that you want to make available.
- Resources in the **Virtual View Manager Data Services** area point to the resources in the work area. In order to access a resource in the **Virtual View Manager Data Services** area, the anonymous user needs permission to read all the folders above that item and have appropriate permission (such as SELECT, INSERT, UPDATE, DELETE, or EXECUTE) on the item to which the resource points.

In order to expose a resource to either Web services or JDBC clients, you should grant the Read privilege to all the folders above the resource and the appropriate

permission to the resource itself. If the resource uses other resources, then you have to repeat the process with those resources as well.

This is similar to what you would do for any other user, except that some folders have the Read privilege by default for the all group and you need to override those folders.

The anonymous user is denied any access to the /users folder and admin cannot change it. This means that all published resources you want anonymous to use must reside in the /shared folder.

Adding users to the Cognos domain

IBM Cognos Virtual View Manager administrators with the Modify All Users and Modify All Resources rights can add users to the Cognos domain.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **User Management**.
2. Click **Add User**.
3. Enter the new user name and password.

   The user name is the login name for the user and can only contain alphanumeric characters and the underscore character.

   The password must be at least six characters long and it may have selected symbols and upper case alphanumeric characters. The following are some examples of valid passwords:

   joe-./:;<=>?@[\]^_`|~

   123_joe

   !23Abcd-+23

   !"#$%&'{}*+
4. Select a base template to begin rights assignment and select rights as appropriate for the local security policy and the expected level of user interaction with the server and underlying data sources.
5. Enter notes in the **Annotation** field to give future administrators an indication of the user's role in the system or organization.
6. Click **OK**.

   The newly added user name is added to the **Cognos** domain.

   Note: LDAP users are managed entirely by the LDAP server. Virtual View Manager adds the LDAP domain and selected groups. Members of those groups inherit Virtual View Manager rights and privileges for tools and resources from the rights and privileges assigned to the group from the Manager Group page and resources.

Removing users from the Cognos domain

Removing a user from the Cognos domain removes the user from IBM Cognos Virtual View Manager.

Removing a user who is derived from an LDAP domain/group does not prohibit the user from logging into the system again. See "Removing a group from an LDAP domain" on page 43 for more information.

Procedure

1. In IBM Cognos Virtual View Manager Administrator, click **Users** > **User Management.**
2. Select check box to the left of the each user to be removed from the domain and click **Remove User(s)**.

Managing group membership

A group must exist in the Cognos domain before you can try to add a user to that group.

See "Adding groups to the Cognos domain" on page 27.

Procedure

1. In IBM Cognos Virtual View Manager Administrator, click **Users** > **User Management**.
2. Select the link in the **# Groups** column for the desired user.

   The **Edit the User's Group Membership** window is displayed.
3. Select the desired groups in which the user will be a member, and click **OK**.

   All rights and privileges are inherited by group definition and user membership in that group. If a user belongs to multiple groups, no special rights and privileges are gained from having duplicate rights and privileges.

   If a user is added to the group named admin, it means that this user obtains administrative privileges in Virtual View Manager. In order to use the new privileges as an administrator, the user must log in again.

View group membership

IBM Cognos Virtual View Manager Administrator displays the groups in which a selected user belongs, and it also provides filtering to see all the members of a single selected group.

View a user' s group membership in the Cognos domain

You can view a user's group membership in the Cognos domain.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **User Management**.
2. If the user belongs to a single group, it is displayed in the **Groups** column listing for that user. Otherwise expand the **Groups** column to show the groups.

View a group's membership

You can view a group's membership.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **User Management**.
2. Select the link in the **# Users** column.

   The users in that group are then listed on the **User Management** page using the appropriate group filter.

   Alternatively, go directly to the **User Management** page and use the Domain and Group filters to show the groups.

Changing a password

IBM Cognos Virtual View Manager administrators with the Modify All Users and Modify All Resources rights can also change any Cognos domain user password, whereas non-admin Cognos domain users may only change their own passwords.

Change your own password

Non-admin Cognos domain users may only change their own passwords.

Procedure

Click **File** > **Change Password**.

Changing a user's password or setting explicit user rights

IBM Cognos Virtual View Manager administrators with the Modify All Users and Modify All Resources rights can change any Cognos domain user password.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **User Management**.
2. Select a user name.

   The resulting window allows you to change the user's password and enables assignment of rights.

   Note: Changes made to the user rights profile take effect nearly immediately as Virtual View Manager checks for appropriate rights every time feature access is attempted.

Changing ownership of resources

Administrators can change the ownership of resources one by one or as a group of resources within a container resource.

Abandoned resources owned previously by a user that no longer exists in the system are re-assigned to the nobody user. The nobody user cannot log in or be removed. The administrator can change the ownership to a valid user.

Procedure

1. In the resource tree, select a resource, and click **Administration** > **Change Owner of <resource>**.

   Or, right-click a resource, and click **Change Resource Owner**.

   The current owner's user name and domain are displayed, and a list of new owners with their user names and domains is presented.

   Ownership cannot be changed for system-owned resources or home folders.
2. Select a new owner's user name from the **User** drop-down list.
3. Optionally, select the **Apply the change recursively** check box.

   The **Apply the change recursively** check box is checked by default if the selected resource is a container. This box is unchecked for leaf resources. It is checked and disabled if the owner cannot be changed or if the resource is a physical data source. All resources within a physical data source and the data source itself are always owned by one user.

   The **Apply the change recursively** check box is enabled when the resource is not in one's home folder.

4. To selectively change the ownership, select the **Change if the current owner** is box.
5. Click **OK** to see the list of resources ready to be transferred to the new owner.
6. View the list of resources, and click **Commit** to apply the changes.

Manage user and group privileges

Resource developers, owners, and users delegated Grant privilege on a resource may also set resource specific privileges for that object. Administrators with Modify All Users or Modify All Resources privileges may also review, set, and revoke privileges for any resources and define rights for any groups and users in Virtual View Manager.

Management of user and group privileges on resources is generally best left to the developer who created and owns the resource, providing access for a security audit by an administrator at any time. It is more efficient to let the resource developers and owners set privileges for the native sources that they define, configure, and manage from within Virtual View Manager.

To facilitate decentralized management of specific resource privileges it is helpful to define groups of users with similar and well-defined roles where possible. Group assignment of privileges on the resource is encouraged for any large deployment.

See the "Privileges" section in the Security chapter of the *User Guide* for more details on access privileges.

Chapter 4. LDAP domain administration

IBM Cognos Virtual View Manager supports the following types of domains: Cognos, LDAP, and dynamic. This chapter focuses on how to configure and administer LDAP domains for use with Virtual View Manager.

The following topics are covered in this chapter:
* "About the LDAP domain"
* "Configuring LDAP properties file"
* "Domain administration" on page 40
* "LDAP user management" on page 45

About the LDAP domain

IBM Cognos Virtual View Manager can leverage enterprise LDAP implementations of Active Directory and iPlanet domains, groups, and users to authorize views and use, creation and management of Virtual View Manager defined resources.

Currently supported LDAP authentication servers are:
* iPlanet 5.1 sp2
* Windows Server 2003 Active Directory

Virtual View Manager configuration and management for use of LDAP domains requires the Read and Modify All Users rights in addition to the Access Tools right to give the administrator the ability to create and modify domains, groups, and users. Those rights also enable specification and modification of rights given to LDAP groups and users.

Manager also enables quick configuration and provisioning of LDAP (Active Directory and iPlanet) domains, groups, and users with rights and privileges enabling use, view, and change of Virtual View Manager defined resources. LDAP configurations and usage are described in the next chapter.

Configuring LDAP properties file

Query searches for retrieving user and group information are controlled by a properties file named ldap.properties.

The ldap.properties file is located in the following installation directory:

`<installation_directory>/conf/server/ldap.properties`

The file contains relevant query parameters for the two supported LDAP directory servers, Active Directory and iPlanet.

If you add LDAP domains to IBM Cognos Virtual View Manager, you should configure the ldap.properties file after a successful installation.

This section describes the structure of the LDAP properties file and also gives an example of the default LDAP properties file.

Structure of the LDAP properties file

The LDAP properties file (ldap.properties) has four sections which are described in the following sections.

These conventions are followed when describing sections in the LDAP properties file.

- TYPE must be replaced with either "iplanet" or "activedirectory".
- Property file variables are designated with a capital letter enclosed by angled brackets: <A>, <B>, ..., <X>.
- A description of each line in the ldap.properties file is provided.

Section 1: (used for querying all users)

Section 1 of the ldap.properties file is used for querying all users.

```
TYPE.all.users.search.context=<A>
```

Search-context used to find all users.

```
TYPE.all.users.filter=<B>
```

Filter to pass to a query for finding all users.

```
TYPE.all.users.username.attribute=<C>
```

Username attribute to retrieve the name of user found from a query.

```
TYPE.all.users.search.timeout=<D>
```

Search timeout value to limit the time for infinite search. 0 (zero) means infinite timeout, timeout is in milliseconds and should be greater than 0 (zero).

Section 2: (used for querying all groups)

Section 2 of the ldap.properties file is used for querying all groups.

```
TYPE.all.groups.search.context=<A>
```

Search-context used to find all groups.

```
TYPE.all.groups.filter=<B>
```

Filter to pass to a query for finding all groups.

```
TYPE.all.groups.groupname.attribute=<C>
```

Group name attribute to retrieve the name of a group found from a query.

```
TYPE.all.groups.search.timeout=<D>
```

Search timeout value to limit the time for infinite search. 0 (zero) means infinite timeout, timeout is in milliseconds and should be greater than 0 (zero).

Section 3: (used for authenticating LDAP users)

Section 3 of the ldap.properties file is used for authenticating LDAP users.

```
TYPE.user.username.comparison.is.case.sensitive=<A>
```

Sets the user name comparison to be case-sensitive or not. By default the value of <A> is true but it may be set to false.

```
TYPE.user.search.context=<B>
```

Search-context used to find the user attempting authentication.

```
TYPE.user.filter=<C>
```

Filter used to authenticate user in LDAP directory server. The USERNAME keyword will be replaced at runtime with the appropriate username.

```
TYPE.user.username.attribute=<D>
```

User name attribute to retrieve the name of the user attempting authentication from a query.

```
TYPE.user.search.timeout=<E>
```

Search timeout value to limit the time for infinite searches. 0 (zero) means infinite timeout, timeout is in milliseconds and should be greater than 0 (zero).

Section 4: (used for querying all groups for a user)
Section 4 of the ldap.properties file is used for querying all groups for a user.

```
TYPE.user.groups.search.context=<A>
```

Search-context used to find all the groups for a user.

```
TYPE.user.groups.filter=<B>
```

Filter to pass to a query for finding the members of a group. The USERDN keyword will be replaced at runtime with the appropriate user distinguished name.

```
TYPE.user.groups.groupname.attribute=<C>
```

Group name attribute for finding the name of a group to which a user belongs.

```
TYPE.user.groups.search.timeout=<D>
```

Search timeout value to limit the time for infinite searches. 0 (zero) means infinite timeout, timeout is in milliseconds and should be greater than 0 (zero).

Example of an ldap.properties file

This section presents an example of the operational lines in the default ldap.properties file. Following the file is a key for the symbols and attributes that can be used.

```
iplanet.all.users.search.context=ou=people
```

```
iplanet.all.users.filter=(&(objectclass=person))
```

```
iplanet.all.users.username.attribute=uid
```

```
iplanet.all.users.search.timeout=0
```

```
iplanet.all.groups.search.context=ou=groups

iplanet.all.groups.filter=(&(objectclass=groupofuniquenames))

iplanet.all.groups.groupname.attribute=cn

iplanet.all.groups.search.timeout=0

iplanet.user.username.comparison.is.case.sensitive=true

iplanet.user.search.context=ou=people

iplanet.user.filter=(&(uid=USERNAME)(objectclass=person))

iplanet.user.username.attribute=uid iplanet.user.search.timeout=1000

iplanet.user.groups.search.context=ou=groups

iplanet.user.groups.filter=(
&(uniquemember=USERDN)(objectclass=groupofuniquenames))

iplanet.user.groups.groupname.attribute=cn

iplanet.user.groups.search.timeout=1000

activedirectory.all.users.search.context=cn=users

activedirectory.all.users.filter=(&(objectclass=user))

activedirectory.all.users.username.attribute=samaccountname

activedirectory.all.users.search.timeout=0

activedirectory.all.groups.search.context=cn=users

activedirectory.all.groups.filter=(&(objectclass=group))

activedirectory.all.groups.groupname.attribute=cn

activedirectory.all.groups.search.timeout=0

activedirectory.user.username.comparison.is.case.sensitive=true

activedirectory.user.search.context=cn=users

activedirectory.user.filter=(&(samaccountname=USERNAME)(objectclass=user))

activedirectory.user.username.attribute=samaccountname

activedirectory.user.search.timeout=1000

activedirectory.user.groups.search.context=cn=users

activedirectory.user.groups.filter=(&(member=USERDN)(objectclass=group))

activedirectory.user.groups.groupname.attribute=cn
```

```
activedirectory.user.groups.search.timeout=1000
```

LDAP properties file symbols and attributes

The following symbols can be used in an ldap.properties file.

LDAP search context symbols

The pipe character may be used to separate multiple search context property values. This may be interpreted as a disjunction: or.

The pipe character is represented by this symbol:

|

LDAP seach filter symbols

Several symbols are available for various types of search filters.

& Conjunction (for example, and - all in the list must be true)

| Disjunction (for example, or - one or more alternatives must be true)

! Negation (for example, not - item being negated must not be true)

= Equality (according to the matching rule of the attribute)

~= Approximate equality (according to the matching rule of the attribute)

>= Greater than (according to the matching rule of the attribute)

<= Less than (according to the matching rule of the attribute)

=* Presence (for example, the entry must have the attribute returning whatever that value is)

* Wildcard (indicates zero or more characters can occur in the position); used when specifying attribute values to match

\ Escape (for escaping "*", "(", ")" when they occur inside an attribute value)

LDAP attribute key

Several symbols are available for specifying LDAP attribute keys.

o = Organization

ou = Organization Unit

cn = Common Name

dn= Distinguished Name

dc = Domain Component

Query examples

This section shows example iPlanet or Active Directory LDAP server query examples.

The iPlanet or Active Directory LDAP server configurations are similar except where the object class values, search contexts, and user/group attribute values may be different where:

```
TYPE = { iplanet | activedirectory }
```

Search for specific groups with a group filter

All group filters can use the search filter syntax described above in the above Search Filter Syntax area.

Example: Find all groups that have a prefix "cs_" in their name where "Y" is a group object class for the domain type, and "Z" is a group name attribute:

Example solution:

```
TYPE.all.groups.filter=(&(objectclass=Y)(Z=cs_*))
```

Note: This same method can be used for finding specific users also.

Specify multiple locations to find users or groups

All search context attributes can support looking for LDAP objects in multiple search contexts. Use the pipe character (|) to separate multiple search contexts.

For example,

```
TYPE.all.groups.search.context=cn=users|cn=users2
```

Disable case sensitivity for LDAP authentication

By default, IBM Cognos Virtual View Manager is case sensitive when used with either an iPlanet or ActiveDirectory domain. However, you can change this in the ldap.properties file.

For example,

```
TYPE.user.username.comparison.is.case.sensitive=false
```

When the LDAP user name comparison is not case sensitive, then the user `cn=foo,ou=users,dc=domain,dc=com` can login to a Virtual View Manager LDAP domain with user name foo or FOO. All variations of the user name used to login to Virtual View Manager tools will map to the actual user name stored in the LDAP server.

Note: If you disable case sensitive mode and have multiple users with the same name (but with variations in capitalization) login will be disabled for that user name. Search context may be specified to find user attribute values that differentiate the users affected. For instance in Active Directory, the samaccountname attribute for a user object is unique in the entire LDAP server, but cn (common name) is not.

iPlanet get all users

To start a search from the root node and get all users use a blank (null) value in the search context.

```
TYPE.all.users.search.context=
```

To find groups that match the objectclass filter use the following:

```
TYPE.all.users.filter=(&(objectclass=person))
```

To get user names from the user object name attribute:

```
TYPE.all.users.username.attribute=uid
```

To perform a search without a timeout:

```
TYPE.all.users.search.timeout=0
```

iPlanet get all users (under container ou=people)
This search context will find only groups under container ou=people:

```
TYPE.all.groups.search.context=ou=people
```

This search will find only groups that match the objectclass filter:

```
TYPE.all.groups.filter=(&(objectclass=person))
```

This search will get group names from this group object name attribute:

```
TYPE.all.groups.groupname.attribute=cn
```

To specify a search that does not have a timeout (infinite search timeout):

```
TYPE.all.groups.search.timeout=0
```

iPlanet get all groups
Using a null value (blank) will start search from the root node and get all groups:

```
TYPE.all.groups.search.context=
```

To find only those groups that match the objectclass filter:

```
TYPE.all.groups.filter=(&(objectclass=groupofuniquenames))
```

To get group names within this group object name attribute:

```
TYPE.all.groups.groupname.attribute=cn
```

To specify a search that does not have a timeout (infinite search timeout):

```
TYPE.all.groups.search.timeout=0
```

iPlanet get all groups (under container ou=groups)
This search context will find only groups under the container ou=groups

```
TYPE.all.groups.search.context=ou=groups
```

To find only groups that match the objectclass filter

```
TYPE.all.groups.filter=(&(objectclass=groupofuniquenames))
```

To get group names from this group object name attribute:

```
TYPE.all.groups.groupname.attribute=cn
```

To specify a search that does not have a timeout (infinite search timeout):

```
TYPE.all.groups.search.timeout=0
```

iPlanet and Active Directory user authentication

IBM Cognos Virtual View Manager LDAP user authentication dependent on either iPlanet or Active Directory servers depends on several configurations prior to successful user authentication through a Virtual View Manager interface.

- The LDAP server must be configured for use
- The LDAP domain must be configured for use in Virtual View Manager Administrator
- Specific Active Directory or iPlanet groups within the specified domain must be authorized to use Virtual View Manager defined resources.

  Note: All members of Virtual View Manager authorized LDAP groups get the basic set of privileges granted to the all group. Other resource privileges and Virtual View Manager rights must be assigned explicitly to the LDAP group or to the individual user.

- Only users who are members of the specified domain and authorized groups may authenticate properly using Virtual View Manager resources.

All LDAP users trying to authenticate against an LDAP server need to use the same username attribute value in the both settings below:

```
TYPE.user.filter=(&(uid=USERNAME)(objectclass=person))
```

```
TYPE.user.username.attribute=uid
```

For example, a Virtual View Manager log in using an iPlanet LDAP domain with group='g1' defined locally on the server will take this login:

```
username=user1 password=password domainname=iplanet_domain
```

and generate an LDAP authentication request as follows:

```
uid=user1,ou=people,dc=DOMAIN_NAME,dc=COM
```

The example above assumes the following:

- Virtual View Manager LDAP domain "iplanet_domain" already exists where LDAP domain "iplanet_domain" has group "g1" in its group list.
- User "user1" exists in LDAP server at:
  ```
  ou=people,dc=DOMAIN_NAME,dc=COM
  ```
- Group "g1" exists in LDAP server at:
  ```
  ou=groups,dc=DOMAIN_NAME,dc=COM
  ```
- User "user1" is a member of the group "g1"

Domain administration

LDAP domain administration involves several tasks.

These tasks include:

- Adding an LDAP domain to the IBM Cognos Virtual View Manager server.
- Adding users and groups to an LDAP domain.

- Removing users or groups from an LDAP domain.
- Changing the LDAP domain connection parameters.
- Removing an LDAP domain.

Adding an LDAP domain

You can add more than one LDAP domain to IBM Cognos Virtual View Manager, provided that each of those domains has a unique name. The names dynamic and cognos are reserved domain names in Virtual View Manager. They cannot be used as LDAP domain names.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **Domain Management**
2. Click **Add Domain**.
3. Enter the **Domain Name**.

   The domain name will be part of the login.

   When the process of adding the domain is complete, this name is displayed in the **Domain Name** column and as part of the login (lower case only).
4. Specify the LDAP directory type as either **Active Directory** or **iPlanet**.
5. Type the path to the LDAP server in the **Server URL** field using the format:

   `ldap://<hostname:port>/< suffix>`
6. Enter an administrative LDAP user name and password.
7. Click **OK**.

Working with groups from an LDAP domain

When you are adding groups from an LDAP domain to IBM Cognos Virtual View Manager, it means that you are selecting groups or users from the LDAP server and adding them to the Virtual View Manager Server. This enables differentiated group and user access, use, creation, and modification of Virtual View Manager resources as LDAP authenticated users.

LDAP domain users must belong to at least one LDAP group selected to use Virtual View Manager Server as an authenticated user enabling implicit assignment of rights, privileges, and ownership of defined resources.

Similarly when an LDAP domain group is deselected from use with Virtual View Manager Server, that group and all users defined exclusively by that group are removed locally from Virtual View Manager disallowing access as an authenticated user. The external LDAP server is unaffected by these Virtual View Manager definition changes.

After adding an LDAP group to Virtual View Manager, members of that group can be authenticated with the LDAP server. Rights may be assigned to members and data sources may define privileges for the group or individual members of the group to use resource definitions and data.

A security check on user rights and privileges is made every time any request is made by Virtual View Manager and any defined resources. Authentication status with the LDAP domain is checked and maintained with a non-persistent session.

Authenticated users may own and use resources as defined by the rights and privileges assigned to them either explicitly as individuals or implicitly by

membership in one or more groups. Members of a group defined for use within Virtual View Manager inherit all the rights and privileges defined for that group.

When the **Edit/Add External Groups** window is displayed, the currently available LDAP groups are displayed and those groups already selected for use within Virtual View Manager are shown with a marked check box.

Adding a group to an LDAP domain

Adding external groups from an LDAP domain gives the IBM Cognos Virtual View Manager system a way to support differentiated access and use of Virtual View Manager defined resources for selected user groups without including the entire domain.

Note: Adding a group is the only way to add users to Virtual View Manager from an LDAP server.

