



TIBCO® Data Virtualization

Administration Guide

Version 8.8.1 | March 2025

Contents

Contents	2
Basic TDV Administration Tasks	17
Overview of TDV Administration	17
Security Features	19
Starting TDV Processes on Windows	19
Starting and Stopping TDV Processes on Windows	20
Keeping the TDV Server Process Running after Logoff	22
Turning Off Automatic Restart of Some Processes	22
Customizing the TDV Server Startup Scripts on Windows	23
Starting TDV Processes on UNIX	24
Starting and Stopping TDV Processes on Unix	25
Setting TDV Server to Start Automatically on UNIX	28
Removing TDV Service Files on UNIX	29
Customizing the TDV Server Startup Scripts on UNIX	30
Working with the Repository Utility	32
Configuring the Java Keystore File	35
Understanding TDV User Templates and Rights	36
Group and User Rights Templates	37
Summary of TDV Rights and Restrictions	38
Changing the Repository Password	43
Managing the Repository Metadata Table Size	44
Validating TDV Software License Compliance and Asset Management	45
Tips for Configuring the Number of TDV Processors	46
TDV Logging Information	48
About TDV Log Files	48
Installation and Uninstallation Logs	48

Server, Monitor, and Studio Log Files	50
Configuring Email Alerts for TDV Events or Actions	54
Configuring and Enabling Event Logging	57
Enabling Logging of System Events	59
Enabling SNMP Traps in TDV	60
About the SNMP Server Events in TDV	61
Configuring an SNMP Trap Receiver	65
Events that Can Be Sent to Custom Event Handlers	66
Enabling Recording of Data Source Usage in the Events Log	72
Customizing Audit Log File Behavior	74
Adjusting Time Limits for Request Events	78
Determining Data Source Type and Version Information	79
Logging Query Execution Statistics	80
Log File Collection for Support	81
How Log File Collection Works	81
About the System Information Files	82
Security with Log File Collection	83
Saving Log Files to Support	83
Using TDV Log Files to Track Resource Privilege Changes	85
Validating TDV Software License Compliance and Asset Management	86
Determining Your TDV Software License Conformance	87
Tips for Configuring the Number of TDV Processors	89
Logging Tips from an Expert	90
Using Detailed Logging	90
Controlling the Size of Files Specified in log4j.properties	91
Hiding Enabling Error Details	91
Configuring TDV Data Connections	92
Installing and Using Preconfigured JDBC Drivers	92
Using the ODBC Driver on Windows	93
Supported ODBC Data Types	93
Adding ODBC Data Sources on Windows	94

Using the ODBC Driver on UNIX	96
Setting the ODBC Environment Variables on UNIX	96
Creating a DSN with driverConfig on UNIX	98
Connecting SAS System to TDV ODBC	101
Configuring TDV for Using a JMS Broker	102
Configure Communications between TDV and the JMS Broker	103
Adding JMS Connectors to the TDV Server	103
Configuring TDV for AIX Platforms	106
Improving Studio Response Times for AIX Connections	107
System Monitoring with Studio Manager	108
Studio Manager Window and Toolbar Overview	108
Using Studio Manager	110
Launching Studio Manager	111
Selecting Columns for Display	112
Viewing Table Row Details	113
Sorting Rows	113
Customizing Filters for Studio Manager	114
Configuring the Columns on the Cached Resources Panel	115
Enabling and Disabling Caches in Studio Manager	116
Modifying the Cache Schedule in Studio Manager	116
Refreshing a Cache in Studio Manager	117
How to Troubleshoot Cache Refresh	117
Configuring Time for Requests to Stay Active on the Studio Manager Request Panel	121
Scheduling Data Source Connection Testing	122
Studio Manager UI Reference	123
Server Overview Panel	123
Cached Resources Panel	125
Data Sources Panel	125
Events Panel	127
I/O Panel	127

Memory Panel	128
Requests Panel	129
Sessions Panel	130
Storage Panel	130
Transactions Panel	131
Triggers Panel	131
System Management with Manager	133
Using Manager	134
Launching Manager	134
Refresh the Current Page	136
Sort with Manager	136
Filter Table Data	137
Creating a New Table Filter	137
Copying an Existing Filter	138
Manager UI Reference	139
MANAGER HOME Page	139
SERVER OVERVIEW Page	142
CACHED RESOURCES Page	145
DATA SOURCES Page	148
REQUESTS page	151
SESSIONS page	155
TRANSACTIONS Page	157
TRIGGERS page	160
TDV Configuration Options	163
Fine Tuning Memory	163
About Paging	164
Configuring the Caching and Data Processing Directory	164
Viewing Usage and Cleaning Up Memory	165
Changing Default Memory Settings	165
Fine Tuning Performance Using Connection Pools	166

Using Pass-through Optimization with Oracle Data Source Clients	167
Using Pass-through Optimization with SQL Server Data Source Clients	168
Using Pass-Through Introspection with Vertica Data Source Clients	169
Enabling Studio Locking	169
Configuring Case Sensitivity and Trailing Spaces Settings	170
Determine Whether Case or Trailing Space Settings Affect Query Performance	171
Setting Server-wide Case and Trailing Space Behavior Using Configuration Parameters	173
Setting Session-wide Case and Trailing Space Behavior Using Connection Properties	175
Configuring Case and Trailing Space Behavior for Built-in Procedures	175
Configuring Case and Trailing Space Behavior for Queries	176
Using STRICT to Control Case and Trailing Space Behavior for Queries	177
Mismatch Effects on String Comparisons	177
Function Overrides	178
LPAD and RPAD return length	178
Management of Data Source Customization	179
Password Storage Options	180
Customizing the Login Screen Default Domain Value	180
Custom Server Settings For External Tools	181
Composite Domain Administration	182
About the Composite Domain	182
About Domain Management	183
Group Management	184
Built-in Groups	184
Adding Groups to the Composite Domain	185
Removing Groups	186
User Management	187
Built-in Users and Their Privileges	188
Adding Users to the Composite Domain	189
Removing Users from the Composite Domain	191

Auditing User Access to TDV Defined Resources	191
Managing Group Membership	193
Viewing Group Membership	194
Editing Group Membership	195
Changing Passwords for Other Composite Domain Users	196
Changing Ownership of Resources	197
Manage User and Group Privileges	198
LDAP Domain Administration	199
About the LDAP Domain	199
LDAP Domain - Active Directory 2003 Limitation	199
Configure the LDAP Properties File	200
Structure of the LDAP Properties File	200
Example of an ldap.properties File	203
LDAP Properties File Symbols and Attributes	206
Query Examples	208
LDAP Domain Administration	212
About Kerberos Configuration Files and LDAP Login Credentials	213
Adding an LDAP Domain	214
Working with Groups from an LDAP domain	217
Editing LDAP Domain Connection Parameters	220
Removing an LDAP Domain	221
LDAP User Management	221
Adding Users to TDV from an LDAP Domain	222
Remove LDAP Users from TDV	222
Add Users to Groups	223
Configuring LDAP for Use with Certificate Authentication	224
Configuring LDAP for Use with Nested Groups	224
Azure Domain Administration	226
About the Azure Domain	226
Azure Domain Administration	226