User and group management is performed on the LDAP server and Virtual View Manager rights and privileges are assigned to LDAP groups and users.

LDAP users are given rights and privileges to use Virtual View Manager resources by explicit addition of the groups in which those members belong. Ensure that appropriate groupings of users will be enabled to use Virtual View Manager resources by management of users on the LDAP servers.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **Domain Management**.
2. Select the LDAP domain., and click **Edit External Groups** at the bottom of the table.

   The **Add External Groups** window displays all groups in the LDAP or iPlanet domain.
3. Select those groups that you wish to grant access to Virtual View Manager resources.

   You can use the navigation arrows and page numbers at the bottom of the window to display additional groups. You can also change the sort order by clicking the sort icon.
4. Click **OK** and refer to the notes that follow.

   Initially no users are shown as members of the selected groups. Users from the groups appear in the system after the initial use of any Virtual View Manager resource.

   Set appropriate rights and privileges for LDAP groups in the same way that Virtual View Manager groups and users get assigned rights and privileges. Pure end-users should receive no rights but get privileges which are assigned at the individual resource level to groups and users to use data via JDBC, ODBC, or Web services clients. Unauthenticated users, anonymous, and dynamic users using pass-through authentication may be given privileges to view, access, and execute procedures on data resources, but they may not receive rights to change Virtual View Manager definitions and settings.

   Groups of developers, operations users, and administrators should get explicit rights to access tools and rights to read and/or modify Virtual View Manager resources at design time.

   After initial Virtual View Manager use, LDAP domain users may be added directly to specifically defined Virtual View Manager groups granting the user implicit rights and privileges assigned by group membership, or they may be

given explicit individual rights and privileges. Managing rights and privileges by group assignment is encouraged as role based access control enables better control of large groups of users.

Refer to "Rights" and "Privileges" in the User Guide for more information.

To add users to LDAP domains and groups, see "Adding users to IBM Cognos Virtual View Manager from an LDAP domain" on page 45, and "Adding users to groups" on page 46.

Removing a group from an LDAP domain

Removing a group from an LDAP domain means that you are removing the LDAP group, all the users within that group, and all the implicit rights and privileges given to that group definition for resources and uses on the IBM Cognos Virtual View Manager server.

Resource definitions for /shared resources owned by users removed by a group deletion retain most configurations and access privilege information for remaining groups after the LDAP based owner is removed. Resource ownership is shifted to a special system user named nobody. Those data sources should be assigned a new owner and connections to those data sources should be tested and re-introspected to ensure that the resource remains accessible.

Group deletion also removes all access privilege information for the deleted group and affected members. Group deletion also clears any personal work space associated with deleted user profiles in the /users node.

As stated earlier, the external LDAP server is unaffected by these Virtual View Manager definition changes.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **Domain Management**.
2. Select the LDAP domain.
3. Click **Edit External Groups** at the bottom of the table.

   The **Add External Groups** window displays all groups in the LDAP or iPlanet domain.
4. Deselect those groups that you wish to remove. You can use the navigation arrows and page numbers at the bottom of the window to display additional groups.
5. Click **OK**.

Viewing group membership

The IBM Cognos Virtual View Manager administrator with Read All Users right may review and monitor user group membership from the Virtual View Manager Administrator.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **User Management.**

   The table of users may be filtered by domain and group, and sorted on multiple attributes.
2. In the **Groups** column click the + icon to expand the list of groups to which the selected LDAP user belongs.

Adding and removing LDAP users from a group

You can add or remove LDAP users from a group.

The Virtual View Manager server and Virtual View Manager Administrator do not manage LDAP group membership. LDAP users may be added to Virtual View Manager groups as described in the following procedure, but LDAP groups are not modifiable from Virtual View Manager Administrator.

Procedure

1. In Virtual View Manager Administrator, click Users > Group Management.
2. Click **Edit Users** ( ) from the group to which a user should be added or removed.
   The **Edit Group Membership** window is displayed.
3. Add or remove users by checking or unchecking the users who should or should not belong to the group.

Results

LDAP users get all the rights and privileges inherited from the groups in which they belong.

Editing LDAP domain connection parameters

Editing an LDAP domain enables change of connection parameters required to connect and read data from an LDAP authentication server. Everything but the nominal, domain name display text may be modified.

Change of the LDAP password in the Virtual View Manager domain profile requires entry of the old password to unlock the encrypted cipher text.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **Domain Management**.
2. Click the domain name.
3. Modify any field as required and click **OK**.

Removing an LDAP domain

Removing an LDAP domain means that you are removing it from use by IBM Cognos Virtual View Manager. When an LDAP domain is removed, all the users, groups, rights, and privileges associated with that domain are deleted and removed from Virtual View Manager. Ownership of shared Virtual View Manager resources that were created and owned by users who will be deleted with the removal of the LDAP instance are moved to ownership by the user named nobody. The LDAP users and groups will be untouched on the LDAP server.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **Domain Management**.
2. Select the domain name.
3. Click Remove Domain.
4. Click **OK**.
   The domain, groups and users from that domain are no longer configured for use of Virtual View Manager resources.

Resources that were created or owned exclusively by a deleted user are assigned to the nobody user. Privileges to utilize resources owned by nobody remain the same for those groups and users who remain after the LDAP domain is removed.

LDAP user management

By default, without additional rights and privileges, all members of LDAP groups selected for use with IBM Cognos Virtual View Manager will be able to log into JDBC, ODBC, and Web services clients configured for use with Virtual View Manager. Rights to use Virtual View Manager tools and to view and use other resources must be added to group definitions or assigned explicitly to the user.

Assignment of rights and privileges to users from an LDAP domain is much the same as assignment of rights to users from a Virtual View Manager domain.

Only a user with Read/Modify All Users, and Access Tools rights may add/modify an LDAP domain, add/remove groups, and clear and reset LDAP users to group settings.

Adding users to IBM Cognos Virtual View Manager from an LDAP domain

Under most conditions, LDAP users are added indirectly by addition of the groups to which they belong. Group management of rights and privileges may be assigned for users depending on the role performed by members of the group. To add users to an LDAP domain, the IBM Cognos Virtual View Manager administrator first has to add that LDAP domain to the Virtual View Manager Server, and then add groups to that domain.

To add a user to an LDAP domain, the user must belong to a group in the LDAP server. See "Adding and removing LDAP users from a group" on page 44 for details.

When adding a user from an LDAP domain, three conditions must be satisfied to add a user to an LDAP domain.
- The entered LDAP username and password for the user are authenticated with the LDAP server successfully. If LDAP authentication fails or the LDAP user does not belong to any local group definitions, the user is not added to the LDAP domain in the Virtual View Manager Server. Consequently, this LDAP user is not allowed to log into Virtual View Manager successfully.
- The LDAP user is already a member of a group defined by the LDAP server.
- That pre-existing LDAP group is defined for use by Virtual View Manager.

If these conditions are met, the user is successfully added to the Virtual View Manager server. This user will be added to each local LDAP group as a member where appropriate. The domain sync process adds or removes LDAP users either to or from the appropriate local LDAP groups.

Procedure

Start Virtual View Manager, and log in with a valid LDAP username, password, and domain.

Removing LDAP users from IBM Cognos Virtual View Manager

Removing a user from a domain and group configured for use in IBM Cognos Virtual View Manager only removes the user locally from Virtual View Manager Server while the user may still exist in the LDAP server and possess implicit rights and privileges given by membership in the LDAP domain and group. Simply removing a user who is derived from an LDAP domain/group does not prohibit the user from logging into the system again.

Removing and preventing an LDAP user from gaining access to resources defined by Virtual View Manager requires one of three courses of action:

- The LDAP group membership can be redefined at the source directory to exclude the undesired user.
- Rights and privileges for the entire LDAP group can be restricted to exclude access to resources and permissions. If other members of that LDAP group require rights and privileges, then those users may be given explicit rights and privileges or they may be made members of a Virtual View Manager group that grants appropriate rights and privileges after they have been initialized into the system.
- The entire LDAP group can be removed from those included in the Virtual View Manager external groups list.

In normal usage individual LDAP users may only be removed from Virtual View Manager superficially. Virtual View Manager services are not normally used as interfaces to manage LDAP users directly, though conceivably Virtual View Manager could be configured and given permissions to change those tables directly and programmatically. Almost always it is more advantageous to manage users and group memberships using the more familiar enterprise ready Active Directory or iPlanet interfaces. For example, if an individual LDAP Active Directory user needs to be locked out to prevent Virtual View Manager access, then a management task must be performed directly on the LDAP server to change the column values for memberOf to exclude that user from the group given permission to authenticate for Virtual View Manager access.

In Virtual View Manager Administrator, users may be removed easily enough, but LDAP users selected for removal are only removed temporarily because LDAP group membership will continue to give implicit rights and privileges. Removing an LDAP user resets rights and privileges to those rights and privileges inherited by group membership. Any user workspace in Virtual View Manager is also deleted by user removal, but that space is recreated when the use logs in again.

Some possible workarounds:

- An LDAP group may be deleted to remove all group users, rights and privileges for that group. See, "Working with groups from an LDAP domain" on page 41 for more information.
- Group rights and privileges may be restricted to initially grant nothing to all members of the group, and then other members of the targeted group may be added to other better defined groups of similar users with the desired set of rights and privileges.

Adding users to groups

LDAP users are added to LDAP groups by the LDAP administrator.

LDAP and dynamic users may be added to IBM Cognos Virtual View Manager groups to enable assignment of rights and privileges by a group membership.

Though in the case of dynamically defined users it is recommended that they are not given any rights, nor any privileges to resources outside of the public domain. For more information refer to the chapter on the Dynamic Domain.

Add LDAP users to Virtual View Manager groups using the Group Management page in Virtual View Manager Administrator to add or remove users to and from the group.

Any users added to a group gain any additional rights and privileges associated with that group. Rights and privileges accumulate based upon explicit assignment and group memberships.

Chapter 5. Dynamic domain administration

IBM Cognos Virtual View Manager supports the following types of domains: Cognos, LDAP, and dynamic. This chapter focuses on how to enable and administer dynamic domains for use with Virtual View Manager.

The following topics are covered in this chapter:
- "Dynamic domains"
- "Domain administration"
- "Group administration" on page 50
- "User administration" on page 51

Chapter 3, "IBM Cognos Virtual View Manager domain administration," on page 25, and Chapter 4, "LDAP domain administration," on page 33 are discussed in the previous chapters.

Dynamic domains

Dynamic domains enable users to negotiate direct access to a secured data source by way of a IBM Cognos Virtual View Manager server pass-through login. The Virtual View Manager system does not store the password of dynamic users; it retains only an ephemeral encrypted copy in memory available during the current user session (the timeout setting is configurable).

When a user requests a view or procedure that requires data from a source that has pass-through login enabled (by using the Virtual View Manager data source driver configuration setting), the user login and the parsed request for data are passed directly to the secured data source. This pass-through allows existing data source security structures to handle the authentication and request authorization. The dynamic domain lets the developer defer security authorization and enforcement to the data source security which is presumed to be more stringent and tightly controlled.

With the dynamic domain, the Virtual View Manager solution may be made more transparent. The end user may use their existing login information for authentication with a data source to gain the same permissions they had in the past without needing to separately log into Virtual View Manager.

Note: Only one login (user name and password) is permitted for dynamic domain pass-through authentication. So, more than one pass-through-enabled data source may be used for federated queries if the data sources are set to authenticate using the same login.

Dynamic domains also enable a potentially large user base that does not require either a Virtual View Manager or an LDAP domain structure.

Domain administration

Aside from enabling the IBM Cognos Virtual View Manager configuration settings to enable the dynamic domain no special user management is required to enable users to access resources given proper privileges on resources that have been selected for exposure to dynamic domain users.

User login specifying the dynamic domain using JDBC, ODBC, or Web services is sufficient to dynamically create a new user profile.

For security reasons, dynamic domain users are blocked from using Virtual View Manager utilities. Dynamic domain users and the dynamic all group are given no rights by default. It is strongly recommended that no rights be assigned to dynamic users or groups so that they remain pure end-users without rights for changing the system.

The dynamic domain is disabled by default Virtual View Manager configuration setting. Once enabled, dynamic users have default Read access to basic resources in Virtual View Manager. See "Granting privileges to dynamic domain users" for the specific resources.

Enabling the dynamic domain

By default, the dynamic domain is disabled and an attempt to log in using this domain will fail as if the domain did not exist. This domain needs to be enabled before it can be used to log in.

Procedure

1. In IBM Cognos Virtual View Manager, click **Administration** > **Configuration**.
2. In the **Configuration** window, expand **Virtual View Manager Server** > **Configuration** > **Security**, and select **Enable Dynamic Domain Login**.
3. In the right panel, set the **Value** to **True** to enable the dynamic domain, and click **OK**.

Group administration

The dynamic domain has only one group named all. All dynamic users belong to the all group. No additional dynamic groups may be created.

The dynamic domain cannot utilize groups for differentiation of user permissions by group assignment of privileges or rights because no password is stored to authenticate who is currently using a given user name. The data sources enabled with pass-through login perform the authentication and authorization security.

Granting privileges to dynamic domain users

Resources may be opened for use by anyone including dynamic domain users by granting privileges to the dynamic all group on published resources.

Note: Dynamic all privileges open published resources to public access.

No rights should be given to dynamically authenticated users because anybody can log in as a dynamic user; they are not authenticated by the IBM Cognos Virtual View Manager system.

When the dynamic domain is enabled, dynamic users have default Read access to the following basic resources in Virtual View Manager:
- /services
- /services/databases
- /services/webservices
- /services/webservices/system
- /shared

- /lib

All other access privileges must be explicitly granted for either the dynamic all group or for the individual dynamic user after initial login.

Note: Dynamic users cannot be authenticated by definition as the password is not stored. Assigning resource privileges to individual dynamic users opens a resource to any user who may utilize that user name.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **Group Management**.
2. In the **# Users** column, select the number representing the count of users in the all group row of the dynamic domain.

   Users in the all group of the dynamic domain are shown.

User administration

Management of dynamic domain users is mostly passive as far as IBM Cognos Virtual View Manager is concerned. Data sources enabled with a pass-through login must be configured to authenticate the user and to authorize access to data.

Initial login of a dynamic domain user with a JDBC, ODBC, or Web services client creates a new user profile on Virtual View Manager. The new user is assigned an ID and may be treated like a normal user given proper caution to avoid inadvertent exposure of sensitive resources. Dynamic domain users do not have a home directory, hence they cannot create or own resources.

Note: Assigning resource privileges to any dynamic user exposes that resource to potential public access by any client using that user name. In sensitive environments dynamic users and the dynamic all group should only get privileges to access public resources, while data sources enabled with pass-through login may independently authenticate and authorize dynamic users to gain access to secured data.

Individual users in the dynamic domain may be deleted, but the all group and the dynamic domain are protected from deletion.

Note: Deletion of a dynamic user does not prevent that user name from being used to log in again.

The password for a dynamic domain user does not persist across sessions for logging purposes. But the password used for the current session is kept in memory and is passed when a request is made to data sources that have the pass-through option enabled.

Adding users to the dynamic domain

Use these steps to add users to the dynamic domain.

Procedure

1. Enable the dynamic domain as described in "Enabling the dynamic domain" on page 50.
2. Connect to Virtual View Manager through JDBC/ODBC specifying dynamic as the domain value in the connection string.

See the "Client Interfaces" chapter in the Virtual View Manager *User Guide* for information on connecting via JDBC/ODBC.

The following sample command uses the JDBCSample.bat program to run from the command line to create a user named newuser in the dynamic domain:

JdbcSample.bat system localhost 9401 newuser password dynamic "SELECT * FROM ALL_USERS"

Removing users from the dynamic domain

Removing users from the dynamic domain is mostly meaningless if the dynamic domain is enabled for use. Dynamic users are not authenticated, nor prevented from accessing all resources provided by privileges granted to the dynamic all group.

If you were to go through the motions to remove a user from the list of users registered by the dynamic domain, then that would remove any group membership that had been assigned to the user, but the user would still be able to use a client with that same user ID for login.

Dynamic users group membership

Dynamic users are created at first login. All dynamic domain users automatically becomes a member of the dynamic all group.

See "Adding users to the dynamic domain" on page 51 for details on adding users to the dynamic domain.

It is not recommended to add dynamic users to Cognos domain groups unless the group privileges are an entirely public set of permissions and no IBM Cognos Virtual View Manager rights are given.

Viewing dynamic user group membership

Dynamic users should not get regular group membership because the user name may be used by any individual using the same login name to gain access to group resources.

Procedure

1. In Virtual View Manager Administrator, click **Users** > **User Management**.
2. Select the dynamic domain just above the table to filter the users.

   If a user belongs to more than the dynamic all, then an integer will be displayed next to the icon in the **#Groups** column.
3. Select the link in the **#Groups** column to view or edit that users group membership.

   The **Edit User's Group Membership Information** window displays a list of the groups to which the user belongs. Though it is not shown, all users in the dynamic domain belong to the dynamic all group.

Chapter 6. System monitoring

IBM Cognos Virtual View Manager Administrator provides system management capabilities.

System management entails monitoring of system activities, system status, events, and system data and performing administrative tasks such as licensing and user management. You access the system monitoring information using Virtual View Manager Administrator. You can also access the same information from the **Administrator** tab in Virtual View Manager.

Virtual View Manager Administrator displays summary views of status, server information, cached resources, data sources, requests, sessions, transactions, triggers, and event logging. It also enables domain, group, and user management including configuring LDAP servers and selected groups for use with Virtual View Manager. License management, and SSL keystore management are also provided.

The following topics are covered in this chapter:
- "Working with IBM Cognos Virtual View Manager Administrator"
- "Administrator home page" on page 56
- "Server Overview page" on page 57
- "Cached resources" on page 60
- "Data sources summary information" on page 62
- "Requests" on page 64
- "Sessions" on page 66
- "Transaction summary information" on page 68
- "Triggers" on page 69
- "Events" on page 71
- "Event and log files" on page 73

Working with IBM Cognos Virtual View Manager Administrator

IBM Cognos Virtual View Manager Administrator enables users with appropriate rights to view, monitor, and update selected summary views and status. Additionally, properly authorized users may perform some server management tasks, and manage domains, groups, and users and their associated rights.

This section describes how to launch Virtual View Manager Administrator and use its basic features.

Launching IBM Cognos Virtual View Manager Administrator

IBM Cognos Virtual View Manager Administrator runs as a Web browser process.

Procedure
1. Do one of the following:
   - From Virtual View Manager, click **Administration** > **Launch Administrator (Web)**.
   - In a Web browser, enter:
     `http://server_name:port_number/manager`

The default URL with default port setting from the localhost would be:

`http://localhost:9400/manager`

2. Enter a username and password.

Using IBM Cognos Virtual View Manager Administrator

Many IBM Cognos Virtual View Manager Administrator pages have features like adjustable page refresh settings, tables with sort functionality, detail buttons to display more information about a table row, row selection check boxes to specify the performance of an action, and table filters that sharpen focus on the rows of the display. This section describes how you can use these features.

Refreshing the current page

Most IBM Cognos Virtual View Manager Administrator pages have a refresh mechanism in the upper right corner of the page. The refresh rate specifies the time interval for an automatic refresh of the data on the currently displayed page.

When you set the refresh rate, keep these things in mind:

- Page refresh rates are set independently of one another.
- Settings persist across sessions
- Different users have different refresh rate settings for each page.

You can manually refresh the page at any time by clicking **Refresh Now** 🔁 .

Sorting

Table column header text displays a small white arrow in the column header row to show if the table is sorted by that column in ascending or descending order. Secondary and higher order sorting is indicated by a gray arrow in the columns used to further organize row display order.

Procedure

Click **Sort**.
The **Advanced Sort** screen appears and allows you to define your sorting preferences.

Selecting a row for an action

Check box selection enables buttons that perform actions on the selected row or rows.
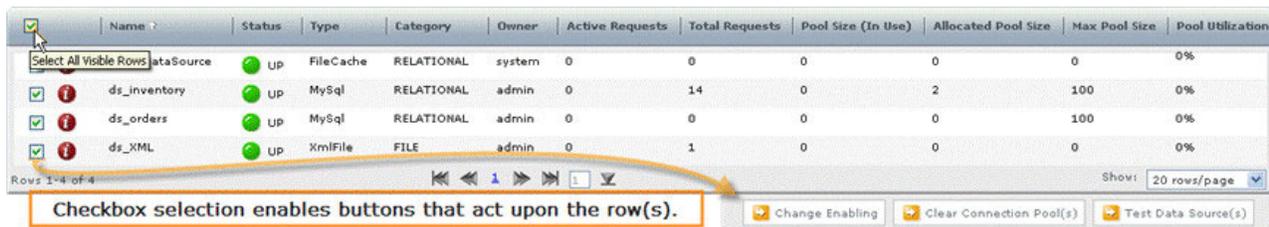
Getting additional row information

You can view more information about a row.

Procedure

Click **Show Row Details** ⓘ .

Filtering table data

A table filter determines what rows are displayed based on the rules you specify for the data in the table columns.

Create a new filter:

You can create a new filter for a table.

Procedure

1. Above any table in Virtual View Manager Administrator, click **Filter** > **Edit Filters**.

   The **Advanced Filter** screen appears.

2. Click **Add Filter** .

3. In the **Name** edit box, type a unique name for your filter.

4. Specify rules for your filter:
   - **Filter By** specifies a column in the table on which to apply the rule. The drop-down menu lists all the columns displayed in the table view.
   - **Operator** specifies the operator for the rule. The drop-down lists all available operators for type of data in the column (numeric, text, list of values, and so on).
   - **Condition** specifies the value or condition the data in the column must match.

5. Click **Add Rule** to add another rule.

6. Click **Remove Rule** to remove a rule.

7. Click **Match Any Rule** or **Match All Rules** to specify how you want the filter to work.
   - **Match Any Rule**—the filter is applied if any one of the rule conditions are met.
   - **Match All Rules**—the filter is applied only if all of the rule conditions are met.

8. Click **OK** to save the filter. The filter is saved with this table and can be applied or edited any time the table is viewed.

Copy an existing filter:

You can copy an existing filter for a table.

Procedure

1. Above any table in Virtual View Manager Administrator, click **Filter** > **Edit Filters**.

   The **Advanced Filter** screen appears.

2. Select the filter in the list box on the left.

3. Click **Copy Filter** .

4. Edit the name of the new filter.

5. Make changes to the filter as desired.

6. Click **OK** to save the filter.

   You cannot edit the default All Errors or **All Warnings & Errors** filters; however you can copy and add to them if so desired.

Remove an existing filter:

You can remove an existing filter from a table.

Procedure

1. Above any table in Virtual View Manager Administrator, click **Filter** > **Edit Filters**. The Advanced Filter screen appears.

2. Select the filter in the list box on the left.

3. Click **Remove Filter** .

   You cannot remove the default **All Errors** or **All Warnings & Errors** filters.

Navigating table data

The arrows at the bottom of many tables enable you to quickly navigate forward and backward through the data.

Controlling the number of rows displayed

The drop-down list in the lower right corner of many pages lets you select the number of rows that are displayed on each page. You can choose to display 10, 20, 50, or 100 rows per page.

Administrator home page

The IBM Cognos Virtual View Manager Administrator home page is the first page displayed when you log in. The page provides a quick summary of the current status.

Server Info panel

The Server Info panel shows summary information and links to the Server Overview page.

The **Server Info** panel includes this information:

- **Server Name** - the HTTP base port is displayed. All other ports are derived from the HTTP base port as follows: base port +1 = JDBC and ODBC base port +2 = HTTP SSL base port +3 = JDBC SSL base port +6 = Monitor base port +8 = Default for Repository

- **Total Memory Used**

  Percentage of total available Java Heap Memory (RAM) currently in use. Java Heap Memory is a configuration setting that requires restart to change. Total memory is further divided into Managed and Reserved memory with a built in margin to prevent OOM errors.

- **Sessions**

  The number of active sessions with a theoretical estimate of the total number of sessions that could be supported.

- **Server Requests**

  Number of active requests made to the server. Active requests have been started but not yet completed. The total count is cumulative of all requests.

- **Datasource Requests**

  Active data source requests sent to other resources and the total number of outgoing requests that the server has made on other resources. The total includes all requests completed or otherwise.

- **Transactions**

Active transactions with the total number of transactions that the server has made on other resources.

Server Status panel

You can view simple indicators showing the current aggregate status of the modules/consoles listed.

The **Server Status** panel displays simple indicators showing the current aggregate status of the modules/consoles listed:

Status can be one of the following: **OK** (green), **Disabled** (grey), **Warnings** (yellow), **Down** or **Errors** (red). A single warning or critical error will change the status from green to yellow, or from yellow to red depending on the failure severity and the module.

Quick Links panel

The Quick Links panel provides direct links to the Event Log and User Management pages.

The **Quick Links** panel provides direct links to the **Event Log** and **User Management** pages.

Server Overview page

In IBM Cognos Virtual View Manager Administrator, you access server information that provides a consolidated overview of the system, and the overall status of all of the other system components.

You access server information from **Monitoring** > **Server Overview**. The **Server Overview** page shows statistics about the server that are grouped into four areas: server status information, sessions and requests information, cache information, and server status indicators. In addition, two server overview buttons are provided at the bottom of the page. Each of these page elements are described below.

Server status information

This section describes the server status information. Where appropriate, the related configuration parameter in IBM Cognos Virtual View Manager Administrator is also described.

This information is available from the upper left section of the **Server Overview** page.

Status

Status reports the presence and count of errors and warnings from all pages under the **Monitoring** tab.

Status can be:
- **OK**—no errors or warnings exist.
- **Warnings**—the total number of warnings on the server. If there are also errors, they are shown in place of showing a count of warnings.
- **Errors**—the total number of errors.

Server name

Server Name with the HTTP base port is displayed. All other Virtual View Manager ports are derived from the HTTP base port.

Virtual View Manager ports are derived from the HTTP base port as follows:

base port +1 = JDBC and ODBC

base port +2 = HTTP SSL

base port +3 = JDBC SSL

base port +4 = Reserved

base port +5 = Reserved

base port +6 = Monitor

base port +7 = Reserved

base port +8 = Default for Repository

Total memory used

Percentage of total available Java heap memory currently in use. Changing the Java heap memory requires the server to be restarted. Total memory is further divided into Managed and Reserved memory with a built in margin to prevent Out of Memory (OOM) errors.

Maximum viewable events

Maximum number of events that may be viewed from the **Events** console.

This number also controls the maximum number of rows of information displayed in the table in each console. Additional entries may be viewed in the log files.

Maximum event entries

Maximum number of events to be stored in the repository. When the number of events reaches this threshold the oldest events are discarded. The log files are generally configured to retain a more expansive archive of event entries.

Session and request information

You can view summary information about sessions and requests.

The **Server Overview** page displays the following summary information.

- **Sessions** - The number of currently connected user sessions, and the total number of sessions started since the server started to run, including closed sessions.
- **Requests** - Active requests, started but not completed, and the total number of incoming requests made to the server. Includes the requests that have been completed.
- **Data Source Requests** - Active data source requests sent to other resources versus the total number of outgoing requests that the server has made on other

resources. The total is the resource cache storing resource meta data number of resources loaded on the server.includes all requests completed or otherwise.

Privilege, user, and repository caches

You can view summary cache information to determine usage of system caches for security privileges on defined data resources, user privileges, and data source metadata repository stores. These system caches are not to be confused with the data source caches which store materialized views specifically configured at the data source level. The system caches store data such as user session values of privileges, recently introspected data source meta data, and execution plans.

Summary cache information is displayed on the **Server Overview** page.

Access (Hits/ Accesses) displays a percentage that for all three rows should be relatively high, as it is a fair indicator of enhanced performance obtained by system cache usage.

Access is any request to access an object in the repository, and **Hits** are inquiries sent to the cache or successful cache usages.

A miss is an access attempt that was required to look beyond the cache for a particular entity, meaning a disk access, LDAP or data source query.

Thus a high percentage of hits to total access attempts is one indicator of enhanced performance meaning that most of the entity access attempts are hitting the cache without requirement of disk or source data retrieval.

The second column in the sub-section is the **Capacity (Entries/Max)** which shows the amount of repository usage by each of the system caches. Each of the system cache sizes is configurable.

- **Privilege Cache** - Privilege cache refers to repository storage of explicit privileges for resources.
- **User Cache** - Current user cache data stored in the repository.
- **Repository Cache** - Repository cache is a resource meta data store enabling quick use of configured resources.

Server status indicators

You can view a summary status for other Administrator pages. The red, green, and yellow indicators give you a quick idea of the status of key system components. You can click on any of the links to go to the related page for detailed information.

The **Server Status** summary box displays a summary status for other Administrator pages which are all available for display from the **Monitoring** and **Logging** tabs.

Working with the Server Overview page

You can start and stop a server or clear the repository cache for the server.

Start and stop a server or clear the repository cache for the server by clicking these buttons:

- **Stop** - Stops the server after acknowledgement of a verification prompt. Stopping the server requires the Modify All Status right.

• **Clear Repository Cache** - Immediately empties the repository cache. Clearing the repository cache requires the Modify All Status right.

Cached resources

In IBM Cognos Virtual View Manager Administrator, you can access cache resources information.

You access this information by clicking **Cached Resources** from the **Monitoring** menu. The **Cached Resources** page is displayed, providing information about the cached views and procedures. Summary information is displayed at the top of the **Cached Resources** page, and information about each individual cache is displayed in the table below.

The following sections describe the data displayed in the **Cached Resources** table and how you can work with this data.

Working with the Cached Resources page

You can enable or disable cached tables and procedures and you can refresh them manually if necessary. Each row has a check box to selectively choose the cache you want to enable, disable, or refresh.

After selecting the cache(s), you can click these buttons:
• **Change Enabling**—toggles the enabled/disabled status of selected caches. You must have the Modify All Status right.
• **Refresh Cache(s)**—refreshes the selected caches. You must have the Write privilege on the selected resource.

You can also change the number of rows displayed at once, quickly navigate through the pages using the arrows at the bottom of the table, and sort and refresh the data. These actions are described in "Using IBM Cognos Virtual View Manager Administrator" on page 54.

The Cached Resources table

You can view summary information with cache details available.

The **Cached Resources** table displays summary information. Click **Show Row Details** button 🛈 to view the **Details** window.
• **Name**—the display name of the view or the procedure.
• **Status**—current status of the cached view. The status of a cached resource can be any one listed in the following table:

Status	Event
Not Loaded	Cache has not been loaded.
Up	Cache has been loaded successfully.
Down	Cache is not loaded and the most recent refresh failed.
Stale	Cache is loaded with valid data, but the most recent refresh failed. However, reads against the cache can succeed.
Disabled	Cache has been disabled.

Status	Event
Config Errors	Cache cannot operate due to a configuration error.

- **Type**—values may be Table or Procedure.
- **Variant**—unique set of procedure input parameter values. Every set of procedure input parameters have a different storage table result set.
- **Owner**—resource owner. The cache is refreshed and cleared using the owner's identity.
- **Last Access**—date and time of the last end-user invocation of a view or procedure, also includes last refresh of data by timer, or last change of meta-data.
- **Last Refresh End**—completion date and time of last query refresh.
- **Last Fail End**—date and time when the last refresh attempt failed.
- **Storage Used**—disk space used to store result of table or procedure variant.

Cached resource details

You can display the fully qualified path and other details about cached resources.

Each row has a **Show Row Details** button which you can use to display this information.

In addition to the information presented in the **Cached Resources** table (and described in"The Cached Resources table" on page 60), these details are provided:

- **Path** — resource location, fully qualified path.
- **Owner domain** — domain of the user who created the resource, or who is currently designated as owner.
- **Total Accesses** — count of the number of times the cache resource is used since last server restart.
- **Last Success End** — last successful completion date and time.
- **Last Success Duration** — time (seconds) required for last successful refresh.
- **Last Fail Duration** — time recorded for last failure.
- **Total Successes** — count of successful refreshes since last sever restart.
- **Total Failures** — count of failed refreshes since last server restart.
- **Message** — error message returned from cache refresh failure.

Data sources

You can access data source information, providing information about all data sources added to the server.

To access data source information, click **Data Sources** from the **Monitoring** menu. The **Data Sources** page is displayed, providing this information:

- A consolidated overview of all the data sources in the repository.
- The overall status of the data sources with a count of warnings if any.
- An aggregated count of the active requests and an accumulated count of the total number of requests handled by the server since the last restart.
- Estimations of the total volume of data passed from the server to all data sources and back to since the last restart.

Data sources summary information

You can access summary information about data sources.

Summary information at the top of the **Data Sources** page includes:

- **Status**—displays the current status which can be **OK** (green), **Disabled** (grey), **Warnings** (yellow), **Down** or **Errors** (red). A single warning or critical error will change the status from green to yellow, or from yellow to red depending on the failure severity and the module.
- **Requests**—displays the number of active requests and the total number of requests since server restart.
- **Bytes**—displays the number of bytes sent to the data sources and number of bytes received from the data sources.

Working with the Data Sources page

You can perform several actions on selected data sources.

You can select one or more data sources by check box and then perform the following actions those data source(s):

- Enable or disable the data source

  The **Change Enabling** button toggles the status of the data source. Enabling makes the data source accessible via IBM Cognos Virtual View Manager definitions and configurations. Disabled takes the data source offline and makes it inaccessible to defined channels.
- Clear the currently allocated pool connections.

  The **Clear Connection Pool(s)** button drops the current connection pool allowing current processes to restart connections when necessary.
- Verify the data source connection using the **Test Data Source(s)** button.

You can also test the current status all data sources by clicking the **Test All** button. An administrator user with the Modify All Status right may use the **Test All** button.

The Data Sources table

For each data source in IBM Cognos Virtual View Manager, several columns of information are available.

The following columns are shown for each data source.

- **Name**—User-defined name for the datasource.
- **Status**—Overall status for the data source, which may be one of the following:

Status	Indicates
Disabled	The data source is disabled; represented by a grey circle.
Down	The data source is inaccessible; represented by a red circle.
Not Tested	The status of the data source has not been tested.
Up	The data source is connected to the server; represented by a green circle.

- **Type**—Native data source type, some of the more common supported data sources categories include: IBM DB2, Custom Java Procedure, Infirmity, FileCache, LDAP, Microsoft Access, Microsoft Excel, Microsoft SQL Server, MySQL, Netezza, Oracle, Sybase, Teradata, Virtual View Manager, Wsdl, XML, and XmlHttp
- **Total Requests**—Cumulative count of all requests made on data sources via the server since the server was started.
- **Active Requests**—Current count of all outstanding data source requests
- **Pool Size (In Use)**—Current connection-pool size for relational data sources
- **Allocated Pool Size**—Current number of connections allocated for the server for a particular relational data source
- **Max Pool Size**—Maximum connection-pool size for a relational data source, zero is unlimited or not applicable.
- **Pool Utilization**—Utilization of pool represented as a percentage where allocated connections is divided maximum connections for a relational data source.

Data source details

You can display read-only detail information for a data source in the repository.

To display the following read-only detail information for a data source in the repository, click **Show Row Details** .

- **Name**—User-defined name for the data source.
- **Path**—Fully-qualified path to the data source. For example, if the data source ds_orders reside in /shared/sources, the path to ds_orders would be: /shared/sources.
- **Status**—Current status, which can be Up, Down, or Not Tested.
- **Category**—Category, which can be File, LDAP, Relational, WSDL, XML/HTTP.
- **Type**—Type of the data source within the category to which it belongs.
- **Total Requests**—Total number of requests, including active requests, made to the server since last startup.
- **Active Requests**—Number of in-progress requests to the server.
- **Bytes From Data Source (Estimated)**—An estimate of the total number of bytes of data received by the server from this data source.
- **Bytes Into Data Source (Estimated)**—An estimate of the total number of bytes of data sent to this data source from the server.
- **Pool Size (In Use)**—Current connection-pool size, if the data source is relational.
- **Allocated Pool Size**—Current number of actual connections both idle and active allocated by the server for a particular relational data source.
- **Max Pool Size**—Configurable setting for maximum connection-pool size used to limit the number of connections allowed to burden a relational data source.
- **Pool Utilization**—Utilization of pool represented in percentage, if the data source is relational.
- **Number of Logins**—The number of times a connection to the data source is made in the connection pool.
- **Number of Logouts**—The number of connections to the data source that were manually destroyed by logout from the connection pool.

Requests

You access information about all current requests for service.

You access requests information by clicking **Requests** from the **Monitoring** menu. The **Requests** page provides information about all current requests for service including:

- Inbound requests through a data service.
- Outbound requests against physical data sources.
- Internal requests against internal views.

Summary information is displayed at the top of the page, and information about each individual request is displayed in the table below. Operational information about queued, in process, and recently completed requests gives the administrative user an idea about what requests are taking inordinate amounts of time or memory resources to complete.

Note that some of the information displayed on the page is controlled by the configuration settings in Virtual View Manager. To view these settings, open Virtual View Manager and click **Administration** > **Configuration**. You can also access request information from:

- **Virtual View Manager Server > Configuration > Events and Logging > Event Generation > Request Events**
- **Virtual View Manager Server > Runtime Processing Information > Requests**

Requests are removed from the table periodically, based upon the purge period configuration setting found at **Virtual View Manager Server > Runtime Processing Information > Requests > Request Purge Period**. The default setting purges requests every 5 minutes.

Requests summary information

You can access summary information about requests.

The **Requests** page provides the following summary information:

- **Status**—Aggregated status of all requests can be OK, Warning, Error, or Unknown. A single warning or error supersedes display of a status of OK.When failed requests or waiting requests are present, a count of those warnings or errors will be shown.
- **Waiting Requests**—Current number of requests waiting in the queue due to memory constraints.
- **Waiting Requests Threshold**—An event trigger threshold that causes an event. The event may be used for notification.
- **Server Requests**—The number of currently active requests, and the total number of requests made to the server since the server was started.
- **Data Source Requests**—The number of currently active requests, and the total number of requests made to the data sources since the server was started.

Working with the Requests page

You can perform several actions for selected requests.

You can select one or more requests in the **Requests** table by check box and then perform the following actions on those requests:

- **Clear Plan Caches**—Clears all query plan caches, in the event that current statistics gathering may have significantly changed information sufficient to change the query execution plan. Forces recalculation of all query plans at next time of execution which will be an initial performance hit.
- **Purge Completed Requests**—Immediately removes all completed and failed requests. This is useful to reset the view, removing any noise, just prior to testing a set of requests.
- **Cancel Requests**—Clears all requests.

The Requests table

For each request, several columns of information are available.

The **Requests** table displays these columns for each request:
- **ID**—Unique request identifier.
- **Status**—Can be any one of the following values:

Status	Description
STARTED	The request was created but not invoked or executed.
RUNNING	The request is executing.
WAITING	The request is in a wait queue.
COMPLETED	The request successfully completed but is not yet closed.
CLOSING	The request is in the process of closing.
SUCCESS	The request successfully executed and is closed.
FAILED	The request execution failed.
TERMINATED	The request was cancelled and closed.
COMMITTED	Changes were committed to the database.
ROLLED_BACK	Changes that might have been made were rolled back.
TOP_TIME	Is in the group of requests that took the longest amount of time to complete. The number of requests in this group is configurable using this property: Virtual View Manager Server > Runtime Processing Information > Requests > Number of Top Requests Tracked The default is 10 requests.
TOP_MEMORY	Is in the group of requests that took the largest amount of managed memory. The number of requests in this group is controlled by the same property described for TOP_TIME above.

- **Owner**—Userid of the user who submitted the request.
- **Parent ID**—Unique ID for the request's parent process.
- **Session Type**—STUDIO, HTTP (Web service), INTERNAL, or client procedure.
- **Session Name**—Name of the component that initiated this request.
- **Start Time**—Time the request started to execute.
- **End Time**—Time the request was completed. Blank if the request is unfinished.

- **Total Duration**—Amount of time elapsed between Start Time and End Time.
- **Rows Affected**—The number of rows affected by this request.
- **Max Memory**—Maximum memory utilized by this request, blocks of 2MB are initially reserved and then if additional memory is required 2MB blocks are incrementally assigned.
- **Max Disk**—The maximum amount of memory ever occupied by the request
- **Summary**—The SQL statement or procedure made by this request.

Request details

Every individual request has additional detailed information that might help in troubleshooting failed requests. You can view read-only details for individual requests.

To view the read-only details, click **Show Row Details** 🛈 for the desired row.

In addition to the information presented in the **Requests** table (and described in "The Requests table" on page 65), these details are provided:

- **Request Type**—Either SQL or Procedure.
- **Owner domain**—Name of the domain to which this owner belongs.
- **Session Type**—The type of session: STUDIO, HTTP (Web service), INTERNAL, or client procedure.
- **Transaction ID**—Unique ID for the request's session.
- **Duration**—Amount of time elapsed between Start Time and End Time.
- **Server Duration**—Represents the actual time spent by the server processing this request. The difference between Server Duration and Total Duration is the overhead on the server.
- **Current Memory**—Memory utilization of this request.
- **Current Disk**—The amount of current memory occupied by the request.
- **Description**—a more complete description of the summary.
- **Message**—Displays an error message if the request caused an error.

Sessions

You can access information about the current and recently active sessions. Summary information is displayed at the top of the page, and information about each individual session is displayed in the table below.

You access sessions information by clicking **Sessions** from the **Monitoring** menu.

Note that some of the information displayed on the **Sessions** page is controlled by the configuration settings in IBM Cognos Virtual View Manager. To view these settings, open Virtual View Manager and click **Administration > Configuration**. The following locations also have configuration settings for transactions:

- **Virtual View Manager Server > Configuration > Events and Logging > Event Generation > Session Events**
- **Virtual View Manager Server > Runtime Processing Information > Sessions**

Sessions are removed from the table periodically, based upon the configuration setting **Virtual View Manager Server > Runtime Processing Information > Sessions**. The default setting purges sessions every 30 minutes.

Sessions summary information

You can obtain summary information about sessions.

The **Summary** page provides the following information:
- **Status**—Aggregated status of all sessions can be OK, Warning, Error, or Unknown. A single warning or error supersedes display of a status of OK.
- **Studio Session Timeout**—The amount of time the session can be inactive before it times out.
- **Sessions**—The number of currently active sessions, and the total number of sessions since the server was started.

Working with the Sessions page

If you want to limit the display of sessions to only active sessions, you can quickly remove the completed sessions. You can also terminate a session.

To quickly remove the completed sessions, click the **Purge Completed Sessions** button.

If you want to terminate a session, you can select one or more sessions in the **Sessions** table by check box and then click the **End Sessions** button.

The Sessions table

For each session, the Sessions table provides several columns of information.

The **Sessions** table displays these columns for each session:
- **ID**—Unique session identifier.
- **Status**—Can be any one of the following values:

Status	Description
ACTIVE	The session is currently active.
TIME_OUT	The session has timed out. The Session Timeout configuration setting determines how long the session can remain idle before it times out. See **Virtual View Manager Server > Runtime Processing Information > Sessions > Session Timeout**.
CLOSED	The session is closed.

- **Name**—The session name.
- **Type**—The session type: STUDIO, HTTP (Web service), INTERNAL, or client procedure.
- **Owner**—Userid of the user who initiated the session.
- **Host**—The IP address or name of the host server.
- **Login Time**—Time the user logged in.
- **Idle Duration**—Amount of time the session has been idle.
- **Total Duration**—Amount of time elapsed since the user logged in.
- **Active Requests**—The number of active requests.
- **Total Requests**—Total number of requests processed for this session.
- **Bytes To Client**—The number of bytes in for all requests during this session.

- **Bytes From Client**—The number of bytes sent out for all requests during this session.

Session details

Every individual session has additional detailed information.

To view the read-only details, click the **Show Row Details** button for the desired row.