Adding an Azure Domain	226
Azure Group Management	227
Editing Azure Domain Connection Parameters	227
Removing an Azure Domain	228
Dynamic Domain Administration	229
About Dynamic Domains	229
About Dynamic Domain Administration	230
Enabling the Dynamic Domain	231
About Group Administration for Dynamic Domains	231
Considerations for Granting Privileges to Dynamic Domain Users	231
Viewing Dynamic User Names	232
About User Administration	233
Adding Users to the Dynamic Domain	234
Remove Users from the Dynamic Domain	234
Dynamic Users Group Membership	234
Viewing Dynamic User Group Membership	235
OAuth Domain Administration	236
About the OAuth2 Domain	236
OAuth2 Architecture	236
OAuth Supported Data Sources	238
OAuth2 Domain Administration	239
Adding an OAuth2 Domain	240
TDV to IdP Field Mapping	244
Working with Claims from an OAuth2 domain	248
Editing OAuth2 Domain Connection Parameters	252
Removing an OAuth2 Domain	252
Defining Security Policies for Claims	253
Hybrid OAuth2 Domain	253
Claim Management for Hybrid OAuth2 Domains	256

TDV and SSL Authentication	258
Overview of TDV and SSL	258
Keystore and Truststore Files for TDV	259
Keys Passed between System Components	259
Default Locations of Keystore and Truststore Files	260
Keystore and Truststore Configuration Parameters	261
Setting Which Protocols to Disable When Creating an SSL Connector	263
Setting Up SSL	263
Using the Keytool Utility	264
Installing a Truststore Certificate	264
Setting Up Authentication between Studio and the TDV Server	267
Setting Up Authentication between Client Applications and TDV Server	269
Creating a JDBC Client Application with SSL Capability	270
Setting Up Authentication between Client Applications and TDV Server over JDBC	272
Setting Up Client Authentication for Web Data Sources	274
Setting Up Client Authentication for Relational Database Sources	275
SSL Mutual Authentication	276
SSL Mutual Authentication Using Certificate	276
Example - How to Obtain a third-party SSL Certificate and install into your Server and Studio Truststore?	278
Configuring Kerberos Single Sign-On	281
About Kerberos Authentication and TDV	281
Supported Platforms and Requirements for Kerberos	282
Data Sources and Kerberos	282
Using Kerberos Authentication with TDV	283
Configuring Kerberos for Use with TDV	283
Configuring TDV for Use with Kerberos Authentication	287
Setting Up SSPI Kerberos SSO	289
Setting Up JGSS Kerberos SSO	292
About Studio and SSO with Remote Desktop	298

Using Kerberos Authentication with Published Resources	298
Configuring New Web Services for Kerberos Authentication	298
Verifying Kerberos for an OData Data Service	301
Using Kerberos SSO Authentication with Data Sources	302
About Configuring Kerberos SSO for Data Sources	302
About JDBC Clients and Kerberos SSO	305
Setting the DSN for ODBC Clients and Kerberos SSO	312
About ADO.NET Clients and Kerberos SSO	313
Configuring Kerberos with Hive and Impala Data Sources	314
Tip from an Expert on SSO Connection Issues	315
Kerberos Test Utility	315
Troubleshooting Tips	317
Managing Security for TDV Resources	319
Overview of TDV Security Features	319
Summary of Password Encryption and Security in TDV	320
Summary of Internet Security Options	321
Rights and Privileges	321
Resource Rights	321
Resource Privileges	324
Configuring Account Security for TDV	340
Configuring Account Lockout for TDV	340
Locking and Unlocking TDV User Accounts	341
Setting IP Restrictions	341
Row-Based Security	342
About Row-Based Security	342
Enabling Row-Based Security on TDV Resources	345
Creating or Editing Row Filter Policies	346
Creating or Editing a Row Filter Policy Group	351
Assigning Row Filter Policies to TDV Resources	353
Testing the Security Filter Policies Within Studio	353
Encryption Settings for TDV Server	355

Defining or Editing Encryption to Protect TDV Server Data	356
Backup of Encryption Settings	357
Exporting or Importing an Encryption Settings File From Manager	357
Configuring Pass-Through Security for HiveServer2	359
Configuring Samba and Winbind for NTLM (Tips from an Expert)	359
Using Version Control and TDV	360
About Version Control (VCS) for TDV	360
Configuring Version Control for TDV Resources	362
Creating a VCS connection with SSL Authentication	365
Change User Credentials for TDV Version Control	366
Attach a TDV Folder to a VCS Instance	366
Committing TDV Resources to the VCS	367
View the History of a Resource	368
View the Full History of a Resource	368
Compare a Resource with Local	369
Revert Changes to a Resource	369
Checking In a Resource to the VCS	370
Checking In Multiple Resources to the VCS	370
Detach a VCS Folder from Your TDV Instance	371
TDV Data Directory Management	371
Manage Connections	371
Managing Column-Based Security	373
About Column-Based Data Obfuscation	373
Column-Based Restrictions and Privileges	373
Enabling Column-Based Security on TDV Resources	374
Creating or Editing Column Filter Policies	375
Mapping Column Filter Policies to TDV Resources	379
Testing the Column Filter Policies Within Studio	379
Importing and Exporting Column Filter Policies	380

System Event and Log Monitoring	381
Configuring Events	381
About Events	381
EVENT LOG Summary Information	382
Work with the EVENT LOG Page	382
The EVENT LOG Table	382
Event Details in Manager	383
I/O Log	384
Memory Log	384
Storage Log	384
TDV Command-Line Utilities	385
The TDV Export and Import Utilities	385
About the Backup Export Utility	386
Rights Required for the Backup Export Utility	387
Using the Backup Export Utility	387
Using the Keystore File from an Exported CAR File	392
Rules for the Backup Import Utility	392
Rights Required for the Backup Import Utility	393
Using the Backup Import Utility	393
Modifying the PostgreSQL Repository Maximum Allowed Packets	402
The TDV Package Import Utility	403
Rules for the Package Import Utility	403
Restrictions for the Package Import Utility	404
Using the Package Import Utility	404
The TDV Package Export Utility	417
The TDV Server Utility Program	427
Using the Server Utility	428
Server_util.sh Examples	431
Using the TDV Server Heap Dump Utility Program	432
Deployment Manager	433