In addition to the information presented in the **Session** table (and described in "The Sessions table" on page 67), these details are provided:
- **Owner domain**—Name of the domain to which the session owner belongs.
- **Data Service**
- **Logout Time**—The time this session logged out.
- **Timeout Duration**—The amount of time this session can be idle before it will time out.
- **Active Transactions**—The number of currently active transactions.
- **Total Transactions**—Total number of transactions processed in this session.

Transactions

You can access sessions information, which includes the current and recently active transactions.

Access sessions information by choosing **Transactions** from the **Monitoring** menu. The **Transaction** page provides information about the current and recently active transactions. Summary information is displayed at the top of the page, and information about each individual transaction is displayed in the table below.

By default, the **Transactions** page displays transactions that occurred within the last 5 minutes.

Note that some of the information displayed on the **Transactions** page is controlled by the configuration settings in IBM Cognos Virtual View Manager. To view these settings, open Virtual View Manager and click **Administration > Configuration**. The following locations have configuration settings for transactions:
- **Virtual View Manager Server > Configuration > Events and Logging > Event Generation > Transaction Events**
- **Virtual View Manager Server > Runtime Processing Information > Transactions**

Transactions are removed from the table periodically, based upon the configuration setting **Virtual View Manager Server > Runtime Processing Information > Transactions > Transaction Purge Period**.

The default setting purges transactions every 5 minutes.

Transaction summary information

You can obtain summary information about transactions.

The **Transactions** page provides the following summary information:
- **Status**—Aggregated status of all transactions can be OK, Warning, Error, or Unknown. A single warning or error supersedes display of a status of OK.

- **Total Transactions Run**—The total number of transactions processed since the server was started.
- **Total Transactions Rolled Back**—The total number of transactions that were rolled back since the server was started.
- **Total Transactions Failed**—The total number of transactions that failed since the server was started.
- **Active Transactions**—The total number of currently active transactions.

Working with the Transactions page

If you want to limit the display of transactions to only active transactions, you can quickly remove the completed transactions.

Remove completed transactions by clicking the **Purge Completed Transactions** button.

If you want to terminate a transaction, you can select one or more transactions in the **Transactions** table by check box and then click the **Cancel Transactions** button.

The Transactions table

For each transaction the Transactions table displays several columns of information.

The **Transactions** table displays these columns for each transaction:
- **ID**—Unique transaction identifier.
- **Status**—Can be any one of the following values:

Status	Description
START	The transaction started.
COMMITTED	The transaction was committed.
FAIL	The transaction failed.
ROLLBACK	The transaction was rolled back.
COMPENSATE	The transaction was compensated.

- **Mode**—Displays the mode for this transaction: AUTO or EXPLICIT.
- **Owner**—User who initiated this transaction.
- **Session ID**—Unique identifier for the transaction.
- **Session Name**—Name of the component that issued the transaction.
- **Start Time**—Time at which the transaction was initiated.
- **End Time**—Time at which the transaction was completed.
- **Duration**—Amount of time for which the transaction has been running or ran.

Transaction details

Every individual transaction has additional detailed information available that you can view.

To view the read-only details for a transaction, click **Show Row Details** for the desired row.

Triggers

You access information about the current and recently active triggers.

You access trigger information by choosing **Triggers** from the **Monitoring** menu. Summary information is displayed at the top of the **Triggers** page, and information about each individual trigger is displayed in the table below.

By default, the page displays triggers that occurred within the last 5 minutes.

Note that some of the information is controlled by the configuration settings in IBM Cognos Virtual View Manager. To view these settings, open Virtual View Manager and choose **Configuration** from the **Administration** menu. The following locations have configuration settings for triggers:

- **Virtual View Manager Server > Configuration > Events and Logging > Event Generation > Trigger Events**
- **Virtual View Manager Server > Runtime Processing Information > Triggers**

Virtual View Manager provides a TestAllDataSources trigger by default.

Trigger summary information

You can obtain summary information about triggers.

The **Triggers** page provides the following summary information:

- **Status**—Aggregated status of all triggers can be OK, Warning, Error, or Unknown. A single warning or error supersedes display of a status of OK.
- **Total Runs**—Total number of trigger executions carried out since the server started.
- **Total Failed Runs**—Total number of trigger executions that failed since the server started.

Working with the Triggers page

You can change the status of a trigger between enabled and disabled.

If you want to change the status of a trigger between enabled and disabled, you can select one or more triggers in the **Triggers** table by check box and then click the **Change Enabling** button.

The Triggers table

The Trigger table provides information about each trigger.

The **Triggers** table displays these columns for each trigger:

- **Name**—The name of the trigger.
- **Status**—Can be any one of the following values:

Status	Description
ACTIVE	The trigger is processing.
DISABLED	The trigger is disabled.
CONFIG ERROR	The trigger is not configured correctly.

- **Condition**—Type of trigger. Time-event, system-event, or user-defined event.
- **Action**—The type of action this trigger generated.
- **Owner**—User who initiated this trigger.
- **Next Time**—Next time when the execution will occur.
- **Frequency**—Recurrence of execution.

- **Last Time**—Last time the execution occurred
- **Last Success**—Last time the execution was successful
- **Total Attempts**—Total number of times the this trigger was invoked.

Events

You access information about the server events that have been logged for server processes.

You access information about the server events that have been logged by choosing **Event Log** from the **Logging** menu. The **Event Log** page shows information about all of the events generated by processes running on the server. Events originate from one of the following types of objects or processes:

- cached views
- data sources
- requests for query executions
- schedules
- sessions
- transactions

Summary information is displayed at the top of the page, and information about each individual event is displayed in the table below.

Note that some of the information displayed on the **Event Log** page is controlled by the configuration settings in IBM Cognos Virtual View Manager. To view these settings, open Virtual View Manager, choose **Configuration** from the **Administration** menu, and go to **Virtual View Manager Server > Configuration > Events and Logging**.

Server event attributes

IBM Cognos Virtual View Manager creates SNMP events that are compliant with SNMPv1 protocol.

The tables in Chapter 9, "SNMP traps," on page 127, contain details of these server events. The attributes listed are: SNMP ID, short name for the event, and description.

The SNMP traps in the tables have the following OID prefix 1.3.6.1.4.1.18439.2.3.

Set the **generic trap** to 6. Virtual View Manager Server Monitor's specific trap numbers range between 10000 and 19999. Virtual View Manager Server's specific trap numbers range between 20000 and 29999.

In the SNMP details tables, the SNMP ID is listed with the event name, and a description that includes the names of the variables whose values are passed in the message with the text description. The variable values in the description columns are passed in the MIB message payload with an incremented, sequential number that follows the concatenation of the OID prefix and SNMP ID. For example, the variables from the first MIB description are trapTime, trapServerHostName, and trapServerPort.

The SNMP variables are sequentially assigned numbers for display in the MIB message payload. Hence trapTime is assigned a numeric representation of 1,

trapServerHostName is 2, and trapServerPort is 3. All SNMP variables appear in the MIB, identified only with their numeric representation.

Using this example, the line pictured below from a SNMP MIB shows an OID prefix, with a SNMP ID designating a csMonitorStart event with a trapServerPort (represented by the numeral 3) opened on port 9406
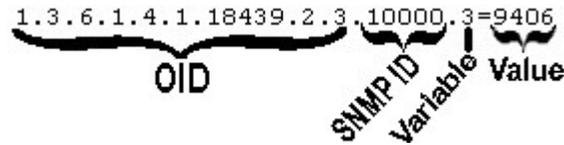
Good SNMP software will parse this message into a human readable format according to the MIB definitions set.

Event log summary information

You can obtain summary information about events.

The **Event Log** page provides the following summary information:
- **Status**—Aggregated status of all events can be OK, Warning, Error, or Unknown. A single warning or error supersedes display of a status of OK.
- **Maximum Viewable Events**—The maximum number of event descriptions maintained in the repository event table. The default maximum is 1000 entries. The number of events is configurable in IBM Cognos Virtual View Manager with the **Administration** menu **Configuration** option by choosing **Virtual View Manager Server > Events and Logging > Logging > Memory > Maximum Viewable Entries** in the **Configuration** window.
- **Maximum Event Entries**—Maximum number of events that can be stored in the server. This number is set in Virtual View Manager with the **Administration** menu **Configuration** option by choosing **Virtual View Manager Server > Events and Logging > Logging > Database Logger > Maximum Log Entries** in the **Configuration** window.

Working with the Event Log page

The Event Log page is mainly an informational page. You can change the sort order, filter the data, or get more details on a specific event, but there are no actions on the data itself.

The Event Log table

The Event Log table provides information about each event.

The **Event Log** table displays these columns for each event:
- **ID**—Unique event identifier.
- **Severity**—The severity level of the event, represented by a colored circle, may be one of DISABLED/OFF (gray circle), INFO (green circle), WARNING (yellow circle), or ERROR (red circle). Refer to Description below.
- **Category**—The type of event, such as REQUEST, SESSION, or TRANSACTION.
- **Type**—The type of event that occurred.
- **Owner**—User who generated this event.
- **Time**—The date and time the event occurred.
- **Description**—A description of the event, such as the request id.

Event and log files

All events in IBM Cognos Virtual View Manager are logged, but not all log entries are tied to system events and visible through the Administrator. Also, there may be cases where an event is associated with multiple log entries.

Virtual View Manager uses a number of log files to store information logged during system and user activities. The sections below describe the installation and uninstallation log files and the server, monitor, and studio log files.

Server, monitor, and studio log files

IBM Cognos Virtual View Manager provides log files for server, monitor, and user activities. These log files can be found in the installation_directory\logs directory.

The following log files are available:

File name	Description
cs_server.out.<TIMESTAMP>	Standard out/error logs
cs_server.log	Main server log.
cs_server_events.log	Server events log.
cs_monitor.out	Server monitor standard out log. Only available on Windows systems.
cs_monitor.out	Server monitor standard out/error log.Only available on UNIX operating systems. UNIX systems have stdout and stderror go to this file, which is not the case for the Windows systems.
cs_monitor.log	Server monitor main log.
cs_monitor_events.log	Server monitor events log.
cs_studio.log	Virtual View Manager application main log.

Chapter 7. Utilities

This chapter describes the command-line utility programs which are available with IBM Cognos Virtual View Manager.

These programs are available in the *installation_location*/bin directory. Generally to use these programs (with the exception of pkg_export, and pkg_import), you need to have administrative rights to Access Tools, Read All Resources, and Modify All Resources.

The following utility programs are described:

backup_export and backup_import

IBM Cognos Virtual View Manager provides two ways to perform a full server backup of an existing server instance. Either the Virtual View Manager interface or the backup_import command line utility may be used to create or restore saved configurations from a full server backup CAR file. By default the backup_export command exports every part of the server configuration, including domains, users, groups, all resources, security settings (ownership of resources and privileges on resources), cache configurations, scheduling, driver configurations, and server-level settings.

Note: For more information on how the Virtual View Manager interface can be used to export and import a full server configuration, see the *IBM Cognos Virtual View Manager User Guide.*

Use the command-line backup programs to backup and save all data source metadata, user-defined resources, (published, shared, and other), and server settings. The file generated by a full server backup export may be used to later restore the entire set of configurations (with the exception of local machine based server configuration settings), data source meta-data, and all derived and child resources.

Local computing environment based settings and configurations are excluded from a full server backup export.

Options may be specified to either include or exclude custom Java and data source statistics regarding cardinality and table boundaries. Other command line utilities, namely pkg_export and pkg_import, give greater flexibility and control over what and how aspects of the server may be exported for purposes other than backup.

The following rights are required to perform a full server backup:
- Access Tools
- Read All Resources

- Read All Users and
- Read All Config

The following rights are required to import an archive:
- Access Tools
- Read and Modify All Resources
- Read and Modify All Users
- Read and Modify All Config

The command line programs are available in the *installation_location* /bin directory:
- backup_export.bat
- backup_import.bat
- backup_export.sh
- backup_import.sh

Backup program commands

The backup program provides several commands and options.

Backup export syntax

This topic provides the syntax for the backup_export command line program.

```
backup_export
 -server <hostname> [-port <port number>] [-encrypt]
 -pkgfile <file name>
 -user <user name> -password <password> [-domain <domain>]
 [-pkgname <name>] [-description <text>]
 [-optfile <file name>] [-excludejars]
 [-includeStatistics] [-verbose]
```

Command line options for backup_export

The backup_export command line program provides several command-line options.

Options in the following table are listed alphabetically. Syntactical order of the options is relevant in most cases.

Notes:
- The export does not include runtime history like log files or probe history.
- The -pkgname and -description flags are optional. They let you include a name and notes within the contents.xml to assist in later identification.

Option	Optional/mandatory	Comments
-description <text>	optional	Description of the exported archive file set as an attribute of the contents.xml file within the exported CAR.
-domain <domain>	optional	User domain. The default value is cognos.

Option	Optional/mandatory	Comments
-encrypt	optional	Encrypts communication between the command line and the server using SSL sent over the dedicated HTTPS port. NOTE: When using the -encrypt option, the HTTPS port is automatically used. Use the -port option to specify the HTTP web services base port and the command line utility automatically switches the SSL connection to the dedicated HTTPS port.
-excludejars	optional	Suppresses export of custom jars. By default. the export does not include binary files such as custom Java procedure JAR files or JDBC driver JAR files.
-includeStatistics	optional	Includes any known resource cardinality statistics about data source table boundaries, column boundaries, and configurations when statistics gathering is enabled.

-optfile <file name>	optional	Specifies an option file providing a way to pass options without using the command line. The options file is useful for scripting and ensuring that sensitive information does not display on the operating system command line.
		
		The contents of the options file are the same strings you would use on the command line, with the addition of new lines being legal.
		
		For example:
		
		`backup_export -server `*`localhost`*
		`-user test -password `*`password`*
		`-pkgfile sample.car`
		
		is the same as:
		
		`backup_export -optfile sample.opt`
		
		where sample.opt contains mandatory arguments and options:
		`-server `*`localhost`*` -user test`
		`-password `*`password`*` -pkgfile`
		`sample.car`

-password <password>	mandatory	Password of the administrative user performing the export.

Option	Optional/mandatory	Comments
-pkgfile <file name>	mandatory	Specifies the path and the file name created as the backup archive file. The file path specified must be accessible to the command line client. The server passes data to the command line client and then the client writes that CAR file to the location specified.
-pkgname <name>	optional	Names an attribute in the contents.xml within the exported backup file. This option is likely to deprecated in future releases.
-port <port_number>	optional	Specifies Web Services Base Port (HTTP) used to communicate with the server. The default value is 9400. Specification is optional is the port setting has not changed. **Note:** Specify the Web Services HTTP port even if the -encrypt option is specified. If -encrypt is specified then the tool will automatically add two to the port number specified here.
-server <hostname>	mandatory	Target server to which the utility will connect.
-user <username>	mandatory	System administrative user name. To properly use backup_export the administrative user must have at least the following rights: Access Tools, Read All Resources, Read All Users, and Read All Config.
-verbose	optional	Generates descriptive output describing the export in the command line window.

Backup import syntax

This topic provides the syntax for the backup_import command line program.

```
backup_import
 -server <hostname> [-port
<port number>] [-encrypt]
 -pkgfile <target/file name>-user
<user name> -password <password>
[-domain <domain>]
-relocate <oldPath>
<newPath> ...
[-optfile <filename>] ...
[-set <path>
```

```
<type>
<attribute>
<value>]
[-printinfo] [-overwrite] [-verbose]
```

Command line options for backup_import

The backup_import command line program provides several command line options.

Option	Optional/mandatory	Comments
`-domain <Domain>`	optional	User domain. The default value is cognos.
`-encrypt`	optional	Encrypts communication between the command line and the server using SSL sent over the dedicated HTTPS port. NOTE: When using the -encrypt option, the HTTPS port is automatically used. Use the -port option to specify the HTTP web services base port and the command line utility automatically switches the SSL connection to the dedicated HTTPS port.

`-optfile <File Name>`	optional	Specifies an option file providing a way to pass options without using the command line. The options file is useful for scripting and ensuring that sensitive information does not display on the operating system command line.
		
		The contents of the options file are the same strings you would use on the command line. Additionally line breaks (new lines) are allowed in the options file, in contrast to the command line which does not normally permit extra line breaks or Enter keystrokes
		
		For example:
		`backup_import`
		`-server localhost -user test -password`
		`password -pkgfile sample.car`
		
		is the same as:
		`backup_import -optfile sample.opt`
		
		where sample.opt contains mandatory arguments and options:
		`-server localhost -user`
		`test`
		`-password password -pkgfile`
		`sample.car`

`-overwrite`	optional	This option does a destructive import overwriting any duplicate resources with the identical resource path.
`-password <Password>`	mandatory	Password of the administrative user performing the import.

Option	Optional/mandatory	Comments
-pkgfile *<FileName>*	mandatory	Specifies the location and file name of the CAR file that is to be read from as the backup archive file. The -pkgfile target file must be a CAR file.
-port *<BasePortNum>*	optional	Specifies Web Services Base Port (HTTP) used to communicate with the server. The default value is 9400. Specification is optional is the port setting has not changed. **Note:** Specify the Web Services HTTP port even if the -encrypt option is specified. If -encrypt is specified then the tool will automatically add two to the port number specified here.
-printinfo	optional	This option disables actual import, and instead allows for inspection of the archive file. PrintInfo displays properties of the CAR file to the command window and exposes some impacts of import with information such as user name, source server with port number, source Java version, user domain, creation date of the archive file, source operating system, and archive type.

-relocate
<old path>
<new path>
	optional	Specify a new location resource name path for top-level items. Specify the old path and the new path using the resource name.
		
		The -relocate option may be used to exclude specified resources from import. Set the -relocate option with the <new path> attribute set to "NOIMPORT" and the resources designated by the <old path> attribute will not be imported.
		
		The following example option stops all resources in the /shared/OldProject container from being imported from the CAR. file.
		
		`-relocate`
		`/shared/OldProject NOIMPORT`

-server *<HostName>*	mandatory	Target server to which the utility will connect to for import of the CAR file defined resources.

Option	Optional/mandatory	Comments

-set
\<path\> \<type\>
\<attribute\>
\<value\>
	optional	Enables resource attribute changes during import.
		
		Use the -set option to change the data source attribute value(s) of classpath, host, port, database, user, or password when deploying a .car file to a new location like a production server.
		
		The resource name is the \<path\>.
		
		The \<type\> is always equal to "DATA_SOURCE" when using the set option with the attributes named above.
		
		The \<attribute\> can be: classpath, host, port, database, user, or password.
		
		Set \<value\> to a valid entry for the selected attribute.
		
		String values may be enclosed with double quotes to allow for spaces in the \<value\>. Example:
		
		`./pkg_import.sh`
		`-pkgfile MyCustJavaDB.car`
		
		`-optfile C:\X.opt -set /shared/datasources/`
		`DBCustom_Java`
		`DATA_SOURCE`
		`classpath "C:\Program Files\JavaDev;`
		`D:\My Documents\Java Joe"`
		
		The set option may be repeated as often as is necessary to set different attributes.
		
		Multiple classpaths can be set with a single statement. When executing backup_import from a Windows system use the semicolon as a delimiter where \<value\> could be:
		
		`C:\DevZone\ATeam\Jars\my.jar;D:\Current`
		`\Ref\classes`
		
		When executing the pkg_import from a UNIX/Linux system use colons as the delimiter. For example:
		
		`/lib/ext/classes:/lib/src/jars`

-user *\<UserName\>*	mandatory	System administrative user name. To properly use backup_import the administrative user must have at least the following rights: Access Tools, Modify All Resources, Modify All Users, and Modify All Config
-verbose	optional	Generates descriptive output describing the export in the command line window. By default the verbose option is not engaged and only error messages are displayed or logged.

Notes on Command line options for backup_import:

- The -pkgfile target file must be a .car file.
- The -verbose option causes informational messages to be displayed after importing. Without this option, only error messages are displayed.
- The -set option allows data source connection information to be explicitly specified or changed from the original during import.

  ```
  [-set<path>
  <attribute> <value>]
  ```

  This option is typically used when deploying a .car file to a production server. The properties most commonly changed are: host, port, database, user, password

Set option for command line options for backup_import:

The following attributes of the -set option may be invoked when importing data sources:

- user <login> or <username> or error depending on source type
- password <password> or error depending on source type
- user2 <appUserName> or error if not Oracle EBS
- password2 <appPassword> or error if not Oracle EBS
- host <urlIP> or <dsn> or <server> or <appServer> or <url> or <root> or error depending on the source type
- port <urlPort> or <port> or error depending on source type
- database <urlDatabaseName> or <enterprise> or <appServer> or error depending on the source type
- path <root> or <url> or error depending on source type
- annotation

Example usage to change the *password* property of a data source:

```
pkg_import mycar.car -set /shared/myDataSource DATA_SOURCE
password myNewPassword
```

In this example, -set is the option; DATA_SOURCE is the type of the resource being imported, and password is the property that is being changed.

Rules for import

Importing follows some rules to resolve conflicts during import.

These rules are:

1. If an imported resource does not exist, it is created. The person performing the import gets appropriate privileges as the creator (such as READ|WRITE for a folder or READ|WRITE|EXECUTE for a procedure). If the user is in the admin group and has requested -includeaccess, then the owner of the resource is set as it was before and any privileges in the import are also put in place.
2. If a resource is imported to a non-existent folder, the folder (and any parent folders of that folder that do not yet exist) is created with the importing user getting READ|WRITE privilege and ownership of these folders. Auto-creation of missing folders is not supported in the *Virtual View Manager Data Services* area.
3. If a resource already exists and it is also being imported, the old version is overwritten (assuming you have the WRITE privilege) in all ways, except the following:

   - The owner is not changed. The original owner retains ownership.

- Privileges for users that are not explicitly changed by the import are left intact. For example, if abe has READ|WRITE and bob has READ|WRITE, and the import lists abe as READ but does not mention bob then abe's privileges are updated but bob's are left intact.
- If the resource is a folder or data source, its children resources are not removed.

4. Restrictions

- The configuration settings (done through the *Administration > Configuration* menu option) are not carried over when you export/import a resource.
- You cannot create a resource in a folder where you do not have the WRITE privilege.
- You cannot overwrite a resource unless you have the WRITE privilege for that resource.
- You cannot export just part of a physical data source. It is either all or nothing.
- You cannot import just part of a physical data source. If you import, you must include the source definition itself.
- The *Virtual View Manager Data Services* area has strict structure rules that are enforced.
- You cannot import anything that was exported from the *Virtual View Manager Data Services* area to outside of that area.
- You cannot import anything that was exported from outside the *Virtual View Manager Data Services* area into that area.

virtualviewmanager

The command-line program virtualviewmanager.bat (for Windows) or virtualviewmanager.sh (for UNIX) can be used to stop or start the server from a command-line interface. Examples are given here using the *.sh command.

Start, stop, or restart the server

The command-line program virtualviewmanager.bat (for Windows) or virtualviewmanager.sh (for UNIX) can be used to stop or start the server from a command-line interface. Examples are given here using the *.sh command.

Procedure

1. Make sure that the monitor is active.
2. Run the following command:

```
<installation_directory>/bin/virtualviewmanager.sh
server <start|stop|restart> -user <username> -password
<password>
```

Start, stop, or restart the monitor

The command-line program virtualviewmanager.bat (for Windows) or virtualviewmanager.sh (for UNIX) can be used to stop or start the monitor from a command-line interface. Examples are given here using the *.sh command.

Procedure

1. Run the following command:

```
<installation_directory>/bin/virtualviewmanager.sh
monitor <start|stop|restart>
```

Starting the monitor automatically starts the server.

Stopping the monitor automatically stops the server.

Restarting the monitor automatically stops and restarts the server.

Note: If you plan to run the monitor process manually and wish to keep the monitor process running even after you log off the system, you need to run the following command:

```
nohup<installation_directory>/bin/virtualviewmanager.sh
monitor <start|restart>
```

2. On Windows, if you plan to run the monitor process manually and wish to keep the monitor process running even after you log off the system, you need to set *Server Ignore Signals* to *True*, as follows:

   - Go to *Administration > Configuration > Virtual View Manager Components > Virtual View Manager Server > Configuration > Monitor > Server Ignore Signals (On Monitor Restart).*

   - Set the *Value* to *True* and the restart the monitor.

Run the server as a foreground process with no monitor

The command-line program virtualviewmanager.bat (for Windows) or virtualviewmanager.sh (for UNIX) can be used to run the server as a foreground process with no monitor from a command-line interface. Examples are given here using the *.sh command.

Procedure

1. Stop the monitor.
2. Run the following commands, as needed:

   ```
   <installation_directory>/bin/virtualviewmanager.sh
   server <run|debug>
   ```

   ```
   <installation_directory>/bin/virtualviewmanager_server.sh
   <run|debug>
   ```

   The commands with the debug option also enable Java debugging on port 8000. All those actions only output to log files. These commands may be used in order to obtain thread dumps or debug a server issue without the monitor getting in the way.

   If you are running the server with no monitor with one of the commands above, a monitor start will kill that server process and restart a new one in the background.

   Note: If you plan to run the server process manually and wish to keep the server process running even after you log off the system, you need to run the following command:

   ```
   nohup <command> &
   ```

   where <command> refers to one of the commands listed above in this section.

3. On Windows, if you plan to run the server process manually and wish to keep the server process running even after you log off the system, you need set *Server Ignore Signals* to *True*, as follows:

   - Go to *Administration > Configuration > Virtual View Manager Information Server Components > Virtual View Manager Server > Configuration > Monitor > Server Ignore Signals (On Monitor Restart).*

   - Set the *Value* to *True* and restart the monitor.

Start, stop, or restart the repository that was installed during the installation

The command-line program virtualviewmanager.bat (for Windows) or virtualviewmanager.sh (for UNIX) can be used to stop or start the repository that was installed during installation from a command-line interface. Examples are given here using the *.sh command.

Procedure

Run the following command:

```
virtualviewmanager.sh
repo <start|stop|restart> -osuser <osusername
> -user <username>
-password <password >
```

where:

- *osusername* is the username of the account that will run the repository process
- *username* is the repository database user name
- *password* is the repository database user password

This command is only supported for a repository that was installed during the installation of Virtual View Manager.
On Windows, the usage is as follows:

```
virtualviewmanager.bat repo <start|stop|restart|install|uninstall>
```

Stopping and starting the server on Windows startup program

On Windows, you can manually start or stop the server through the startup program without a command-line interface.

- To stop Virtual View Manager server manually
  - From the *Start* menu, click IBM Cognos Virtual View Manager, *Server*, *Stop Virtual View Manager Server*.
- To start Virtual View Manager server manually
  - From the *Start* menu, click IBM Cognos Virtual View Manager, *Server Start Virtual View Manager Server*.

install_services

The install_services.sh script can be used from a command-line interface to automatically re-start the server and the repository on UNIX.

If at anytime after installing the software you re-start the installation machine, IBM Cognos Virtual View Manager and the metadata repository will not start automatically. In order for the installation machine to automatically start or stop the server during a re-start, run the script named install_services.sh, which would install two service files (csw.repository and csw.server) and configure them.

Running install_services.sh does not interrupt any repository or server processes that are currently running.

Procedure

1. Log into the server computer with root privileges.
2. Change to the *installation_directory* /bin directory.
3. Run the following command:

```
install_services.sh
```
You are prompted for a user name and other details to configure the service files.

pkg_export and pkg_import

Use the pkg_export and pkg_import command line utilities to export and import specified resources from the server.

These command line utilities provide two ways of exporting and importing specified resources from the server:

- Export to and import from a CAR file (-pkgfile option)
- Export to and import from a hierarchical directory (-pkgdir option)

Selected directories or individual resources may be exported to a single discrete CAR file or they may be exported to a hierarchical directory that mirrors the IBM Cognos Virtual View Manager resource tree directory structure.

Advantages of exporting and importing to a single CAR file include:

- Aggregates all metadata files and objects into a single file
- Multiple resource files are packaged in a single compressed archive.
- All resource and metadata objects may be exported/imported.

Disadvantages of CAR files include:

- No incremental update available. Any changes to a resource backed up to a CAR file requires re-creation of the CAR file to capture modifications.
- Exported CAR file objects are in a proprietary format that is not easily modified.

Advantages of exporting and importing to a directory structure include:

- All incremental changes (additions, modifications, and deletions) made on the server may be captured with a single pkg_export execution that updates only those files affected by the changes.
- Exported files are stored in a hierarchical directory structure that mimics the structure created, using the server.
- SQL and metadata files are stored in file formats that may be modified.
- Directory and file updates compatible with many source code control systems

Disadvantages of directory package exports:

- JARs, users, groups, domains, and rights in the directory export are not supported.

The utilities, pkg_export and pkg_import, are available in the *installation_location*/ bin directory.

On Windows, the utilities are named pkg_export.bat and pkg_import.bat. On UNIX, they are named pkg_export.sh and pkg_import.sh.

Note: The Virtual View Manager menu option to export selected resources to CAR files is equivalent to using the pkg_export with the -pkgfile option. For details on using the Virtual View Manager menu option, see the Resource Management Basics chapter in the Virtual View Manager *User Guide*.

Package export command line utility

This utility exports specified directories and resources. The package export command-line utility can be thought of as if it were two different utilities because some of the options that are available to each of the two primary types (CAR file -pkgfile and directory -pkgdir) of export are not interoperable.

Two sections follow to describe the various mandatory switches and options available for use with each option of the package export command-line utility:

Package export to a CAR file (-pkgfile option)

For export of specified resources to a CAR file:

```
pkg_export -pkgfile <FileName>
 <NamespacePath> [...]
 -server <hostname> [-port <port>]
[-encrypt]
 -user <user> -password <password> [-domain <domain>]
 [-pkgname <name>] [-description <text>]
 [-optfile <filename>]
[...]
 [-rebindable <path>
<description>]
...
 [-includeaccess] [-includecaching]
 [-includesourceinfo] [-includejars]
 [-includeAllUsers] [-includeUser <domain>
<user>]
...
 [-includeGroup <domain>
<group>]
...
 [-includeDomain <domain>]
...
 [-includeRequiredUsers] [-includeDependencies]
 [-includeStatistics]
 [-verbose] [-quiet]
```

Examples:

In the example above, myParameterizedQuery is exported with dependencies that in this case includes the products table from the orders data source.

```
pkg_export -pkgfile MyExport.car
shared/procedures/myParameterizedQuery
-server localhost
-user admin -password AdminPassword
 -includeDependencies
-rebindable shared/sources/ds_orders/products "This needs
rebinding to the production data source."
```

The -rebindable option is specified to notify or remind the user during an import that the products resource will need to be rebound. During import a message prompt and the <description> get displayed within the command prompt and they also appear when the resulting export CAR file is imported in Virtual View Manager. Rebinding must be done manually after the import unless the -rebind option specifies the new resource for rebinding during import.

Another example:

```
pkg_export -pkgfile Sources_Backup.car
shared/sources
-optfile C:/BackupScripts/Sources/weekly.opt
 -includeDependencies
 -includesourceinfo
```

In this example shared/sources is backed up to Sources_Backup.car.

The option file, weekly.opt, must contain any mandatory arguments that were missing in the original command:

```
-server localhost
-user DBASecure1
-password Password
-domain EnterpriseLDAP
```

An options file could be dynamically be generated to specify options and arguments including the user name, password, and domain so that programmatic backups may be set while obfuscating the DBA login from display on either the application window or in the file prior to running the scheduled script.

The following table of pkg_export -pkgfile parameters and arguments is organized by the order of appearance.

pkg_export -pkgfile parameters and args	Optional/mandatory	Comments
-pkgfile *<FileName>*	mandatory	Specifies new CAR file name.
-server *<HostName>*	mandatory	Server host to which the utility will connect.
-port *<BasePortNum>*	optional	Specifies Web Services Base Port (HTTP) used to communicate with the server. The default value is 9400. Specification is optional if the port setting has not changed from the default. **Note:** Specify the Web Services HTTP port even if the -encrypt option is specified. If -encrypt is specified then the tool will automatically add two to the port number specified here.
-encrypt	optional	Encrypts communication between the command line and the server using SSL sent over the dedicated HTTPS port. **Note:** When using the -encrypt option, the HTTPS port is automatically used. Use the -port option to specify the HTTP web services base port and the command line utility automatically switches the SSL connection to the dedicated HTTPS port.

pkg_export -pkgfile parameters and args	Optional/mandatory	Comments
-user *\<UserName\>*	mandatory	Virtual View Manager system administrative user name. To properly use pkg_export the user must have ownership of the specified resource or at least read privilege on all the specified resources with: Access Tools and Read All Config rights
-password *\<Password\>*	mandatory	Password for user profile used to export package.
-domain *\<Domain\>*	optional	User domain. The default value is cognos and does not have to be stated unless the user profile used for export has another value.
From this row down the Parameters are listed alphabetically.		
-description "*\<Text\>*"	optional	Package description of the archive file. Notation appears when Virtual View Manager is used to import the CAR.
-includeAllUsers	optional	Exports all domains, groups, and users to the export file. Requires the Read All Users right.
-includeUser *\<DomainName\>* *\<UserName\>*	optional	Includes the specified user in the export file. This option may be repeated to export multiple users. Repeat the option keyword -includeUser with arguments for the new domain and user as many times as is necessary.
-includeGroup *\<DomainName\>* *\<GroupName\>*	optional	Exports group information about the specified group in the export file.
-includeDomain *\<DomainName\>*	optional	Export the specified domain metadata to the export file.
-includeaccess	optional	Includes the current user access controls (privilege specifications) on the resources in the export file. Default setting does not include access control even when exported by an administrator. An error may occur if this option is used and the exporting user is not a member of the admin group.
-includecaching	optional	Includes the details of caching on views and procedures in the export file.

pkg_export -pkgfile parameters and args	Optional/mandatory	Comments
-includeDependencies	optional	Gathers and includes all dependent resources for the resources you choose to export.
-includejars	optional	Includes the .JAR files in the export file
-includeRequiredUsers	optional	Includes the information about the required users in the export file.
-includesourceinfo	optional	Data source connection details such as user name, password, host name, and port are not included by default. Specify -includesourceinfo to include these.
-includeStatistics	optional	Includes any resource stats known about the table boundaries, column boundaries, etc.

pkg_export -pkgfile parameters and args	Optional/mandatory	Comments

`-optfile`
`<FileName>`
	optional	Options file.
		
		The options file feature provides a way to pass options without using the command line. Use the options file for scripting and ensuring that sensitive information is not displayed on the operating system command line.
		
		The contents of the options file are the same strings you would use on the command line. Unlike most command prompt windows, new lines in the text file are legal.
		
		For example,
		`backup_export`
		`-server localhost`
		`-user test -password password`
		`-pkgfile sample.car`
		
		is the same as:
		`backup_export -optfile sample.opt`
		
		where sample.opt contains either:
		`-server localhost`
		`-user test`
		`-password password`
		`-pkgfile sample.car`
		
		or (this one with new lines):
		`-server localhost`
		`-user test`
		`-password password`
		`-pkgfile sample.car`
		
		This option is also available for use in pkg_import.

`-pkgname` `"<Name>"`	optional	Package name may be set, but this string is not currently used. Spaces are allowed but punctuation is not.

`-rebindable`
`<Path>`
`<Description> ...`
	optional	Marks a resource dependency for rebinding on import. When a rebindable resource is imported a reminder and the <Description> are displayed on the command line. The -rebind option must be specified on import for that action to take place on import. That message is also displayed in Virtual View Manager to prompt designation of a new resource (path) as the resource dependency.

pkg_export -pkgfile parameters and args	Optional/mandatory	Comments
-verbose	optional	This option causes verbose reporting of information much like the GUI provides instead of the normal terse reporting.

Note: The pkg_export command exports each of the listed resources (as specified using a namespace path such as /users/manager/sources or /shared/projects/pegasus). It does not include any domains, users, groups, or server settings.

- This command can be used by any user. If you do not have the READ privilege to any of the specified resources, the export will fail. If you do have the READ privilege for all of them, then all resources that are children of those resources will also be included as long as you have the READ privilege on those children.

  Users who attempt export without the READ privilege on a child resource will not be able to export the child resource, and they will not be notified of the omission of the child from the export package.

- By default, access information (privilege settings) are not included in the export. The -includeaccess option must be provided if you want privilege settings included in the archive file. The -includeaccess option is ignored if you are not logged in as an administrator with the Read All Resources right.

- By default, no caching configurations are included in the export. The -includecaching option must be specified to include cached data from materialized views, configurations that include scheduling for cache refreshes.

- By default, physical data source connection details such as usernames, passwords, host names, and ports are blanked out during an export. The -includesourceinfo option must be provided if you want those details included. Note that if passwords are included, they are encrypted.

- The -pkgname and -description flags are optional. They let you include your specific information in the archive file for later identification.

- The -verbose option reports if problems are encountered during the export

Package export to a directory (-pkgdir option)

The -pkgdir option exports resources to a directory for incremental update. The resource definitions and file objects are externalized from the server and they may be modified given certain precautions.

Since privilege settings and resource metadata are exported as files that may be modified, they must be protected from intentional or inadvertent changes.

One of the primary advantages of externalizing resources with the -pkgdir option is that exports may be quickly performed that incrementally update only those files that have been created, changed, or deleted. This is great for source control systems that safeguard files and track version changes.

- Exporting to a new directory

  When a package export is made to a new directory all defined resources in the specified directories are exported to a hierarchical directory structure matching the Virtual View Manager resource tree.

- Exporting to an existing externalized directory

  To incrementally update an existing externalized directory use the pkg_export -pkgdir command and exclude the resource namespace path. By default, the package export will incrementally update only those files that have changed

since the last export. Make sure to specify the existing -pkgdir
<existing_externalized_directory> without specifying the resource namespace
path to engage.

When exporting to an existing externalized directory, the resource namespace
path should not be specified as the directory already has a manifest file that
specifies what resources will be updated in the export directory. During the first
-pkgdir export, the manifest file is initialized and set with a specific set of
resources and/or directories that will be updated in that -pkgdir directory. The
manifest file will always define those resources that will be stored in that
particular physical location. The manifest file should not be modified under
normal conditions, but the timestamp may be changed to enable an export of
resources changed after a particular time.

Once defined, the manifest file is inspected and only those objects that have
changed will be exported and saved to the external directory.

- Exporting all with the -full option

  Use the -full option to export everything (all resources changed or not) in the
  resource namespace path. The -full option forces all defined resources named by
  the root directory in the manifest file to be updated. All resources are exported,
  overwriting everything in the specified path ensuring that the externalized files
  and the server are synchronized.

Note: When a lot of resources must be exported at once (or when exporting the
entire directory) on a Windows server under high load conditions, one must first
modify the Windows server to allow more sockets per user and reduce the TCP
timed wait delay. Windows can safely accommodate 65534 sockets per user in the
TCP/IP scheme and TCP timed wait can be decreased to 30 seconds to release
socket resources more quickly.

By default, most Windows operating systems do not have the following keys so
they must be created to tune the system for high-volume package exports.

Navigate to key HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\
Tcpip\Parameters

And create the following two keys:
- DWORD Value Name="MaxUserPort", Value=65534 (Decimal)
- DWORD Value Name="TcpTimedWaitDelay", Value=30 (Decimal)

Once you've set the parameters, restart the computer for them to take affect.

Note: Limitation - Externalization (package export to a directory) does not support
export and import of JARs, users, groups, domains, and rights in the packaged
directory export.

Usage of package export to a directory (-pkgdir option):

Windows and UNIX specific versions of the pkg_export utility are available.

As stated previously, the pkg_export utility is available in the *installation_directory*
/bin directory.

```
pkg_export
-pkgdir <directory> <NamespacePath>
[...]
 -server <hostname> [-port
<port>] [-encrypt]
```

```
 -user <user> -password <password> [-domain
<domain>]
 [-pkgname <name>] [-description
<text>]
 [-optfile <filename>] ...
 [-full]
 [-messagesOnly]
 [-fileEncoding <encoding>]
 [-listIgnored]
 [-verbose] [-quiet]
```

Examples:

The default administrative user profile is used to execute the following command which externalizes all resources defined by the server instance on the localhost in the resource path /shared/examples/ds_orders. All the metadata object definitions are exported as XML files to the C:/temp directory.

```
pkg_export.sh -pkgdir C:/temp /shared/examples/ds_orders
-server localhost -user admin -password admin
```

The above example establishes a manifest file in C:/temp that will reserve the directory for sync of all resources in /shared/examples/ds_orders path.

Later, to see what files have been changed and what would be exported the -messagesOnly option may be specified so that nothing is actually exported on execution of the command:

```
pkg_export.sh -pkgdir C:/temp -server localhost -user
admin -password admin -messagesOnly
```

To incrementally synchronize that externalized directory execute this command:

```
pkg_export.sh -pkgdir C:/temp -server localhost -user
admin -password admin
```

Use the -encrypt option to securely backup a directory to some local or locally mapped drive from across an unsecured network. For example,

```
pkg_export.sh -pkgdir Z:/Secure /shared/examples/ds_orders
-server 111.123.456.99 -port 2220 -encrypt -user admin999 -password
Secured999 -domain SomeLDAPDomain
```

With the above example command, an SSL/HTTPS connection with a Virtual View Manager Server at the specified IP address is established. The specified port is the designated Virtual View Manager Server HTTP base port, but the -encrypt option changes the actual connection to use the HTTPS port. The encrypted connection enables a secured login authentication. The Virtual View Manager Server exports the requested XML files securely over the SSL connection to the client application and the client writes those files to the designated target directory Z:/Secure.

The following table of pkg_export -pkgfile parameters and arguments is organized by the both the order of appearance and alphabetically.

pkg_export -pkgdir parameters and args	Optional/mandatory	Comments
-pkgdir *<Directory>*	mandatory	Specifies the absolute location for creation of a directory and placement of the export files. Example: -pkgdir C:/SourceControl/ProjectI Mapped directories must be available to the user on the command line client
<NamespacePath>	mandatory	Uses the Virtual View Manager resource name convention to specify one or many resource directories for export. Example: /shared/examples/ds_orders All resources within the specified Virtual View Manager directory are exported to -pkgdir C:/Directory/*specified*
-server *<HostName>*	mandatory	Virtual View Manager Server host to which the utility will connect. Default name is localhost.
-port *<BasePortNum>*	optional	Specifies Web Services Base Port (HTTP) used to communicate with the Virtual View Manager Server. The default value is 9400. Specification is optional if the port setting has not changed. **Note:** Specify the Web Services HTTP port even if the -encrypt option is specified. If -encrypt is specified then the tool will automatically add two to the port number specified here.
-encrypt	optional	Encrypts communication between the command line and the server using SSL sent over the dedicated HTTPS port. **Note:** When using the -encrypt option, the HTTPS port is automatically used. Use the -port option to specify the HTTP web services base port and the command line utility automatically switches the SSL connection to the dedicated HTTPS port.
-user *<UserName>*	mandatory	User name for authorization of export. User must have Access Tools and ownership of the resources to be exported or the Read All Resources rights for proper export of all configurations.
-password *<Password>*	mandatory	Password for user profile used to export package.
-domain *<DomainName>*	optional	User domain. The default value is cognos and does not have to be stated.

pkg_export -pkgdir parameters and args	Optional/mandatory	Comments
-description *<text>*	optional	Textual description of the archive directory. Visible when selected for import with the Virtual View Manager interface.