About Deployment Manager	433
Deployment Manager Architecture	434
Limitations	436
Basic Deployment Manager Concepts and Definitions	436
User Roles and Workflows	437
Starting Deployment Manager	438
Defining Sites	440
Accessing Sites	441
Adding a New Site	442
Editing Site Properties	443
Refreshing Site Resources	444
Deleting a Site	445
Defining Resource Bundles	446
Defining Principal Bundles	454
Defining Mappings	458
Defining Deployment Plans	471
Accessing Plans	472
Creating a New Deployment Plan	473
Defining the Resource Bundle to Migrate	473
Defining the Principal Bundles to Migrate	477
Removing Resources from a Target Site	478
Removing Principals from a Target Site	479
Extending Deployments with Procedure Call Operations	480
Previewing a Deployment Plan	481
Editing Deployment Plan Properties	482
Refreshing Deployment Plans	483
Exporting a Deployment Plan	484
Deleting a Deployment Plan	485
Executing Deployment Plans	485
Executing a Deployment Plan	486
Executing a Deployment Plan Remotely	487
Viewing Source Site Updates Since the Last Plan Execution	488

Viewing the Execution Log Results	489
Purging the Execution Logs	489
Importing and Executing a Deployment Plan	490
Backing Up and Restoring the Deployment Manager Server	491
Backing Up the Deployment Manager	491
Restoring the Deployment Manager Server	492
Using TDV Workload Management	495
Limitations	496
About Rule Precedence	496
Setting Up and Configuring Workload Management	498
Viewing Your Workload Management Rules in Studio	501
Viewing Rule Precedence	502
Testing Workload Rules	502
Configuring Email Alerts for Workload Management	503
Request Limit Rule using Workload Management	505
Example - Creating a trigger to limit memory use using Workload Management	509
Configuring NTLM Authentication	510
NTLM Authentication and TDV	510
Implementing NTLM Authentication for Windows	511
Configuring TDV as the Server	512
Configuring SQL Server in the TDV Server to use Windows Authentication	514
Configuring TDV as the Client	515
Implementing NTLM Authentication for UNIX	522
Verifying NTLM for an OData Data Service	525
Pluggable Authentication Modules	528
About Pluggable Authentication Modules	529
Minimum Elements of a PAM	530
Working with TDV and PAM	531

The Manifest File	532
Creating a Principal Authentication Module	533
Deploying Pluggable Authentication Modules	535
What Happens at Deployment and Run Time	536
Verifying that the PAM Deployed	536
Enabling PAM Features for TDV	538
Adding a Module	539
Ordering Module Execution Sequence	540
Assigning Users to TDV Groups or Identities	540
Undeploying Pluggable Authentication Modules	542
Example	543
Collecting TDV and Data Usage Metrics	548
About Data Usage Metrics	548
Setting Up and Configuring Metrics Collection	549
Pre-Creating the External Database and Tables for Metrics Data Storage	549
Using Studio to Create the Database and Tables for Metrics Data Storage	556
Configuring TDV Metrics Collection	557
Configuring Email Alerts for Metrics Notification	559
Publishing and Reporting on TDV Metrics Data	561
About Using MAXMEMORY in Your Reports	562
SNMP Trap Message Reference	563
SNMP Details for Monitor Events	564
SNMP Details for Server Events	565
SNMP Details for Requests	567
SNMP Details for Transactions	568
SNMP Details for Cached Resources	569
SNMP Details for Triggers	570
SNMP Details for Data Sources	571
SNMP Details for Sessions	575
SNMP Details for Resources	576

SNMP Details for Storage	578
SNMP Details for Server Events	578
SNMP Details for Security	581
SNMP Details for Workload	582
SNMP Details for KPI	583
SNMP Details for Audit Startup and Shutdown	585
TDV Event Log Message Reference	587
TDV Trap Messages and Variables	587
TDV Monitor Events	587
TDV Server Events	588
Active Cluster Events	602
TIBCO Documentation and Support Services	604
Legal and Third-Party Notices	607

Basic TDV Administration Tasks

This guide describes the configuration and administration tasks for TIBCO® Data Virtualization (TDV). After you finish using the installation guides to install your TDV products, you must use this guide to customize and configure your TDV environment. This guide discusses basic tasks, configuration parameters, connection configuration, system monitoring, and repository setup options.

The following topics are covered:

- [Overview of TDV Administration](#)
- [Security Features](#)
- [Starting TDV Processes on Windows](#)
- [Starting TDV Processes on UNIX](#)
- [Working with the Repository Utility](#)
- [Configuring the Java Keystore File](#)
- [Understanding TDV User Templates and Rights](#)
- [Changing the Repository Password](#)
- [Managing the Repository Metadata Table Size](#)
- [Validating TDV Software License Compliance and Asset Management](#)

Overview of TDV Administration

TDV forms the core of the Data Virtualization Platform. TDV can use data from many different data sources, recombine it, and make it available to many types of clients. The administrator needs to manage data sources, multiple connection and security protocols, complex data access methods, multi-level security, and internal repositories.

You need to perform most administration tasks periodically, but not necessarily daily. The table lists the main administration tasks and where to find information about them.

Administration Task	For More Information
Installing and setting up TDV	Installation details are in the <i>TDV Installation and Upgrade Guide</i> .
Starting and stopping the TDV Server, testing the TDV repository, and dealing with the Java keystore file that is enabled by your TDV installation.	Starting TDV Processes on Windows Starting TDV Processes on UNIX Working with the Repository Utility Configuring the Java Keystore File
Configuring access protocols, including ODBC, JDBC, Hive, and ADO.NET for data sources and client interfaces.	Configuring TDV Data Connections Refer to “Working with Data Sources” and “Configuring Relational Data Sources” in the <i>TDV User Guide</i> . Also refer to the <i>TDV Client Interfaces Guide</i> .
Performing security-related tasks associated with Studio use at the domain, group, or user level.	Configuring the Java Keystore File Group and User Rights Templates Summary of TDV Rights and Restrictions Overview of TDV Security Features Resource Rights Resource Privileges Row-Based Security Configuring Kerberos Single Sign-On Composite Domain Administration LDAP Domain Administration Dynamic Domain Administration Configuring NTLM Authentication
Tuning and configuring TDV behavior.	TDV Configuration Options Fine Tuning Memory

Administration Task	For More Information
Although this guide explains many configuration parameters, others are explained in other documents alongside the features whose behavior they modify.	Enabling Studio Locking Configuring Case Sensitivity and Trailing Spaces Settings Also refer to “Performance Tuning” in the <i>TDV User Guide</i> . For a complete list of configuration parameters, refer to the <i>TDV Reference Guide</i> .
Monitoring your TDV system.	System Monitoring with Studio Manager System Management with Manager System Event and Log Monitoring
Setting up and testing TDV repositories.	Working with the Repository Utility
Importing and exporting Studio metadata and resources.	The TDV Export and Import Utilities

Security Features

Security features are discussed in the table in [Overview of TDV Administration](#), and also in these topics:

- [Managing Security for TDV Resources](#)
- [Configuring Kerberos Single Sign-On](#)
- [TDV Command-Line Utilities](#)

Starting TDV Processes on Windows

By default, a Windows installation of TDV creates and registers two TDV services that are set to restart automatically whenever the host machine is restarted. For example,

after a manual computer restart. Typically, the TDV Server runs on the host server and is ready to accept connections for design or runtime requests.