-fileEncoding
 <encoding>
	optional	Encoding defaults to UTF-8, but a different encoding option may be specified. Display a list of valid encoding by entering an invalid encoding value, or try to read the 8 point below. Valid encoding values include the following:
		
		`Big5,Big5-HKSCS,Cp1252, EUC-JP,`
		`EUC-KR, GB18030, GB2312, GBK,`
		`IBM-Thai, IBM00858,IBM01140,`
		`IBM01141, IBM01142, IBM01143,`
		`IBM01144, IBM01145, IBM01146,`
		`IBM01147, IBM01148, IBM01149,`
		`IBM037, IBM1026, IBM1047, IBM273,`
		`IBM277, IBM278,IBM280, IBM284,`
		`IBM285, IBM297, IBM420, IBM424,`
		`IBM437, IBM500, IBM775, IBM850,`
		`IBM852, IBM855, IBM857, IBM860,`
		`IBM861, IBM862, IBM863, IBM864,`
		`IBM865, IBM866, IBM868, IBM869,`
		`IBM870, IBM871, IBM918,`
		`ISO-2022-CN, ISO-2022-JP,`
		`ISO-2022-KR, ISO-8859-1,`
		`ISO-8859-13, ISO-8859-15,`
		`ISO-8859-2, ISO-8859-3,`
		`ISO-8859-4, ISO-8859-5,`
		`ISO-8859-6, ISO-8859-7,`
		`ISO-8859-8, ISO-8859-9,`
		`JIS_X0201, JIS_X0212-1990,`
		`KOI8-R, Shift_JIS, TIS-620,`
		`US-ASCII, UTF-16, UTF-16BE,`
		`UTF-16LE, UTF-8, windows-1250,`
		`windows-1251, windows-1252,`
		`windows-1253, windows-1254,`
		`windows-1255, windows-1256,`
		`windows-1257, windows-1258,`

pkg_export -pkgdir parameters and args	Optional/mandatory	Comments
		windows-31j, x-Big5-Solaris, x-euc-jp-linux, x-EUC-TW, x-eucJP-Open, x-IBM1006, x-IBM1025, x-IBM1046, x-IBM1097, x-IBM1098,x-IBM1112, x-IBM1122, x-IBM1123, x-IBM1124, x-IBM1381, x-IBM1383, x-IBM33722, x-IBM737, x-IBM834, x-IBM856, x-IBM874, x-IBM875, x-IBM921, x-IBM922, x-IBM930, x-IBM933, x-IBM935, x-IBM937, x-IBM939, x-IBM942, x-IBM942C, x-IBM943, x-IBM943C, x-IBM948, x-IBM949, x-IBM949C, x-IBM950, x-IBM964, x-IBM970, x-ISCII91, x-ISO-2022-CN-CNS, x-ISO-2022-CN-GB, x-iso-8859-11, x-JIS0208, x-JISAutoDetect, x-Johab, x-MacArabic, x-MacCentralEurope, x-MacCroatian, x-MacCyrillic, x-MacDingbat, x-MacGreek, x-MacHebrew, x-MacIceland, x-MacRoman, x-MacRomania, x-MacSymbol, x-MacThai, x-MacTurkish, x-MacUkraine, x-MS950-HKSCS, x-mswin-936, x-PCK, x-windows-50220, x-windows-50221, x-windows-874, x-windows-949, x-windows-950, x-windows-iso2022jp
-listIgnored	optional	Lists those items that are not exported
-messagesOnly	optional	Preview an export using this option. No files are created, modified, or deleted but instead command line messages will list those files that will be added, modified, and deleted.

pkg_export -pkgdir parameters and args	Optional/mandatory	Comments

`-optfile` `<filename>`	optional	Options file.
		
		The options file feature provides a way to pass options using a file instead of the command line. The options file is useful for scripting and ensuring that sensitive information does not display on the operating system command line.
		
		The contents of the options file are the same strings you would use on the command line, with the addition of new lines being legal.
		
		For example,
		`pkg_export -pkgdir`
		`-server localhost -user`
		`MyAdmin -password`
		`CleartextPassword -pkgfile`
		`Backup.car`
		
		is the same as:
		`backup_export -optfile sample.opt`
		
		where sample.opt contains either:
		`-server localhost`
		`-user MyAdmin`
		`-password CleartextPassword`
		`-pkgfile Backup.car`
		`or (this one with new lines):`
		`-server localhost`
		`-user test`
		`-password password`
		`-pkgfile Backup.car`
		
		This option is also available for use in pkg_import.

`-pkgname` `<name>`	optional	Package name may be set to report a name attribute in the manifest file. Otherwise this field is not used. Spaces are allowed but punctuation is not.
`-quiet`	optional	Suppresses messaging in the command line.
`-rebindable` `<Path>` `<Description> ...`	optional	Marks a resource dependency for rebinding on import. When a rebindable resource is imported using Virtual View Manager import the <Description> is displayed in a dialog that enables selection of a new resource (path) as the resource dependency.
`-verbose`	optional	Generates descriptive output describing the export in the command line window.

Package import command line utility

The package import (pkg_import) utility imports directories and resources from either a zipped archive (CAR) file or from an externalized directory of Virtual View Manager resources rendered as XML files.

The pkg_import utility is located in the *installation_directory* /bin directory.

The package import command-line utility can be thought of as if it were two different utilities because some options for import from a CAR file -pkgfile are different from the options available for import from a directory using -pkgdir.

Use of the package import (pkg_import) command-line utility requires the following rights to properly restore defined resources: Access Tools, Read and Modify All Config, Read and Modify All Resources, Read and Modify All Users. Import generally requires the most administrative rights as it will overwrite many types of resource objects, directories, user privileges, and other resource definitions.

The package or directory being imported and the options employed to import them will dictate what rights are actually needed to import, but by nature of the ability to overwrite resource definitions, import generally requires full admin rights.

Two sections follow to describe the various mandatory switches and options available for use with each option of the package export command-line utility:

Package import from a CAR file (-pkgfile option)

Package import from a CAR file (-pkgfile) has the following generic form:

```
pkg_import -pkgfile <D:/directory/Path/and/File_Name.car>...
-server <HostName> [-port
<PortNumber>] [-encrypt]
-user <UserName> -password
<Password> [-domain <Domain>]
[-relocate <OldPath>
<NewPath>]
...
[-rebind <OldPath>
<NewPath>]
...
[-set <path>
<DataType>
<Attribute>
<Value>]
...
[-optfile <D:/directory/Path/and/File_Name>]
[-printinfo] [-printroots] [-printusers]
[-printcontents] [-printreferences]
[-includeaccess] [-nocaching]
[-excludejars] [-nosourceinfo]
[-overwrite] [-overrideLocks] [-messagesonly]
[-verbose] [-quiet]
```

Examples

In this example, set is the option; DATA_SOURCE is the type of the resource being imported, and password is the attribute being changed:

```
pkg_import.sh -pkgfile C:/Store/EnterpriseArchive/mycar.car
-server localhost -user ProdAdmin -password AdminPassWeird -set
/shared/myDataSource DATA_SOURCE password MyNewPassword
```

In the next example, only shared/procedures/myParameterizedQuery is imported with its associated dependencies from the specified CAR file. The relocate option is used to move the query to a new directory.

```
pkg_import -pkgfile Z:/Archive/QA_Image99999.car

shared/procedures/myParameterizedQuery -server localhost
-user admin -password AdminPassword -includeDependencies-relocate
shared/procedures/myParameterizedQuery
shared/RestrictedUse_Procedures/myParameterizedQuery
```

Package import (-pkgfile) mandatory switches and options

This table describes the mandatory switches and options for pkg_import -pkgfile.

pkg_import -pkgfile parameters and args	Optional/mandatory	Comments
-pkgfile <FileName>	mandatory	Specifies the import CAR file. The file name should be specified as an absolute path from a mapped directory. Multiple CAR files may be specified.
-server <HostName>	mandatory	Virtual View Manager server host to which the utility will connect.
-port <BasePortNum>	optional	Specifies Web Services Base Port (HTTP) used to communicate with the server. The default value is 9400. Specification is optional if the port setting has not changed. Note: Specify the Web Services HTTP port even if the -encrypt option is specified. If -encrypt is specified then the tool will automatically add two to the port number specified here.
-encrypt	optional	Encrypts communication between the command line and the server using SSL sent over the dedicated HTTPS port. **Note:** When using the -encrypt option, the HTTPS port is automatically used. Use the -port option to specify the HTTP web services base port and the command line utility automatically switches the SSL connection to the dedicated HTTPS port.
-user <UserName>	mandatory	User name of profile used to import. User rights specified by the target server instance grant permission to import and may restrict write privilege to import into designated directories only.
-password <Password>	mandatory	Password for user profile used to export package.

pkg_import -pkgfile parameters and args	Optional/mandatory	Comments
`-domain <domain>`	optional	Enter domain for the user performing the import. If it is omitted the assumed value is cognos.

`-optfile <FileName>`	optional	Options file.
		
		The options file feature provides a way to pass options using a file instead of the command line. The options file is useful for scripting and ensuring that sensitive information does not display on the operating system command line.
		
		The contents of the options file are the same strings you would use on the command line, with the addition of new lines being legal.
		
		For example,
		`backup_export`
		`-server localhost -user test`
		`-password password -pkgfile`
		`sample.car`
		
		is the same as:
		`backup_export -optfile sample.opt`
		
		where sample.opt contains either:
		`-server localhost -user`
		`test -password password -pkgfile`
		`sample.car`
		
		or (this one with new lines):
		`-server localhost`
		`-user test`
		`-password password`
		`-pkgfile sample.car`
		
		This option is also available for use in pkg_export.

`-relocate <OldPath> <NewPath>`	optional	This option can be used multiple times. It is used to specify a new location for a top-level item. An error will occur if you try to relocate something that is not in the package or is not a top-level item.

pkg_import -pkgfile parameters and args	Optional/mandatory	Comments
-rebind *\<OldPath\>* *\<NewPath\>*	optional	Sets a new resource path for a dependency resource. The option may be repeated for as many dependencies as may be necessary. Migration from a development or testing environment to a production server deployment makes use of the -rebind option.

-set *\<path\>* *\<type\>* *\<attribute\>* *\<value\>*	optional	Changes the data source attribute value(s) of classpath, host, port, database, user, or password when deploying a CAR file to a new location like a production server.
		
		The Virtual View Manager resource name is the \<path\>.
		
		The \<type\> is always equal to DATA_SOURCE when using the set option with the attributes named above.
		
		The \<attribute\> can be: classpath, host, port, database, user, or password.
		
		Set \<value\> to a valid entry for the selected attribute.
		
		String values may be enclosed with double quotes to allow for spaces in the \<value\>. Example:
		
		`./pkg_import.sh`
		`-pkgfile MyCustJavaDB.car -optfile`
		`C:\X.opt -set /shared/datasources/`
		`DBCustom_Java`
		`DATA_SOURCE classpath`
		`"C:\Program Files\JavaDev;`
		`D:\My Documents\Java`
		`Joe"`
		
		The set option may be repeated to set different attributes.
		
		Multiple classpaths can be set with a single statement. When executing pkg_import from a Windows system use the semicolon as a delimiter where \<value\> could be something like:
		
		`C:\DevZone\ATeam\Jars\my.jar;`
		`D:\Current\Ref\classes`
		
		When executing the pkg_import from a UNIX/Linux system use colons as the delimiter. For example:
		
		`/lib/ext/classes:/lib/src/jars`

pkg_import -pkgfile parameters and args	Optional/mandatory	Comments
-printinfo	optional	Prints information such as user name, source server with port number, source Java version, user domain, creation date of the archive file, source operating system, and archive type.
-printroots	optional	Prints the new path(s) to the imported resources.
-printusers	optional	Prints the user names of the owners of the imported resource(s) and their associated user groups.
-includeusers	optional	Imports all users present in the exported package car. By default domain, groups and user information are not included in either export or import. This option is also available for use in pkg_export.
-printcontents	optional	This option disables actual import, and instead allows for inspection of the archive file. Print contents displays properties of the CAR file to the command window and exposes some impacts of import with information such as user name, source server with port number, source Java version, user domain, creation date of the archive file, source operating system, and archive type.
-printreferences	optional	Prints a list of resources referred to by the imported resource(s).
-includeaccess	optional	Includes the current user access controls on the resources in the export file.

It defaults to not including access control even if you are an administrator. You get an error if you use this flag and are not logged in as a member of the admin group.

This option is also available for use in pkg_export. |
| -nocaching | optional | Omits caching data and all caching configurations as if the caching of the exported resource were never set. |
| -excludejars | optional | Does not import the jar files. |

pkg_import -pkgfile parameters and args	Optional/mandatory	Comments
-nosourceinfo	optional	Suppresses import of the following pre-existing connection attributes when an otherwise identical resource is already present at the target: Driver, ConnectionURL, Port, Database Name, Login, Password, and Pass-through Login. Supports re-import without need for explicit "-set" options and without altering original data source attributes.
-overwrite	optional	This option overwrites duplicated resources.
-messagesonly	optional	Displays the messages generated in a package import without actually performing the import.
-port <PortNumber>	optional	Port for the target server instance. Default=9400
-verbose	optional	This option causes verbose reporting of information much like the GUI provides instead of the normal terse reporting.

- The pkg_import command imports the resources in the archive file into the server following the import rules (). If this command is used on an archive created using *backup_export* , only the resource information will be used.
- By default, access information (ownership and privileges) are not imported. Any resources that are created as a result of the import are owned by you. Existing resources follow the import rules (). The -includeaccess option must be used if you want to preserve ownership and privilege information. This option is ignored if you are not logged in as a member of the admin group.
- By default, caching configuration is imported. The -nocaching option must be used if you want to ignore cache configurations.
- By default, scheduling configuration is imported. The -noscheduling option must be used if you want to ignore scheduling configurations. Note that if you are logged in as a member of the admin group, then all scheduling is imported. If not, only the schedules owned by you are imported.
- By default, physical data source connection details such as usernames, passwords, host names, and ports are imported. The -nosourceinfo option must be used if you want to ignore those details. When ignoring connection details and importing over an existing physical data source, the previous values will be left unmodified instead of being overwritten. When creating a new data source, these values are set to an empty string.
- By default, all resources import into the same location in the namespace they were in when exported. The -relocate option can be used to change the location for the import. You can issue as many -relocate flags as desired and they are processed in order, with the first matching oldPath being applied to each

resource. Relocating a resource will modify (rebind) references made by other resources being imported, but will not modify references on resources that are not part of the import.

- As an additional option, resources can be rebound as a group during the import process using the -rebind option. You can issue as many -rebind flags as desired and they are processed in order. All imported resources will have the rebind performed. In case of conflict between rebinds caused by the -relocate and -rebind flags, the -rebind ones are performed first.

- The -verbose option causes informational messages to be printed after importing. Without this option, only error messages are displayed.

- The -printinfo option causes the archive file to be examined and for information about it to be printed to the display. The archive file is not imported when this option is given.

- The -printcontents option causes the archive file to be examined and for the list of resources in the file to be printed to the display. The archive file is not imported when this option is given.

- The -printroots option causes the archive file to be examined and for the list of top level resources (generally the list provided to pkg_export on the command line) to be printed to the display. This is useful for understanding what is in the archive without having to see the complete list. The archive file is not imported when this option is given.

- The -printreferences option causes the archive file to be examined and for the list of resources that are referenced within the import but are not actually part of the import to be printed to the display. This is useful for identifying external dependencies for the archive. The archive file is not imported when this option is given.

- The -printusers option causes the archive file to be examined and for the list of users that are referenced as owners or in privileges on the resources to be printed to the display. This is useful for identifying the users required to import with access information intact. The archive file is not imported when this option is given.

- The -messagesonly option causes the archive file to generate all the messages it would generate if you actually imported it, but does not commit the import to the server's storage.

- The -set option allows data source connection information to be explicitly specified or changed from the original during import.

  `[-set<path> <attribute> <value>]`

  This option is typically used when deploying a CAR file to a production server. The properties most commonly changed are host, port, database, user, password.

  The following attributes of the `-set` option may be invoked when importing data sources:

  - `user` <login> or <username> or error depending on source type
  - `password` <password> or error depending on source type
  - `user2` <appUserName> or error if not Oracle EBS
  - `password2` <appPassword> or error if not Oracle EBS
  - `host` <urlIP> or <dsn> or <server> or <appServer> or <url> or <root> or error depending on the source type
  - `port` <urlPort> or <port> or error depending on source type
  - `database` <urlDatabaseName> or <enterprise> or <appServer> or error depending on the source type

- path <root> or <url> or error depending on source type
- annotation

Example usage to change the password property of a data source:

```
pkg_import mycar.car -set /shared/myDataSource DATA_SOURCE
password myNewPassword
```

In this example, -set is the option; DATA_SOURCE is the type of the resource being imported, and password is the property that is being changed.

Rules for import

Importing follows some rules to resolve conflicts during import.

These rules are:

1. If an imported resource does not exist prior to import, then it is created. The user performing the import gets all privileges as if they were the original creator (such as READ|WRITE for a folder or READ|WRITE|EXECUTE for a procedure) unless the -includeaccess option is specified.

   If an administrative user who has the Modify All Users right imports a resource using the -includeaccess option, then the original owner of the resource is set as the owner and any pre-existing privileges in the import package are also set for the newly imported resource.

2. If a resource is imported to a non-existent location, the new directory path and folder are created with the import of the new resource. The user who performs the import is assigned ownership of the new folders and gets READ|WRITE privilege on the resource and new container folders.

   The *Virtual View Manager Data Services* area does not support auto-creation of new folders.

3. If a resource already exists and it is also being imported, the old version is overwritten (assuming you have the WRITE privilege) in all ways, except the following:

   - The owner is not changed. The original owner retains ownership.
   - Privileges for users that are not explicitly changed by the import are left intact.

     For example, two users, Abe and Bob, have READ|WRITE privileges on an existing resource that will be overwritten by an import. If the imported resource overwrites Abe resource privileges to READ only, but does not mention Bob, then Abe's privileges are updated but Bob's are left intact.

   - Child resources are not removed when the imported resource is a new folder or data source.

4. Restrictions

   - Server configuration settings (available in Virtual View Manager from *Administration > Configuration*) are not ported over when a resource is exported or imported.
   - Import into a folder requires the WRITE privilege on the destination folder and READ privileges on the parent directory path. Importing cannot create a resource in a folder where the user does not have the WRITE privilege.
   - Overwrite or deletion of a resource requires WRITE privilege on that target.
   - Export of a physical data source is either all or nothing. Partial export of a data source configuration is not permitted.
   - You cannot import just part of a physical data source. If you import, you must include the source definition itself.

- The *Virtual View Manager Data Services* area has strict structure rules that are enforced.
  - You cannot import anything that was exported from the *Virtual View Manager Data Services* area to outside of that area.
  - You cannot import anything that was exported from outside the *Virtual View Manager Data Services* area into that area.

Package import from an externalized directory (-pkgdir option):

Package import from an externalized directory (*pkg_import -pkgdir*) takes the XML files that represent the exported output of a directory and uses those files to restore that namespace path to the state captured by the exported files. Importing an externalized directory will delete, overwrite, and synchronize the server directory so that it has the same resource definitions as the externalized directory.

The manifest file in the top level of the externalized directory will dictate what namespace path will get updated with the XML rendered representations of metadata.

Overwrite of a resource requires *WRITE* privilege on that resource. If the importing user does not have *WRITE* privileges on a resource or a folder then those resources in that path can not be deleted or overwritten by an import action performed by that user. Resource locks do not impede resource overwriting by package import.

Package import to a directory (-pkgdir option):

As stated previously, the pkg_import utility is available in the *installation_location* /bin directory. Windows and UNIX specific versions of the utility are available.