The server and the server repository service use the following naming convention:

- TDV Server <x.y.z>
 - TDV Repository <x.y.z>
- x = major version, y = minor version, and z = service pack number

The TDV Server process runs the TDV Server with the Monitor Daemon option enabled. The Monitor process tracks performance and maintains audit information. Typically, you only run TDV without the Monitor Daemon if you have been prompted to do so by Support.

Note: If you installed TDV on Microsoft Windows Vista Business Edition or Windows 7, see the *TDV Installation and Upgrade Guide* for instructions on preparing Windows.

Depending on what you want to do in your OS environment, see the following sections:

- [Starting and Stopping TDV Processes on Windows](#)
- [Keeping the TDV Server Process Running after Logoff](#)
- [Turning Off Automatic Restart of Some Processes](#)
- [Customizing the TDV Server Startup Scripts on Windows](#)

Starting and Stopping TDV Processes on Windows

You can start and stop a variety of TDV processes.

To start or stop TDV on Windows, use one of the following

- From the Windows Task Manager Services tab, start or stop the TDV Server and Repository processes.
- In a command window, navigate to the TDV installation directory using the command:

```
cd <TDV_install_dir>\bin
```

Enter one of the following.

To start, stop, or restart all TDV processes	Enter the composite command with one of these options: <pre>composite.bat monitor [start stop restart]</pre>
To start the TDV Server and repository processes without the monitor	Enter the command: <pre>composite_server.bat run</pre> This command runs a Java process only; no Windows service is called. Or, use the Server Auto Restart configuration parameter.
To start, stop, restart, install, or uninstall the repository	Enter the command: <pre>composite.bat repo [start stop restart install uninstall]</pre> Note: The stop command on UNIX shuts the repository down, but leaves the monitor process running. On Windows, the stop command stops both the repository and the monitor process.

Tips From an Expert on the Sequence to Start/Stop TDV Processes on Windows

- In general, to stop the TDV processes - Stop the TDV Server first, followed by TDV Repo and then TDV Cache. When Starting the TDV processes, start the TDV Cache first, followed by TDV Repo and then TDV Server.

Run these scripts in the given sequence:

Windows (TDV stop):

```
composite.bat monitor stop
```

```
composite.bat repo stop
```

```
composite.bat cache stop
```

Windows (TDV stop):

```
composite.bat cache start  
  
composite.bat repo start  
  
composite.bat monitor start
```

Note: If you stop cache before stopping the server, it is possible that the server could still try to access the cache or repo and you may encounter errors.

- To stop TDV Service in Windows, open the Windows services and stop the Windows Service "TDV Server <version>" followed by "TDV Repository <version>" and then the "TDV Database Cache <version>". To start the TDV Service, open the Windows services and start the Service "TDV Server <version>"

Keeping the TDV Server Process Running after Logoff

On Windows, if you plan to run the server process manually, you can keep the server process running even after you log off the system.

To keep the server process running after you log off of TDV

1. Go to Administration > Configuration to open the configuration window, and navigate to Server > Configuration > Monitor > Server Ignore Signals (On Monitor Restart).
2. Set the Value to True and restart the monitor.

Turning Off Automatic Restart of Some Processes

Turning the automatic restart of the monitor process off can help diagnose problems with the server restart.

To turn off the automatic restart of the monitor process

1. Go to Administration > Configuration to open the configuration window, and find the Server Auto Restart configuration parameter.

2. Set the Value to False.

The value is effective immediately; a TDV Server restart is not required.

3. Continue with the activities necessary to diagnose the issues with your system.

Customizing the TDV Server Startup Scripts on Windows

The TDV installation provides a startup script for Windows that you can customize for your own purposes; but to maintain customizations you might have made to this script across hotfix or patch updates, you must activate an environment variable.

Note: You can also use this functionality to add JRE VM arguments or run commands automatically before launching TDV.

The procedure below describes how to configure the startup script for Windows.

To configure the TDV startup scripts for Windows

1. Stop the TDV Server.
2. In Windows Explorer, navigate to <TDV_install_dir>\conf.
3. Copy script_env.bat.sample to script_env.bat.
4. Open script_env.bat with a text editor and uncomment (remove “rem” from) the line that contains:

```
rem set CIS_SERVER_VM_ARGS
```

5. Change the value of CIS_SERVER_VM_ARGS to include all the values in VM_ARGS from <TDV_install_dir>\conf\server\server.properties.

Retrieve your platform specific value of VM_ARGS from <TDV_install_dir>\conf\server\server.properties. Find the first “set VM_ARGS” line in that script and use the values from that line.

If you want to run an executable or command-line utility at this point, make sure the command returns control to composite_server.bat. If it does not, TDV does not start correctly.

Each command you add should be on a new line.

6. Start the TDV Server.

7. Review the newest <TDV_install_dir>\logs\cs_server.out.<timestamp> file to make sure the script environment is functioning properly.

Examples

CIS_SERVER_VM_ARGS without quotes and no spaces before or after the equal sign (=):

```
set CIS_SERVER_VM_ARGS=-server -XX:NewRatio=6 -XX:-UseGCOverheadLimit -
XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=C:/TDV702/logs -
XX:PermSize=64m -XX:MaxPermSize=256m -XX:-ReduceInitialCardMarks -
XX:+ExplicitGCInvokesConcurrent -XX:+UseConcMarkSweepGC -
Duser.timezone=GMT
```

CIS_SERVER_VM_ARGS with a space before and after the equal sign (=) when double quotes are used around the value:

```
set CIS_SERVER_VM_ARGS ="-server -XX:NewRatio=6 -XX:-UseGCOverheadLimit -
XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=C:/TDV702/logs -
XX:PermSize=64m -XX:MaxPermSize=256m -XX:-ReduceInitialCardMarks -
XX:+ExplicitGCInvokesConcurrent -XX:+UseConcMarkSweepGC -
Duser.timezone=GMT"
```

Starting TDV Processes on UNIX

The UNIX platform does not install any automatic services to start and stop TDV and the repository processes. All UNIX shell scripts run under the Bourne shell (/bin/sh).

The Monitor process tracks performance and maintains audit information. Typically, you would run without the Monitor only if you have been prompted to do so by Support.