The utility follows the following generic form:
```
pkg_import -pkgdir <directory>
-server <hostname> [-port <port>]
[-encrypt]
-user <user> -password <password> [-domain
<domain>]
[-optfile <filename>]
...
[-messagesonly]
[-fileEncoding <encoding>]
[-listIgnored][-includeOwnership]
[-verbose] [-quiet]
```

Example:

In the following example the user, BackupAdmin, assesses the impact of an import from externalized directory that was saved to a disk. The -messagesonly, -listIgnored, and -verbose options are employed to get a full accounting of what files would be replaced or ignored if the package import were to actually get executed. The -messagesonly option disables the import and no changes are actually made.
```
pkg_import -pkgdir E:/DiskArchive/Jan2008 -server localhost
-user BackupAdmin -password becarefulnow -messagesonly -listIgnored
-verbose
```

Verify Package Import Impact using the -messagesonly Option:

It is always a best practice to execute the pkg_import -pkgdir using the -messagesonly option to verify what resources will be overwritten by the import action. The messages only option of package import reports what files will be imported into the server from the externalized directory without actually deleting or overwriting any resources.

pkg_import -pkgdir Parameters and Args	Optional/ mandatory	Comments
-pkgdir *<directory>*	mandatory	Specifies the externalized directory that will be imported to overwrite and synchronize the cognos namespace path(s) named in the manifest.
-server *<HostName>*	mandatory	Server host to which the utility will connect.
-port *<BasePortNum>*	optional	Specifies Web Services Base Port (HTTP) used to communicate with the server. The default value is 9400. Specification is optional if the port setting has not changed. NOTE: Specify the Web Services HTTP port even if the -encrypt option is specified. If -encrypt is specified then the tool will automatically add two to the port number specified here.
- encrypt	optional	Encrypts communication between the command line and the server using SSL sent over the dedicated HTTPS port. **Note:** When using the -encrypt option, the HTTPS port is automatically used. Use the -port option to specify the HTTP web services base port and the command line utility automatically switches the SSL connection to the dedicated HTTPS port.
-user *<UserName>*	mandatory	User name of profile used to import. User rights specified by the target server instance grant permission to import and may restrict write privileges on any particular directory or resource path.
-password *<Password>*	mandatory	Password for user profile used to import the directory.
-domain *<domain>*	optional	Enter domain for the user performing the import. If it is omitted the assumed value is cognos.

pkg_import -pkgdir Parameters and Args	Optional/ mandatory	Comments

`-FileEncoding` `<encoding>`	optional	Encoding defaults to UTF-8, but a different encoding option may be specified. Use the command line to display a list of valid encodings by entering an invalid encoding value, or try to read the 8 point below. Valid encoding values include the following:
		
		`Big5, Big5-HKSCS, Cp1252,`
		`EUC-JP, EUC-KR, GB18030,`
		`GB2312, GBK, IBM-Thai,`
		`IBM00858, IBM01140,`
		`IBM01141, IBM01142,`
		`IBM01143, IBM01144,`
		`IBM01145, IBM01146,`
		`IBM01147, IBM01148,`
		`IBM01149, IBM037, IBM1026,`
		`IBM1047, IBM273, IBM277,`
		`IBM278, IBM280, IBM284,`
		`IBM285, IBM297, IBM420,`
		`IBM424, IBM437, IBM500,`
		`IBM775, IBM850, IBM852,`
		`IBM855, IBM857, IBM860,`
		`IBM861, IBM862, IBM863,`
		`IBM864, IBM865, IBM866,`
		`IBM868, IBM869, IBM870,`
		`IBM871, IBM918,`
		`ISO-2022-CN, ISO-2022-JP,`
		`ISO-2022-KR, ISO-8859-1,`
		`ISO-8859-13, ISO-8859-15,`
		`ISO-8859-2, ISO-8859-3,`
		`ISO-8859-4, ISO-8859-5,`
		`ISO-8859-6, ISO-8859-7,`
		`ISO-8859-8, ISO-8859-9,`
		`JIS_X0201, JIS_X0212-1990,`
		`KOI8-R, Shift_JIS,`
		`TIS-620, US-ASCII,`
		`UTF-16, UTF-16BE,`
		`UTF-16LE, UTF-8,`
		`windows-1250, windows-1251,`
		`windows-1252, windows-1253,`
		`windows-1254, windows-1255,`
		`windows-1256, windows-1257,`
		`windows-1258, windows-31j,`
		`x-Big5-Solaris,`
		`x-euc-jp-linux,`
		`x-EUC-TW, x-eucJP-Open,`
		`x-IBM1006, x-IBM1025,`
		`x-IBM1046, x-IBM1097,`
		`x-IBM1098, x-IBM1112,`
		`x-IBM1122, x-IBM1123,`
		`x-IBM1124, x-IBM1381,`
		`x-IBM1383, x-IBM33722,`
		`x-IBM737, x-IBM834,`
		`x-IBM856, x-IBM874,`
		`x-IBM875, x-IBM921,`
		`x-IBM922, x-IBM930,`

pkg_import -pkgdir Parameters and Args	Optional/ mandatory	Comments
| | | ```
x-IBM933, x-IBM935,
x-IBM937, x-IBM939,
x-IBM942, x-IBM942C,
x-IBM943, x-IBM943C,
x-IBM948, x-IBM949,
x-IBM949C, x-IBM950,
x-IBM964, x-IBM970,
x-ISCII91,
x-ISO-2022-CN-CNS,
x-ISO-2022-CN-GB,
x-iso-8859-11,
x-JIS0208,
x-JISAutoDetect, x-Johab,
x-MacArabic,
x-MacCentralEurope,
x-MacCroatian, x-MacCyrillic,
x-MacDingbat, x-MacGreek,
x-MacHebrew, x-MacIceland,
x-MacRoman, x-MacRomania,
x-MacSymbol, x-MacThai,
x-MacTurkish, x-MacUkraine,
x-MS950-HKSCS, x-mswin-936,
x-PCK, x-windows-50220,
x-windows-50221,
x-windows-874,
x-windows-949, x-windows-950,
x-windows-iso2022jp
``` |
`-includeOwnership`	optional	Imports resources with the original ownership if the importing administrator has the Modify All Users right.
`-listIgnored`	optional	Displays a list of folders and resources that are not updated by the import.
`-messagesonly`	optional	Displays the messages generated in a package import without actually performing the import.

pkg_import -pkgdir Parameters and Args	Optional/ mandatory	Comments
-optfile `<FileName>`	optional	Options file.
		The options file feature provides a way to pass options using a file instead of the command line. The options file is useful for scripting and ensuring that sensitive information does not display on the operating system command line.
		The contents of the options file are the same strings you would use on the command line, with the addition of new lines being legal.
		For example,

		`backup_import`
		`-server`
		`localhost -user test`
		`-password password -pkgfile`
		`sample.car`

		is the same as:
		`backup_import -optfile sample.opt`
		where sample.opt contains either:

		`-server localhost -user`
		`test -password password -pkgfile`
		`sample.car`

		or (this one with new lines):
		`-server localhost`
		`-user test`
		`-password password`
		`-pkgfile sample.car`
		This option is also available for use in pkg_export.
-quiet	optional	Suppresses notification messages.
-verbose	optional	This option causes verbose reporting of information.

remove_services

The remove_services.sh script can be used from the command-line interface to un-install the Virtual View Manager services files that are used to automatically re-start the server and the repository on UNIX.

Procedure

1. Log into the server computer with root privileges.
2. Change to the installation_directory/bin directory.
3. Run the following command:

   remove_services.sh

Results

This command does not interrupt any repository/server processes that are running, but removes the service files.

repo_util

This section describes the usage of Virtual View Manager's command-line program repo_util which you can use to change the repository database.

The repo_util scripts, namely repo_util.bat and repo_util.sh, are available in the *installation_location* /bin directory.

You can use this program to perform several tasks including the following:
- Test the connection to the repository database
- List the current repository configuration information
- Export the repository configuration
- Update the repository configuration
- Create/Drop the repository schema
- Print diagnostic information about the metadata repository
- Print debugging messages

The syntax to use repo_util on Windows:

```
repo_util.bat command [options] [database_configuration_options]
```

The following tables list the commands and options for running the repo_util program. of commands are given at the end.

Commands for using repo_util

Commands for repo_util	Description
-createSchema	Creates the repository schema.
-dropSchema	Drops the repository schema and all data contained within it. This will permanently delete all of the server's data. Use with caution.
-dumpDiagnosticInfo	Prints diagnostic information about the repository database.
-exportConfig	Exports the repository database configuration in the Java property file format. The output is suitable for use as a repository configuration file. See -configFile for details.
-help	Prints this help information.
-listConfig	Lists the repository database configuration in a human readable format.

Commands for repo_util	Description
`-testConn`	Tests the connection to the repository database.
`-updateConfig`	Changes the repository database configuration. Use this command to change one or more configuration options. You may specify new configuration options individually using command line arguments or collectively, using the `-configFile` option. Configuration options that are not specified are left unchanged.

Debug options for using repo_util

Debug options	Description
`-debug`	Prints debug messages.
`-force`	Do not prompt for confirmation. Use with caution.

Database options for using repo_util

Database options	Description
`-configFile`	Read database configuration options from this Java property file. The property names in the file must match the database option names defined in this section. Run `repo_util.bat -exportConfig` for an example of the contents of this file.

`-connectionUrl`	The JDBC URL that is used to connect to the external database.
	
	For example:
	`jdbc:mysql://localhost:3406/`
	`cs030101?continueBatchOnError=`
	`false&useUnicode=true`

`-databaseCatalog`	Database catalog that contains the schema. It may be blank if the database does not support catalogs.
`-databaseSchema`	Database schema that contains the schema. It may be blank if the database does not support schemas.
`-databasePassword`	Database password.

Database options	Description
`-databaseUser`	Database user name.
`-driverClass`	The fully qualified class name of a JDBC compliant driver. For example: `com.mysql.jdbc.Driver`
`-driverClassPath`	A semi-colon or colon separated list of JAR files and directories. For example: `/tmp/oracle40.jar:/tmp`
`-driverName`	Name of the datasource driver name. This is required for proper operation of the system tables.
`-driverType`	Name of the data source driver type. This is required for proper operation of the system tables.
`-repositoryClass`	Repository class name.
`-schemaCreateScript`	Schema creation script. This script contains the SQL commands to create the Virtual View Manager schema.
`-schemaDropScript`	Schema drop script. This script contains the SQL commands to drop the Virtual View Manager schema.
`-schemaInitializeScript`	Initialization script. This script contains the SQL commands to initialize the Virtual View Manager tables.

Sample commands to use repo_util

Here are some examples of commands for using the repo_util program:

- To list the server configuration information:
  ```
  repo_util.bat
  -listConfig
  ```
- To export a repository configuration file:
  ```
  repo_util.bat
  -exportConfig > repo.properties
  ```
- To update the repository database user name and password:
  ```
  repo_util.bat
  -updateConfig -databaseUser someuser -databasePassword somepassword
  ```
- To update the repository configuration using a repository configuration file, overriding the database password:
  ```
  repo_util.bat
  -updateConfig -configFile repo.properties -databasePassword somepassword
  ```

server_util

The server_util program is used for getting the server performance profile report and resetting the system namespace.

Usage syntax:
```
server_util -server <hostname> [-port
<port>] [-encrypt]
-user <username> -password
<password> [-domain <domain>]
<command> [-verbose]
```

where <command> refers to one of the following options:
* -resetNamespace resets the server namespace to show changes to the system namespace, such as a system table change after application of a patch.
* -profile displays current server profile performance profiling data.

  Server profiling is always working in the background by default. Some of the report contents will display performance metrics once that code path has been exercised. Some pertinent performance profiling is detailed below.
* -clearProfile clear all existing profiling data.
* -regenerateFiles regenerate files that are based on configuration settings. Use of the Configuration window Apply button and server restart generally make this command obsolete.
* -createMemorySizeFile calculates and saves object memory sizes.
* -getServerName retrieves the server name.
* -setServerName -serverName <server name> sets server name, which must be a unique display name.

Sample commands to use server_util

The following command would list the server profile for a default installation:
```
server_util -server HostName -user
admin -password admin
-profile
```

The Server Performance Profile Report contains many different metrics of which only a few are described here:
* request.setup.sql -- Time to construct the SQL sent to outside data sources.
* components.archive -- Import and Export detail
* request.data -- Aggregated amount of time required to send data to the requesting client after SQL processing was underway.
* internal.repository -- Repository response time for metadata information gathering.
* components.operator -- Query engine processing time for SQL that could not be pushed to the underlying data sources during.
* ds -- Aggregated time required to communicate with external data sources.

The following -clearProfile command clears the server profile information to start aggregate and average statistics data with a clean slate.
```
server_util -server localhost -user admin -password admin
-clearProfile
```

The following -resetNamespace command would reset the system namespace.

```
server_util -server localhost -user admin -password admin
-resetNamespace
```

Chapter 8. Setting up a metadata repository

You can set up an IBM Informix, MySQL, Oracle or Sybase database as the IBM Cognos Virtual View Manager metadata repository. IBM Informix is the default database for the Virtual View Manager metadata repository.

The following topics are covered in this chapter:
- "System requirements"
- "Creating a database for the metadata repository" on page 118
- "Additional database configuration requirements" on page 121
- "Configuring the IBM Cognos Virtual View Manager metadata repository" on page 123

System requirements

Ensure that enough space is allocated for the database repository.

The following table lists the recommended space to allocate according to the size of your database repository.

Database size	Space
Small-sized repository database	250MB
Medium-sized repository database	800MB
Large-sized repository database	2000MB (or 2GB)

Pre-requisites and limitations for database types

There are a couple of things to consider if you plan to use Sybase or Oracle as the metadata repository for Virtual View Manager.

Using Sybase ASE 12.5 as a repository database

Indexes are created on the IBM Cognos Virtual View Manager system tables in order to enhance performance of querying these tables. These indexes consist of metadata ID-s and metadata names. By default, Sybase ASE 12.5 has page size of 4K so it only allows each index to have a maximum size of 1250 bytes. Therefore, you must limit the length of the metadata names.

Metadata names can be a maximum of 200 characters. If the metadata name is longer than 200 characters, the name will be trimmed to 200 characters when it is persisted into Sybase repository. For this reason, when you create an index you must define each metadata name as UNIVARCHAR(200).

Example - Index sizing in a system table

In this example, an index on the system columns table comprises the following:
- DATASOURCE_ID
- ORDINAL_POSITION

- COLUMN_ID
- CATALOG_NAME
- SCHEMA_NAME
- TABLE_NAME

If you set the name columns to UNIVARCHAR(200), then each UNIVARCHAR character requires two bytes of storage. Since three of these name fields add up to 600 bytes (3 * 200 * 2 bytes = 1200 bytes), if you include the bytes for DATASOURCE_ID, COLUMN_ID, and ORDINAL_POSITION, this index barely fits 1250 bytes. If no more columns are added to the index, the total bytes in the index will exceed the maximum allowable size of an index. As a result, this index will not be created in ASE 12.5.

Using Oracle Call Interface (OCI) as a repository database

To use the Oracle 10g type 2 database as a metadata repository, the Oracle Call Interface (OCI) client must be installed on the computer where Virtual View Manager is installed.

The PATH (Windows) or LD_LIBRARY_PATH (for most UNIX systems) to OCI must be specified for Virtual View Manager to properly recognize and use the OCI client.

For more detailed information, see the Oracle Instant Client documentation.

Creating a database for the metadata repository

The following sections describe how to create an external database repository to work with IBM Cognos Virtual View Manager.

Create an IBM Informix metadata repository

You can create an IBM Informix metadata repository for IBM Cognos Virtual View Manager.

Before you begin

If you are using an external Informix repository, you must open <vvm_install_location>//conf/repository/informix/composite_schema.sql and add a forward slash (/) after the first three SQL statements:

```
CREATE FUNCTION PartialString(path LVARCHAR(2000)) RETURNS
VARCHAR(255)
WITH (NOT VARIANT);
RETURN SUBSTR(path,1,255);
END FUNCTION;
/
CREATE FUNCTION PartialString32(path VARCHAR(255)) RETURNS VARCHAR(32)
WITH (NOT VARIANT);
RETURN SUBSTR(path,1,32);
END FUNCTION;
/
CREATE FUNCTION PartialString64(path VARCHAR(255)) RETURNS VARCHAR(64)
WITH (NOT VARIANT);
RETURN SUBSTR(path,1,64);
END FUNCTION;
/
```

Procedure

1. Connect to an Informix server instance by setting the following informix environment parameters:

```
set CLIENT_LOCALE=EN_US.utf8
set DB_LOCALE=EN_US.utf8
set SERVER_LOCALE=EN_US.utf8
set DBLANG=EN_US.utf8
```

2. Create a database called VVM_REPO by using the following command:

```
create database "VVM_REPO" in VVM_REPO with log
```

   Note: You need to allow log transactions.

3. Create a user for the VVM_REPO database, for example, you can create a user called cognos.

4. Grant DBA access to the cognos user by using the following command:

```
grant DBA to cognos
```

Create a MySQL metadata repository

You can create a MySQL metadata repository for IBM Cognos Virtual View Manager.

Procedure

1. Create a MySQL database for the metadata repository by using the following command:

```
CREATE DATABASE "VVM_REPO"
```

2. Create a user for the VVM_REPO database, for example, you can create a user called cognos.

3. Grant access to the cognos user by using the following command:

```
grant usage on *.* to cognos@localhost identified by
'password'
```

4. Grant privileges to the cognos user by using the following command:

```
grant all privileges on VVM_REPO.* to cognos@localhost
```

Create a Sybase metadata repository

You can create a Sybase metadata repository for IBM Cognos Virtual View Manager.

Procedure

1. Create a Sybase database for the metadata repository. Use the following syntax to create a database in Sybase ASE 12.5. Substitute the items in italics with appropriate values.

```
create database VVM_REPO
    [on {default | database_device} [= size]
        [, database_device [= size]...]
    [log on database_device [ = size ]
        [, database_device [= size]]...]
    [with {override | default_location = "pathname"}]
    [for {load | proxy_update}]
```

   For detailed information about the above syntax, refer to Sybase ASE 12.5 documentation.

2. To configure the database after creating it, run the following system procedures supplying appropriate values for the items in italics:

```
sp_dboption VVM_REPO, 'abort tran on log full', true
go
sp_dboption VVM_REPO, 'trunc log on chkpt', true
go
```

3. Increase the size of the tempdb database to be about 20 to 25 per cent of the main database using ALTER DATABASE command.

4. Create a user to log into the newly created database, supplying appropriate values for the items in italics.

   Syntax for adding a new account to log into ASE 12.5:

```
sp_addlogin loginame,
password [, defdb]
        [, deflanguage [, fullname]]]
```

   Syntax for adding the created login to a specific database so that the newly created user can access the database:

```
sp_adduser login_name [, name_in_db [, group_name]]
```

Results

The Sybase repository is created, and ready to be configured.

Creating an Oracle metadata repository

This section describes how to create a tablespace and user to use an Oracle repository database. The examples in the steps use VVMRepo for the SID, VM_REPO for the tablespace, cognos or the user.

For a complete syntax to create a tablespace in Oracle that fits your needs, refer to Oracle documentation.

Procedure

1. Contact your database administrator to obtain the Oracle server SID, since the SID is used as the value of databaseName. You need to provide the database name and SID to create an Oracle repository database. In this example, the SID is VVMRepo.

2. Ensure that the Oracle system supports the international character set. Set the database character set to Unicode AL32UTF8, and the national character set to AL16UTF16.

3. Create a tablespace, named VVM_REPO in this example, with the following command:

```
CREATE TABLESPACE "VVM_REPO"
    LOGGING
    DATAFILE 'C:\ORACLE\PRODUCT\10.2.0\ORADATA\VVMRepo\VVMRep.ora'
SIZE 250M
```

4. Create a user named cognos using the following command:

```
CREATE USER "cognos" PROFILE "DEFAULT"
    IDENTIFIED BY "password" DEFAULT TABLESPACE "VVM_REPO"
    TEMPORARY TABLESPACE "TEMP"
    QUOTA UNLIMITED
    ON "VVM_REPO"
    ACCOUNT UNLOCK;
GRANT ALTER SESSION TO cognos
GRANT CREATE ANY INDEX TO cognos
GRANT CREATE PROCEDURE TO cognos
GRANT CREATE TABLE TO cognos
GRANT CREATE SESSION TO cognos
GRANT CREATE SYNONYM TO cognos
```

```
GRANT CREATE VIEW TO cognos
GRANT DROP ANY INDEX TO cognos
GRANT DROP ANY PROCEDURE TO cognos
GRANT DROP ANY TABLE TO cognos;
```

This command creates a user named cognos that is granted access to the newly created tablespace, VVM_REPO, with unlimited space quota. The user is also granted the privileges of creating and dropping indexes, procedures, and tables.

5. Copy the ojdbc.jar file to intallation_location/VVM/apps/server/lib directory of the computer where Virtual View Manager is installed.

Results

The cognos user can be used as the value for the properties databaseUser and databaseSchema described in "Configuring the IBM Cognos Virtual View Manager metadata repository" on page 123.

Additional database configuration requirements

If you are using MySQL or Sybase as the metadata repository, then you need to perform additional configurations for your database.

Configure the MySQL repository

You must set up an external MySQL repository to work with Virtual View Manager.

Procedure

1. Set the following variables in the my.ini file that MySQL system uses for initialization:

```
set-variable=max_connections=100
set-variable=lower_case_table_names=1
set-variable=max_allowed_packet=16M
default-character-set=utf8
```

If you modify the my.ini file to match the required settings, you need to restart the MySQL server for the new settings to be effective.

2. The program MySQL Administrator can be used to set/change MySQL's startup variables.

Enable Use concurrent inserts in MyISAM parameters.

Enable Activate InnoDB in InnoDB parameters.

Configure the Sybase repository

The default installation of Sybase ASE 12.5 is not configured for enterprise applications. For example, the tempdb size is 2MB. The maximum number of connections defaults to 25, and the default maximum network package size is only 512 bytes. If you use such configuration to run IBM Cognos Virtual View Manager, certain Sybase limits will be hit, causing the Sybase system to hang. If you are not sure of your ASE 12.5 configuration, ask your Sybase DBA about it.

In order to run Virtual View Manager efficiently, Sybase ASE 12.5 needs to be configured appropriately.

To reconfigure ASE 12.5, you must have system administration privileges so you can execute a set of system stored procedures.

Configure memory and cache

You must configure memory and cache in Sybase ASE 12.5.

Procedure

1. Run the following system procedures in Sybase ASE 12.5:

```
sp_configure 'max memory', 128000
go
sp_configure 'lock shared memory', 1
go
sp_cacheconfig 'default data cache', '128M'
go
sp_cacheconfig 'procedure cache', '16M'
go
sp_cacheconfig 'cache01', '80M'
go
```

2. Restart Sybase ASE 12.5 after these configuration parameters are applied.

Configure processing limits

You must configure processing limits in Sybase ASE 12.5.

Procedure

1. Run the following system procedures in Sybase ASE 12.5:

```
sp_configure 'number of user connections', 256
go
sp_configure 'number of worker processes', 100
go
sp_configure 'max parallel degree', 3
go
sp_configure 'max scan parallel degree', 3
go
sp_configure 'number of locks', 100000
go
sp_configure 'number of open objects', 50000
go
sp_configure 'number of open databases', 32
go
sp_configure 'number of devices', 32
go
```

2. Restart Sybase ASE 12.5 after these configuration parameters are applied.

Configure network parameters

You must configure network parameters in Sybase ASE 12.5.

Procedure

1. Run the following system procedures in Sybase ASE 12.5:

```
sp_configure 'additional network memory', 8192
go
sp_configure 'max network packet size', 4096
go
sp_configure 'default network packet size', 2048
go
sp_configure 'heap memory per user', 4096
go
```

2. Restart Sybase ASE 12.5 after these configuration parameters are applied.

Configure the character set

You must configure the character set in Sybase ASE 12.5.

Procedure

1. In order for the multi-byte character set to persist correctly in the Sybase repository, configure the Sybase server to use the utf-8 character set by opening the Server configuration tool for Sybase.

2. Click **Configure Adaptive Server**.

3. Select the repository Sybase server, and log into that server.

4. In the **Configure Adaptive Server** window, click the **Language** button.

   To change the character set to utf-8, you need to make sure that the utf-8 character set is installed. If UTF-8 is not already installed, you need to add it.

5. To add the utf-8 character set, click the **Add/Remove** button in the **Character Set** group in the section **Change Options**.

6. Select **Unicode 3.1 UTF-8 Character Set** in the **Available** character set section, and click **Add**.

7. When **Unicode 3.1 UTF-8 Character Set** is added to the section **Selected**, click **OK** to accept the selection.

8. In the **Language Options** window, chooses **Set Default** button in the **Character Set** group in the section **Change Options**. This command opens the **Change Default Character Set** window.

9. Make sure that **Unicode 3.1 UTF-8 Character Set** is selected, and click **OK**.

10. Click **Save** to make the UTF-8 character set to be the default.

    The **Save** command requires a server restart. So, make sure that the server is not running any critical tasks, and restart the server.

Configuring the IBM Cognos Virtual View Manager metadata repository

You can use the command line program repo_util to configure the IBM Cognos Virtual View Manager metadata repository.

Use the command line program repo_util to:

- list the configurations of the current repository database
- modify the configuration for the existing repository database
- change the repository database

The repo_util program is available in the *installation_location*/bin directory where IBM Cognos Virtual View Manager is installed. You use this utility program to create the repo.properties file.

Sample repo.properties File

The following is a sample of a repo.properties file for an IBM Informix repository.