Depending on what you want to do in your environment, see the following sections:

- [Starting TDV Server on UNIX](#)
- [Setting TDV Server to Start Automatically on UNIX](#)
- [Removing TDV Service Files on UNIX](#)
- [Starting, Stopping, or Restarting the Repository on UNIX](#)
- [Starting, Stopping, or Restarting the Cache Database on UNIX](#)
- [Starting TDV without the Monitor on UNIX](#)

- [Customizing the TDV Server Startup Scripts on UNIX](#)

Starting and Stopping TDV Processes on Unix

You can start and stop a variety of TDV processes.

The UNIX platform does not install any automatic services to start and stop TDV and the repository processes. All UNIX shell scripts run under the Bourne shell (/bin/sh).

The Monitor process tracks performance and maintains audit information. Typically, you would run without the Monitor only if you have been prompted to do so by Support.

Starting TDV Server on UNIX

To start TDV and the Monitor at the same time for a UNIX session, use the following procedure. If you plan to run the server process manually and want to keep the server process running even after you log off the system, you need to run the following command:

```
nohup <command>
```

The command argument is one of those listed in the procedure below.

To start TDV and Monitor on UNIX

1. Log in as the user who installed TDV. This should be a user that is not the root user.
2. Navigate to the <TDV_install_dir>/bin directory.

```
cd <TDV_install_dir>/bin
```

3. Type the following command:

```
composite.sh monitor [start | stop | restart] [-user <user_name> -password <password>]
```

Starting, Stopping, or Restarting the Cache Database on UNIX

This command is only supported for a cache database that was installed during the installation of TDV Server.

To start, stop, or restart the cache database

1. Login as the user who installed TDV. This should be a user who is not the root user.
2. Navigate to the TDV installation directory.

```
cd <TDV_install_dir>/bin
```

3. Type the following command:

```
composite.sh cache [start | stop | restart]
```

Starting, Stopping, or Restarting the Repository on UNIX

This command is only supported for a repository that was installed during the installation of TDV Server.

To start, stop, or restart the repository

1. Login as the user who installed TDV. This should be a user who is not the root user.
2. Navigate to the TDV installation directory.

```
cd <TDV_install_dir>/bin
```

3. Type the following command:

```
composite.sh repo [start | stop | restart]
```

Starting TDV without the Monitor on UNIX

Typically, you only run without the Monitor if you have been prompted to do so by Support. You can start TDV for your current UNIX session without starting the Monitor process, but it is not recommended. These actions only output to log files. If you run the server with no Monitor, a Monitor stops that server process and restarts a new one in the background.

To start TDV without the Monitor on UNIX

4. Login as the user who installed TDV. This should be a user that is not the root user.
5. Navigate to the TDV installation directory.

```
cd <TDV_install_dir>/bin
```

6. Type the following command:

```
composite_server.sh run
```

Tips From an Expert on the Sequence to Start/Stop TDV Processes on Unix

- In general, to stop the TDV processes - Stop the TDV Server first, followed by TDV Repo and then TDV Cache. When Starting the TDV processes, start the TDV Cache first, followed by TDV Repo and then TDV Server.

Run these scripts in the given sequence:

UNIX(TDV stop):

```
composite.sh monitor stop
```

```
composite.sh repo stop
```

```
composite.sh cache stop
```

UNIX (TDV start):

```
composite.sh cache start
```

```
composite.sh repo start  
  
composite.sh monitor start
```

Note: If you stop cache before stopping the server, it is possible that the server could still try to access the cache or repo and you may encounter errors.

Setting TDV Server to Start Automatically on UNIX

If after installing the software you restart the UNIX installation machine, TDV Server and the metadata repository do **not** start automatically (unlike when they start automatically after a successful installation of the software). To configure them to start automatically upon UNIX restart, use the following procedure.

To configure the TDV service files `cis.repository` and `cis.server`

1. Log into the installation machine as root.
2. Navigate to the TDV installation directory:

```
cd <TDV_install_dir>/bin
```

3. Run the following command as the root user:

```
cis_install_services.sh
```

This command prompts for a user name and other details needed to install and configure the service files `cis.repository` and `cis.server`.

4. Enter the name of the user who starts TDV (not the root user) and the other information requested.

The script then installs `cis.repository` and `cis.server` into an appropriate location on the installation machine and configures them. The location is displayed on your screen when configuration is successful, so **make note of this location**, because you need it in the verification step below.

Note: Do not run the `cis.repository` or `cis.server` scripts directly. These are template files for `cis_install_services.sh` only.

Running `cis_install_services.sh` does not interrupt any repository or server processes, but prepares the machine to start these processes automatically when a UNIX computer is restarted.

5. Run the following commands as the root user:

```
cd <init_directory>  
  
chmod 550 cis.repository  
  
chmod 550 cis.server
```

To verify the TDV service files configuration

The value of `init_directory` depends on the operating system:

- Linux: `/etc/rc.d/init.d` or `/etc/rc.d`
- AIX: `/etc/rc.d/init.d`

1. Go to the location noted previously from running `cis_install_services.sh`.

Note: The console output of the script `cis_install_services.sh` displays the exact location.

2. Enter these commands:

```
./cis.repository restart  
  
./cis.server restart  
  
./cis.cache restart
```

Now if the machine is rebooted, the monitor, server, and repository processes should automatically start once the machine is ready to go.

Removing TDV Service Files on UNIX

You can use the `cis_remove_services.sh` script from a command line to uninstall the TDV services files that are used to restart the server and repository automatically on UNIX.

This command does not interrupt any repository or server processes that are running, but removes the TDV service files.

To remove the TDV service files

1. Log onto the installation machine as root.
2. Navigate to the TDV installation directory.

```
cd <TDV_install_dir>/bin
```

3. Run the following command:

```
cis_remove_services.sh
```

Customizing the TDV Server Startup Scripts on UNIX

The TDV installation provides a startup script for UNIX that you can customize for your own purpose; but to maintain customizations you might have made to this script across hotfix or patch updates, you must activate an environment variable.

You can also use this functionality to add JRE VM arguments or run commands automatically before launching TDV.

The procedure below describes how to configure the startup script for UNIX.

To configure the TDV startup scripts for UNIX

1. Stop the TDV Server.
2. Navigate to conf under the TDV installation directory.

```
cd <TDV_install_dir>/conf
```

3. Copy script_env.sh.sample to script_env.sh.
4. Open script_env.sh with a text editor and uncomment the last two lines:

```
# CIS_SERVER_VM_ARGS=
```

```
# export CIS_SERVER_VM_ARGS
```

5. Change the value of CIS_SERVER_VM_ARGS to include all the values in VM_ARGS from <TDV_install_dir>/conf/server/server.properties

Retrieve the platform-specific value of VM_ARGS from <TDV_install_dir>/conf/server/server.properties. Locate the line for your platform:

- SunOS and Linux platforms should use the VM_ARGS definition on the "Linux"|"SunOS" line.
- AIX platform should use the VM_ARGS definition on the "AIX" line.