```
#Virtual View Manager Informix Repository Configuration
for repo.properties
#Property Name            Value
#--------------------------------------------------------
connectionUrl=jdbc:informix-sqli://localhost:9408/cognos080401:informixserver=
VVMRepository;DB_LOCALE=EN_US.utf8
connectionValidationQuery=SELECT * FROM metadata_version
databaseCatalog=cognos080401
databaseHost=localhost
databaseName=cognos080401
databasePassword=ibm@123#123
databasePort=9408
databaseSchema=informix
databaseUser=informix
driverClass=informix.jdbc.IfxDriver
```

```
driverClassPath=
driverLibraryPath=
driverName=Informix 9.x
driverType=Informix
poolInitialSize=5
poolMaxSize=50
poolMinSize=5
repositoryClass=com.compositesw.server.repository.internal.InformixRepository
schemaCreateScript=C:\Program
Files\\cognos\\vvm\\8.4.0\\conf\\repository\\informix\\composite_schema.sql
schemaDropScript=C:\Program Files\\cognos\\vvm\\8.4.0\\conf\\repository\\informix\\
composite_clean.sqlschemaInitializeScript=C:\Program
Files\\cognos\\vvm\\8.4.0\\conf\\repository\\informix\\composite_data.sql
```

Procedure

1. Stop Virtual View Manager Server, if it is running.
2. Use one of the following commands to list the current repository configuration:

   repo_util.bat -listConfig (Windows command)

   repo_util.sh -listConfig (UNIX command)
3. Export the current configurations into the repo.properties file using one of the following commands:

   repo_util.bat -exportConfig > repo.properties (Windows command)

   repo_util.sh -exportConfig > repo.properties (UNIX commands)
4. Modify the repo.properties file. Use the following table to help you modify the properties.

Property	Description
connectionUrl	Enter a value for the JDBC connection.
	For IBM Informix:
	jdbc:informix-sqli://host_name:port_number/ databasename:INFORMIXSERVER=VVMRepository;
	DB_LOCALE=en_us.utf8
	For MySQL:
	jdbc\:mysql\://hostname\:3306/ \<database_name>?continueBatchOnError\=false&useUnicode\ =true&characterEncoding\
	=utf8&clobberStreamingResults\=true&useReadAheadInput\ =true
	For Sybase:
	jdbc\:sybase\:Tds\://localhost\:5000/pubs2
	For Oracle OCI:
	jdbc\:oracle\:oci\:@localhost\:1521\:orcl
	For Oracle Thin:
	jdbc\:oracle\:thin\:@localhost\:1521\:orcl

Property	Description
databaseCatalog	Enter a value for the database catalog. It may be blank if the database does not support catalogs.
databasehost	Enter the name of the database host. For example, localhost.
databaseName	Enter the name of the database for your metadata repository.databasePassword Enter the password for the database.
databaseSchema	Enter the database that contains the schema. It may be blank if the database does not support schemas.
driverClass	Enter the driver class for your database as follows: • For IBM Informix, enter com.informix.IfxDriver • For MySQL, enter com.mysql.jdbc.Driver • For Sybase, enter net.sourceforge.jtds.jdbc.Driver • For Oracle OCI and Oracle Thin, enter oracle.jdbc.driver.OracleDriver
driverClassPath	Enter location of the driver file.
driverName	Enter the name of the driver.
driverType	Enter the type of database driver.
poolInitialSize	Enter the initial size for the buffer pool. The default value is 5.
poolMaxSize	Enter the maximum size for the buffer pool. The default value is 50.
poolMinSize	Enter the minimum size for the buffer pool. The default value is 25.
repositoryClass	Enter the repository class for your repository using the following format: com.compositesw.server.repository.internal. <database_type>Repository For example, com.compositesw.server.repository.internal.InformixRepository
schemaCreateScript	Enter the location of the initialization script. For example: installation_location\:\\database_type\composite_schema.sql
schemaDropScript	Enter the location of the script to delete the schema from the repository. For example: installation_location\:\\database_type\composite_clean.sql

Property	Description
schemaInitializeScript	Enter the location of the initialization script. For example: *installation_location\:\\database_type\composite_data.sql*

5. Update the repository configuration. Supply the real password in a command line as follows:

   repo_util.bat -updateConfig -configFile repo.properties -databasePassword *password* (Windows)

   repo_util.sh -updateConfig -configFile repo.properties -databasePassword *password* (UNIX)

6. After the repository database configuration has been updated, restart the Virtual View Manager server.

7. Verify that the repository database is running on the database by running the repo_util command with -listConfig option.

   Note: You can view the server information in the cs_server.log file, located in the *installation_location/* logs directory.

Chapter 9. SNMP traps

IBM Cognos Virtual View Manager system supports SNMP v1 traps. This appendix provides a complete list of events and their corresponding SNMP traps. The server generates traps for monitoring the events that occur in the server.

See "Server event attributes" on page 71 for more information.

SNMP log settings

You can modify SNMP log settings.

Modify SNMP log settings from the **Configuration** window. You open the **Configuration** window by selecting the **Administration** > **Configuration** menu option, and navigating to the SNMP folder.

SNMP details

The SNMP details in the tables below are grouped into several categories.

These categories include:
- Monitor and server events
- Requests
- Transactions
- Cached resources
- Triggers
- Data Sources
- Sessions
- Resources
- Storage

SNMP details for monitor events

SNMP ID	Short name for the event	Description
10000	csMonitorStart	Variables: { trapTime, trapServerHostName, trapServerPort } Description: "This trap is generated when a server monitor was started."
10001	csMonitorStop	Variables: { trapTime, trapServerHostName, trapServerPort } Description: "This trap is generated when a server monitor was stopped."

SNMP ID	Short name for the event	Description
10002	csMonitorFail	Variables: { trapTime, trapServerHostName, trapServerPort }

Description: "This trap is generated when a server monitor had a failure." |
| 10003 | csServerStopUnplanned | Variables: { trapTime, trapServerHostName, trapServerPort }Description: "This trap is generated when a server had an unplanned stop." |
| 10004 | csServerStopPlanned | Variables: { trapTime, trapServerHostName, trapServerPort }

Description: "This trap is generated when a server had a planned stop." |
| 10005 | csServerRestart | Variables: { trapTime, trapServerHostName, trapServerPort }

Description: "This trap is generated when a server was restarted." |
| 10006 | csServerRestartFail | Variables: { trapTime, trapServerHostName, trapServerPort }

Description: "This trap is generated when a server had a restart failure." |
| 10007 | csRepositoryUp | Variables: { trapTime, trapServerHostName, trapServerPort }

Description: "This trap is generated when a server repository was started." |
| 10008 | csRepositoryDown | Variables: { trapTime, trapServerHostName, trapServerPort }

Description: "This trap is generated when a server repository was stopped." |

SNMP details for server events

SNMP ID	Short name for the event	Description
20000	csServerStart	Variables: { trapTime, trapServerHostName, trapServerPort }

Description: "This trap is generated when a server was started." |

SNMP ID	Short name for the event	Description
20001	csServerStop	Variables: { trapTime, trapServerHostName, trapServerPort } Description: "This trap is generated when a server was stopped."
20002	csUserCreate	Variables: { trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName } Description: "This trap is generated when a user was created in a domain."
20003	csGroupCreate	Variables: { trapTime, trapServerHostName, trapServerPort, trapGroupName, trapDomainName } Description: "This trap is generated when a group was created in a domain."
20004	csUserDelete	Variables: { trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName } Description: "This trap is generated when a user was deleted from a domain."
20005	csGroupDelete	Variables: { trapTime, trapServerHostName, trapServerPort, trapGroupName, trapDomainName } Description: "This trap is generated when a group was deleted from a domain."
20006	csUserAddToGroup	Variables: { trapTime, trapServerHostName, trapServerPort, trapUserName, trapUserDomainName, trapGroupName, trapGroupDomainName } Description: "This trap is generated when a user was added to a group."
20007	csUserRemoveFrom Group	Variables: { trapTime, trapServerHostName, trapServerPort, trapUserName, trapUserDomainName, trapGroupName, trapGroupDomainName } Description: "This trap is generated when a user was removed from a group."
20008	csDomainCreate	Variables: { trapTime, trapServerHostName, trapServerPort, trapDomainName } Description: "This trap is generated when a domain was created."

SNMP ID	Short name for the event	Description
20009	csDomainDelete	Variables: { trapTime, trapServerHostName, trapServerPort, trapDomainName } Description: "This trap is generated when a domain was deleted."
20010	csUserPasswordModify	Variables: { trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName } Description: "This trap is generated when a user password was modified."

SNMP details for requests

SNMP ID	Short name for the event	Description
20100	csRequestStart	Variables: { trapTime, trapServerHostName, trapServerPort, trapRequestId, trapOptionalRequestParameter1, trapOptionalRequestParameter2, trapOptionalRequestParameter3, trapOptionalRequestParameter4 } Description: "This trap is generated when a request was started."
20101	csRequestWait	Variables: { trapTime, trapServerHostName, trapServerPort, trapRequestId, trapTransactionId, trapSessionId } Description: "This trap is generated when a request was waiting to run."
20102	csRequestEnd	Variables: { trapTime, trapServerHostName, trapServerPort, trapRequestId, trapTransactionId, trapSessionId } Description: "This trap is generated when a request was completed."
20103	csRequestFail	Variables: { trapTime, trapServerHostName, trapServerPort, trapRequestId, trapOptionalRequestParameter1, trapOptionalRequestParameter2, trapOptionalRequestParameter3 } Description: "This trap is generated when a request has failed."

SNMP ID	Short name for the event	Description
20104	csRequestCancel	Variables: { trapTime, trapServerHostName, trapServerPort, trapRequestId, trapTransactionId, trapSessionId } Description: "This trap is generated when a request was cancelled."
20105	csRequestWaitQueue ThresholdPass	Variables: { trapTime, trapServerHostName, trapServerPort, trapRequestId, trapTransactionId, trapSessionId } Description: "This trap is generated when a request passed the wait queue threshold."
20106	csRequestWaitQueue ThresholdReset	Variables: { trapTime, trapServerHostName, trapServerPort, trapRequestId, trapTransactionId, trapSessionId } Description: "This trap is generated when a request reset the wait queue threshold."
20107	csPreparedStatement Success	Variables: { trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapOptionalRequestParameter1, trapOptionalRequestParameter2 } Description: "This trap is generated when a prepared statement was successfully executed."
20108	csPreparedStatementFail	Variables: { trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapSqlQuery, trapOptionalRequestParameter1, trapOptionalRequestParameter2 }Description: "This trap is generated when a prepared statement has failed during execution."

SNMP details for transactions

SNMP ID	Short name for the event	Description
20200	csTransactionStart	Variables: { trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapSessionId } Description: "This trap is generated when a transaction was started."
20201	csTransactionCommit	Variables: { trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapSessionId } Description: "This trap is generated when a transaction was committed."

SNMP ID	Short name for the event	Description
20202	csTransactionFail	Variables: { trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapMessage, trapStackTrace, trapSessionId } Description: "This trap is generated when a transaction has failed."
20203	csTransactionRollBack	Variables: { trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapSessionId } Description: "This trap is generated when a transaction was rolled back."
20204	csTransaction Compensate	Variables: { trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapMessage, trapSessionId } Description: "This trap is generated when a transaction was compensated for."

SNMP details for cached resources

SNMP ID	Short name for the event	Description
20300	csCacheEnable	Variables: { trapTime, trapServerHostName, trapServerPort, trapCacheName } Description: "This trap is generated when a cache was enabled."
20301	csCacheDisable	Variables: { trapTime, trapServerHostName, trapServerPort, trapCacheName } Description: "This trap is generated when a cache was disabled."
20302	csCacheClear	Variables: { trapTime, trapServerHostName, trapServerPort, trapCacheName, trapCacheParameters } Description: "This trap is generated when a cache was cleared."
20303	csCacheRefreshStart	Variables: { trapTime, trapServerHostName, trapServerPort, trapCacheName, trapCacheParameters } Description: "This trap is generated when a cache refresh was started."

SNMP ID	Short name for the event	Description
20304	csCacheRefreshEnd	Variables: { trapTime, trapServerHostName, trapServerPort, trapCacheName, trapCacheParameters } Description: "This trap is generated when a cache refresh was completed."
20305	csCacheRefreshFail	Variables: { trapTime, trapServerHostName, trapServerPort, trapCacheName, trapCacheParameters, trapOptionalMessage } Description: "This trap is generated when a cache refresh has failed."

SNMP details for triggers

SNMP ID	Short name for the event	Description
20400	csTriggerStart	Variables: { trapTime, trapServerHostName, trapServerPort, trapTriggerName, trapTriggerType, trapTriggerAction } Description: "This trap is generated when a trigger was started."
20401	csTriggerEnd	Variables: { trapTime, trapServerHostName, trapServerPort, trapTriggerName, trapTriggerType, trapTriggerAction } Description: "This trap is generated when a trigger was completed."
20402	csTriggerFail	Variables: { trapTime, trapServerHostName, trapServerPort, trapTriggerName, trapTriggerType, trapTriggerAction, trapOptionalMessage } Description: "This trap is generated when a trigger has failed."

SNMP details for data sources

SNMP ID	Short name for the event	Description
20500	csDataSourceOn	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType } Description: "This trap is generated when a data source is enabled."

SNMP ID	Short name for the event	Description
20501	csDataSourceOff	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType } Description: "This trap is generated when a data source is disabled."
20502	csDataSourceUp	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType } Description: "This trap is generated when a data source was started."
20503	csDataSourceDown	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType } Description: "This trap is generated when a data source was stopped."
20504	csDataSourceModify	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType } Description: "This trap is generated when a data source was modified."
20505	csIntrospectStart	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType } Description: "This trap is generated when a data source introspection was started."
20506	csIntrospectEnd	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType, trapDataSourceReport } Description: "This trap is generated when a data source introspection was completed."
20507	csIntrospectCancel	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType } Description: "This trap is generated when a data source introspection was cancelled."
20508	csIntrospectFail	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType, trapMessage} Description: "This trap is generated when a data source introspection has failed."

SNMP ID	Short name for the event	Description
20509	csTestStart	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType } Description: "This trap is generated when a data source test was started."
20510	csTestSuccess	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType } Description: "This trap is generated when a data source test was successful."
20511	csTestFail	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType } Description: "This trap is generated when a data source test has failed."
20512	csConnPoolSizeIncrease	Variables: { trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId } Description: "This trap is generated when the size of a connection pool has increased."
20513	csConnPoolSizeDecrease	Variables: { trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId } Description: "This trap is generated when the size of a connection pool has decreased."
20514	csConnCheckOut	Variables: { trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId } Description: "This trap is generated when a connection was checked out a connection pool."
20515	csConnCheckIn	Variables: { trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId } Description: "This trap is generated when a connection was checked in a connection pool."
20516	csConnInvalid	Variables: { trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId } Description: "This trap is generated when a connection pool had an invalid connection."
20517	csConnFail	Variables: { trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId } Description: "This trap is generated when a connection pool had a failed connection."

SNMP ID	Short name for the event	Description
20518	csConnPoolExhaust	Variables: { trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId } Description: "This trap is generated when a connection pool has exhausted its connections."
20519	csStatisticsProcessingStart Process	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourcePath } Description: "This trap is generated when a data source started the statistics processing process."
20520	csStatisticsProcessing Complete	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourcePath } Description: "This trap is generated when a data source completed the statistics processing process."
20521	csStatisticsProcessing CompletePartial	This event/message is deprecated. Statistics processing with ID # is partially completed.
20522	csStatisticsProcessing Failed	Variables: { trapTime, trapServerHostName, trapServerPort, trapDataSourcePath, trapMessage } Description: "This trap is generated when a data source failed to complete the statistics processing process."
20523	csStatisticsProcessing Update	This event/message is deprecated. Statistics processing with ID # updated the information.

SNMP details for sessions

SNMP ID	Short name for the event	Description
20700	csSessionLoginFail	Variables: { trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName } Description: "This trap is generated when a session login has failed for a user."
20701	csSessionStart	Variables: { trapTime, trapServerHostName, trapServerPort, trapSessionId, trapUserName, trapDomainName } Description: "This trap is generated when a session was started for a user."

SNMP ID	Short name for the event	Description
20702	csSessionEnd	Variables: { trapTime, trapServerHostName, trapServerPort, trapSessionId, trapUserName, trapDomainName } Description: "This trap is generated when a session was ended for a user."
20703	csSessionTerminate	Variables: { trapTime, trapServerHostName, trapServerPort, trapSessionId, trapUserName, trapDomainName } Description: "This trap is generated when a session was terminated for a user."
20705	csSessionMaxConnections Exhaust	Variables: { trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName, trapHostName, trapLocalHostName, trapLocalHostIP } Description: "This trap is generated when a session creation request was denied for a user."

SNMP details for resources

SNMP ID	Short name for the event	Description
20800	csResourceCreate	Variables: { trapTime, trapServerHostName, trapServerPort, trapResourceName, trapDataSourcePath, trapResourceType } Description: "This trap is generated when a resource was created."
20801	csResourceDelete	Variables: { trapTime, trapServerHostName, trapServerPort, trapResourceName, trapDataSourcePath, trapResourceType } Description: "This trap is generated when a resource was deleted."
20802	csStatisticsResource ProcessingStartProcess	Variables: { trapTime, trapServerHostName, trapServerPort, trapResourcePath } Description: "This trap is generated when a resource started the statistics processing process."
20803	csStatisticsResource ProcessingComplete	Variables: { trapTime, trapServerHostName, trapServerPort, trapResourcePath } Description: "This trap is generated when a resource completed the statistics processing process."

SNMP ID	Short name for the event	Description
20804	csStatisticsResource ProcessingFailed	Variables: { trapTime, trapServerHostName, trapServerPort, trapResourcePath, trapMessage } Description: "This trap is generated when a resource failed to complete the statistics processing process."
20805	csResourceLock	Variables: { trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName, trapLockTime, trapResourcePath, trapResourceType, trapResourceSubType } Description: "This trap is generated when a resource was locked." Only the top-most parent node is reported as locked.
20806	csResourceUnlock	Variables: { trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName, trapUnlockTime, trapResourcePath, trapResourceType, trapResourceSubType, trapComment } Description: "This trap is generated when a resource was unlocked." Only the top-most parent node is reported as unlocked.

SNMP details for storage

SNMP ID	Short name for the event	Description
21000	csStorageLowWarning	Variables: { trapTime, trapServerHostName, trapServerPort } Description: "This trap is generated when a storage low warning has occurred on a machine."
21001	csStorageLowCritical	Variables: { trapTime, trapServerHostName, trapServerPort } Description: "This trap is generated when a storage low critical event has occurred on a machine."

Notices

This information was developed for products and services offered worldwide.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. This document may describe products, services, or features that are not included in the Program or license entitlement that you have purchased.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Software Group
Attention: Licensing
3755 Riverside Dr
Ottawa, ON K1V 1B7
Canada

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, ibm.com, and Cognos are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at " Copyright and trademark information " at www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies:

- Netezza is a registered trademark or trademark of Netezza Corporation, an IBM Company.
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Index

Special characters

-pkgdir 92

A

Access
    Privilege Cache 59
Access (Hits/ Accesses)
    Manager Server Overview 59
Active Requests
    data sources 62
Active Server Requests 56
admin
    group, cognos domain 26
    user 28
administration
    dynamic domain 49
Advanced Filter dialog 55
all
    group, cognos domain 26
Allocated Pool Size
    data sources 62
anonymous 28
audience of document vii

B

backup_export 75
backup_import 75
buttons
    enabling 54
Bytes summary
    data sources 62

C

Cached Resources
    Change Enabling button 60
    Last Access 60
    Last Fail Duration 61
    Last Fail End 60
    Last Refresh End 60
    Last Success Duration 61
    Last Success End 61
    Owner 60
    Owner domain 61
    Path 61
    Refresh Cache button 60
    Status 60
    Storage Used 60
    Total Accesses 61
    Total Failures 61
    Total Successes 61
    Type 60
    Variant 60
Capacity (Entries/Max)
    Server Overview 59
Change Enabling button 70
changing ownership 31

(Column 2)

changing password 31
cognos domain
    built-in groups 26
    built-in users, privileges 28
    group membership 30
    user 29
configuration settings
    triggers 70
configuring
    databases 121
    MySQL repository 121
    Sybase repository 121
    Virtual View Manager metadata repository 123
connection URL 2
connection URL format 6
Connection URL Pattern 8
copy table filter 55
create table filter 55
creating
    database 118
    Informix metadata repository 118
    MySQL metadata repository 119
    Oracle metadata repository 120
    Sybase metadata repository 119
creating with driverConfig 13

D

data sources 2
    removing 9
Data Sources
    Active Requests 63
Data Sources console 61
database repository
    system requirements 117
DB2 type 2 JDBC driver 2
DB2 z/OS 3
DB2 z/OS JDBC driver 3
description of product vii
digital certificate 2
Driver Class Name 8
driverConfig, running 13
drivers 2, 6
dymanic domain
    about 49
    enabling 50
    group administration 50
dymanic domain administration 49
dymanic domain user administration 51
    adding users 51
    removing users 52
dymanic domain user group membership 52
dymanic domain user privledges 51

E

enabling buttons 54
End Sessions button 67
environment variables
    setting on UNIX and Linux operating systems 12