Make sure you add double-quotes around the value specified for CIS_SERVER_VM_ARGS.

If you want to run an executable or command-line utility at this point, make sure the command returns control to composite_server.sh. If it does not, TDV does not start correctly.

Each command you add should be on a new line.

6. Start the TDV Server.
7. Check the end of the newest <TDV_install_dir>/logs/cs_server.out.<timestamp> file to ensure the script environment functionality is working.

In the following example of a cs_server.out.<timestamp> file, script_env.sh added a new JRE VM option "-Dtest=false" and ran script commands to print the contents of a directory.

The output shows you the “before” and “after” VM_ARGS settings so you can see what is being used in TDV.

```
> Tue Mar 20 08:19:01 PDT 2012
> Detected /opt/<installdir>/conf/script_env.sh, sourcing it...
>
> ls logs
> cluster
> cs_server.out.20120320081901
>
> Done sourcing /opt/<installdir>/conf/script_env.sh
> Default VM_ARGS: -server -XX:NewRatio=6 -XX:-UseGCOverheadLimit -
XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/opt/<installdir>/logs -
XX:PermSize=64m -XX:MaxPermSize=256m -XX:-ReduceInitialCardMarks -
XX:+ExplicitGCInvokesConcurrent -XX:+UseConcMarkSweepGC
> Detected CIS_SERVER_VM_ARGS environment variable overrides VM_ARGS
> Changing VM_ARGS to: -server -XX:NewRatio=6 -XX:-UseGCOverheadLimit -
XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/opt/<installdir>/logs -
XX:PermSize=64m -XX:MaxPermSize=256m -XX:-ReduceInitialCardMarks -
XX:+ExplicitGCInvokesConcurrent -XX:+UseConcMarkSweepGC -Dtest=false
> chmod 755 "/opt/<installdir>/bin/init_patch_script.sh" exit code=0
```

```
> "/opt/<installdir>/bin/init_patch_script.sh"
> "/opt/<installdir>" "/opt/<installdir>" exit code=0
> Starting the TDV Server process
```

Working with the Repository Utility

Use `repo_util` to change the repository database. The `repo_util` scripts (`repo_util.bat` and `repo_util.sh`) are available in the `<TDV_install_dir>/bin` directory.

You can use this program to perform the following tasks:

- Test the connection to the repository database
- List the current repository configuration information
- Export the repository configuration
- Update the repository configuration
- Create or drop the repository schema
- Print diagnostic information about the TDV metadata repository

Syntax

```
repo_util.bat
< -createSchema | -dropSchema | -dumpDiagnosticInfo | -exportConfig | -help |
-listConfig | -testConn | -updateConfig >

[ -debug | -force ]

[ -configFile | -connectionUrl | -databaseCatalog | -databaseSchema |
-databasePassword | -databaseUser | -driverClass | -driverClassPath | -driverName |
-driverType | -repositoryClass | -schemaCreateScript | -schemaDropScript |
-schemaInitializeScript ]
```

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.

Options	Description
	Run the repository utility with the <code>-exportConfig</code> option for an example of this file's contents.
<code>-connectionUrl</code>	The JDBC URL that is used to connect to the external database. For example: <pre>jdbc:PostgreSQL://localhost:3406/cs030101?continueBatchOnError=false&useUnicode=true</pre>
<code>-createSchema</code>	Create the repository schema.
<code>-databaseCatalog</code>	Catalog that contains the TDV schema; blank if catalogs not supported.
<code>-databasePassword</code>	Database password.
<code>-databaseSchema</code>	Database schema that contains the TDV schema; blank if schemas not supported.
<code>-databaseUser</code>	Database user name.
<code>-debug</code>	Print debug messages.
<code>-driverClass</code>	Fully-qualified class name of a JDBC-compliant driver. For example: <code>com.PostgreSQL.jdbc.Driver</code> .
<code>-driverClassPath</code>	A semicolon- or colon-separated list of JAR files and directories. For example: <code>/tmp/oracle40.jar:/tmp</code> .
<code>-driverName</code>	Name of the TDV datasource driver name. Required for operation of the system tables.
<code>-driverType</code>	Name of the TDV data source driver type. Required for operation of the system tables.
<code>-dropSchema</code>	Drops the repository schema and all data contained within it.

Options	Description
	Permanently deletes all of the server's data. (Use with caution.)
-dumpDiagnosticInfo	Print diagnostic information about the repository database.
-exportConfig	Export the repository database configuration in Java property file format. The output is suitable for use as a repository configuration file. See -configFile for details.
-force	Do not prompt for confirmation. (Use with caution.)
-help	Print this help information.
-listConfig	List the repository database configuration in a human readable format.
-repositoryClass	Repository class name.
-schemaCreateScript	Script containing the SQL commands to create the TDV schema.
-schemaDropScript	Script containing the SQL commands to drop the TDV schema.
-schemaInitializeScript	Script containing the SQL commands to initialize the TDV tables.
-testConn	Test the connection to the repository database.
-updateConfig	Change options in the repository database configuration. Specify new configuration options individually using command-line arguments, or collectively using the -configFile option. Unspecified options are left unchanged.

Sample Uses of the Repository Utility

Here are some uses of 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 <user>  
-databasePassword <password>
```

- To update the repository configuration using a repository configuration file, overriding the database password:

```
repo_util.bat -updateConfig -configFile repo.properties  
-databasePassword <password>
```

Configuring the Java Keystore File

TDV includes a generic Java KeyStore (JKS) file so that you can use it for development and testing of Web services and for JDBC secured over HTTPS ports. If you plan to build secure Java programs, it is recommended that all TDV instances be configured with their own JKS certificate prior to deployment.

You need to configure the JKS digital certificate that you intend to use for secured Web services and secured JDBC communications. The JKS digital certificate initiates and establishes SSL communication over HTTPS and LDAP ports.

You need Read All Resources and Modify All Resources rights to change the JKS digital certificate file location, file type, or the password.

If a trust store location is not specified, a keystore file is searched for in the following locations:

- \$JAVA_HOME/lib/security/jssecacerts
- \$JAVA_HOME/lib/security/cacerts

To configure the JKS digital certificate for TDV using the SSL MANAGEMENT page

1. Obtain a JKS digital certificate from a Certificate Authority (CA), or generate your own using the keytool command-line utility available in the following directory:

```
<TDV_install_dir>/jdk/bin
```

2. Open Studio, and select Administration > Launch Manager (Web) to open the TDV Manager Web interface.
3. Log in to Manager.
4. In Manager, choose CONFIGURATION > SSL to display the SSL MANAGEMENT page.
5. In the New Value column on the Java Keystore File Location page, enter the absolute path to the new JKS file on the server.

Set Server Values on Restart					
Name	Value on Restart	New Value		Details	
Java Keystore File Location	C:/Program Files/Com	/opt/Composite_Softw	APPLY	REVERT	The location of the Java keystore file used by the HTTP/S and HTTP/S With X.509 Certificate Client Authentication listeners to establish the identity of the server to external clients.
Java Keystore File Type	JKS		APPLY	REVERT	The type of the Java keystore file. It must be a valid Java keystore type such as "JKS" or "PKCS12"
Java Keystore Password	*****		APPLY	REVERT	The password of the Java keystore file and the entries within it. All password protected entries in the keystore file must use the same password as the file itself.

6. Click APPLY.

A dialog warns that the new value takes effect only after server restart. Another prompt notifies you of a successful change and then refreshes the page.

The REVERT button recovers the current value until TDV restart.

7. Change the Java Keystore File Type and the Java Keystore Password values so that the values when the server restarts match the digital certificate being installed.
8. Restart the TDV Server to load the keystore and apply the changes.

Understanding TDV User Templates and Rights

Part of customizing and configuring your TDV environment involves user templates and rights. The following sections describe them:

- [Group and User Rights Templates](#)
- [Summary of TDV Rights and Restrictions](#)

For more information about user and group rights and privileges, see [Managing Security for TDV Resources](#).

Group and User Rights Templates

When you create new groups and users in Manager, the Group Rights and User Rights templates let you assign prearranged sets of rights based on user categories. The rights granted by the templates progress from End User (with no rights) to the Administrator (with all rights). You can use these templates as a starting point, and change the rights to suit your purposes:

Template Name	Description
End User	Starts with no TDV rights; however, the user can request data through ODBC, JDBC, and Web service clients. Data is still protected by privileges set at the data source. The end user template does not include rights to use Studio or TDV application tools.
Developer	Grants access to tools, and lets the developer view all status.
Operations	Grants access to tools; lets the operations person read server configurations, and read and modify all status.
Operations Administrator	Grants access to tools; lets the operations administrator read and modify server configurations, view all status, use all Monitor functionality, and modify all server configurations.
Backup	Grants access to tools; lets the backup-user read server configurations, view all status, and read (but not write) resource and user data for backup purposes.
Restore	Grants access to tools; lets the restore-user view and modify server configurations, all resources and all users.
Backup & Restore	Grants all Backup template and Restore template rights.
KPI	Grants access to tools; lets a user managing KPI policies read and modify all resources, and read all status.

Template Name	Description
RBS and CBS	Grants access to tools; lets a user managing security policies read and modify all resources, read and modify all server configurations, and read all users.
Workload Management	Grants access to tools; lets a user managing workload policies read and modify all resources, read and modify all server configurations, and read all users.
Administrator	Grants complete access, and rights to change everything in the TDV system, except system tables that are locked to ensure system functionality. All rights are required to gain access to and use Manager.

Security of the TDV Server can disallow package export by users who are not members of a specified group. Package export can also be restricted to the resource owner, or to users with Write privilege on the resource. The admin user can perform package export for any resource.

Summary of TDV Rights and Restrictions

The rights that can be granted to a user or group include the following.

TDV Right	Description	Templates Where Right Occurs by Default
Access Tools ACCESS_TOOLS	Gives end-users access to TDV tools (like Studio), command-line utilities (like backup_import), and APIs that connect with TDV. All Administrators, Developers, IT Operations, and personnel responsible for backup and restore must have this right to view and change TDV. Having this right is	Administrator, Backup&Restore, Restore, Backup, Operations, Developer

TDV Right	Description	Templates Where Right Occurs by Default
	implicit in all discussions of access to or manipulation of TDV resources. Additional rights are required for full export or import of a TDV instance. Without this right, the user can only use JDBC, ODBC and Web Services to access the server and underlying native sources.	
Modify All Config MODIFY_ALL_CONFIG	Lets the user modify all TDV configurations, perform full-server backup and restore, write CAR files; create, join, or leave a TDV cluster; and use the Cluster_util command line utility.	Administrator, Backup&Restore, Restore
Modify All Resources MODIFY_ALL_RESOURCES	Gives full (Grant, Write, Select, Insert, Update, Delete, Execute) privileges on all resources, including the right to change privileges on any resource; change owner of a resource; import privileges; create copies of resources that retain original owner and privileges (also requires Modify All Users right); restore/import (also requires other rights). This right lets the user modify all resources and privileges on resources, and change the data source owner, even if the user has not explicitly been given privileges for that resource.	Administrator, Restore
Modify All Status MODIFY_ALL_STATUS	Lets the user perform Manager and Server Overview actions (clear pool and test all data sources); view and clear	Administrator

TDV Right	Description	Templates Where Right Occurs by Default
	data sources on the Manager panel; and synchronize domains.	
Modify All Users MODIFY_ALL_USERS	Gives the user full administrative powers: lets the user create or modify domains, groups, and users and their rights; change resource owners; import resources with associated users and privileges (also requires Modify All Resources); paste while preserving user privileges. This right can be used to grant any other rights. Making changes on the Manager - Users pages requires this right.	Administrator, Backup&Restore, Restore
Read All Config READ_ALL_CONFIG	Lets users browse TDV configuration settings by means of Studio, Manager, or a Web services operation. This includes the configuration panels. With the Read All Users right, gives view access to the Resource Management pages. Without the Read All Users right, manager-users can see only their own privileges and the privileges held by the groups to which they belong. With the Modify All Resources right, lets the user add, remove, and automatically correct dependency privilege settings. This right is appropriate for developers, although the Developer template does not include this right by default.	Administrator, Backup&Restore, Restore, Backup, Operations
Read All Resources	Lets the user view all resources; read all resources (even without explicit Read	Administrator, Backup&Restore,

TDV Right	Description	Templates Where Right Occurs by Default
READ_ALL_RESOURCES	privileges); perform full server backup; execute backup_import; use Manager panels; execute any resource procedure; browse and edit resource services. Developers are not granted this right by default with the Developer template.	Restore, Backup
Read All Status READ_ALL_STATUS	Lets the user view TDV current state, sessions, transactions, requests, caches, support diagnostics, query plan view, cluster status; view event, server, and storage logs (accessible from the Administration -> Studio Logs menu); use the -profile option with server_util; view resource tables such as SYS_CACHES and SYS_DATASOURCE. The Active Resource tables are visible to users with this right, showing sessions, transactions, requests, caches, data sources, clusters, and so on, on Manager panels. This right is useful for developer, operations, and monitoring roles.	Administrator, Backup&Restore, Backup, Operations, Developer
Read All Users READ_ALL_USERS	Lets the user browse all lists of domains, groups, and users using User Services or Manager; perform full server backup (along with Read All Resources and Read All Config); back up and restore the system (along with Read All Resources); reset the system namespace. It does not grant the ability to see any domain or user passwords.	Administrator, Backup, Restore, Backup&Restore

TDV Right	Description	Templates Where Right Occurs by Default
	Viewing the Manager - Users pages requires this right.	
	Developers are not granted this right by default with the Developer template.	
Unlock Resources UNLOCK_RESOURCE	This right is created for releasing locks set by another user, the use case is for a designer who sets locks on resources, but for some reason the lock owner is not available to release the locks when change of those resources must be made by another developer. Only the lock owner or an administrator with the UNLOCK_RESOURCE right should be able to release the lock.	Administrator
No Direct Datasource Data Access NO_DIRECT_DATASOURCE_DATA_ACCESS	This access control option is given to prevent direct data access to the data source asset. The restricted user is subject to direct data source data access restriction and cannot access the data in any of the TDV tools.	Administrator
No Direct Datasource Callable Invocation NO_DIRECT_DATASOURCE_CALLABLE_INVOCATION	This access control option is given to prevent stored procedure calls to the data source asset. The restricted user is subject to direct data source invocation restriction and cannot make stored procedure calls to the data source from any TDV tools.	Administrator

Changing the Repository Password

After installation you might periodically need to change your TDV or Business Directory repository password.

In these instructions, <install_dir> means <BD_install_dir> or <TDV_install_dir>.

To change the repository password

1. Stop the repository.
2. Locate and open the ph_hba.conf file. The file is typically at:

```
<install_dir>\repository\data\pg_hba.conf
```

3. Find and change all lines with "password" to "trust" for the METHOD column. For example:

```
TYPE DATABASE USER ADDRESS METHOD
```

"local" is for Unix domain socket connections only

```
local all all password
```

IPv4 local connections:

```
host all all 127.0.0.1/32 password
```

IPv6 local connections:

```
host all all ::1/128 password
```

4. Start the repository. For example, on Windows:

```
composite.bat repo start
```

5. Login to the PostgreSQL database using one of the following commands:

Platform	Command	Notes
Windows	<pre>./bin/psql -hlocalhost -p9508 -Uroot -dpostgres</pre>	
UNIX	<pre>cd <install_dir>/repository;</pre> <pre>export LD_LIBRARY_PATH=<install_dir>/repository/lib;</pre> <pre>./bin/psql -hlocalhost -p9508 -Uroot -dpostgres</pre>	Use SHLIB for HP-UX and LIBPATH for AIX platforms instead of LD_LIBRARY_PATH, which is only for Linux platforms.

6. Run the psql ALTER USER command.

```
postgres=# ALTER USER root with password '<NEW_DBA_PASSWORD>';
```

```
postgres=# \q
```

7. Stop the repository.

8. Locate and open the ph_hba.conf file. The file is typically at:

```
<install_dir>\repository\data\pg_hba.conf
```

9. Find and change all lines with "trust" to "password" for the METHOD column.

10. Start the repository.

11. Log in to the PostgreSQL database with the new password.

Managing the Repository Metadata Table Size

After installation you might need to change your TDV or Business Directory repository metadata table size. The Repository metadata table size will increase or decrease in size as metadata is added or deleted.

To manage the repository metadata table size

1. Stop the TDV repository.
2. Stop the TDV Server.
3. Locate Metadata Table Size configuration parameter.
4. Use the read only value to determine if you need to make a change to the value for the metadata database table.
5. Adjust the value of Metadata Cache Size (On Server Restart) based on the metadata table size.
6. Adjust block and segment sizes if necessary.
7. Restart the TDV repository.
8. Restart the TDV Server.

Garbage Batch Size

TDV garbage collection process makes use of a temp table in the Postgres repository, during the execution of INSERT/SELECT statements on the metadata tables.

To manage the size of the temp tables created, there is a configuration setting that can be used. This setting can be tuned in TDV Studio using the option Administration > Configuration > Debugging > Metadata > Garbage Batch Size. By default, this setting is turned off (defaults to 0). When set to a value greater than 0, the server will perform garbage collection on certain metadata tables in batches of size no greater than the specified size.

Validating TDV Software License Compliance and Asset Management

Diligently following application licensing compliance can prevent legal or standards infringement problems. You can gather information from the TDV log files to determine your compliance for auditing and renewal purposes.

This section contains:

- [Tips for Configuring the Number of TDV Processors](#)

Tips for Configuring the Number of TDV Processors

Configuring the number of TDV processors can help you take control of compliance to your TDV license terms. Because the environments at different companies varies so widely, you will need to research and perform testing to determine the best method for your particular environment.

Tips for configuring the number of TDV processors

1. Review documents and instructions for how to set CPU affinity.

For example, navigate to and review:

- <http://www.cyberciti.biz/tips/setting-processor-affinity-certain-task-or-process.html>
- <http://pundiramit.blogspot.com/2010/07/how-to-disable-cpu-cores-in-multicore.html>
- <http://stackoverflow.com/questions/628057/how-to-set-processor-affinity-on-an-executable-in-windows-xp>
- http://www.experts-exchange.com/OS/Unix/AIX/Q_27263123.html

2. Determine your number of available CPUs and their unique identifications.
3. Determine the names of the TDV processes that need to be associated with the specific CPUs.
4. Determine if one of the following command can help you configure your number of TDV processors. Some key commands to help you configure processors depending on operating system are as follows.

Platform	Command	Description
AIX	<code>bindprocessor 1234 1</code>	Bind the kernal threads to the process of a processor.
Windows	<code>start java startServer</code>	Modifies the startup script to provide

Platform	Command	Description
	/affinity:1,2,3,4	affinity to 4 CPUs.
Windows	imagecfg -a 0x3 <xxx>.exe	Limits the executable to CPU0 and CPU1.

5. Test your configuration changes and determine if further changes are needed.

TDV Logging Information

All events in the TDV system are logged, but not all log entries are tied to system events and visible through Manager. Also, there can be cases where an event is associated with multiple log entries.

The following topics are discussed:

- [About TDV Log Files](#)
- [Configuring Email Alerts for TDV Events or Actions](#)
- [Configuring and Enabling Event Logging](#)
- [Determining Data Source Type and Version Information](#)
- [Logging Query Execution Statistics](#)
- [Using TDV Log Files to Track Resource Privilege Changes](#)
- [Validating TDV Software License Compliance and Asset Management](#)
- [Logging Tips from an Expert](#)

About TDV Log Files

TDV uses a number of log files to store information logged during installation, uninstallation, and other system and user activities.

- [Installation and Uninstallation Logs](#)
- [Server, Monitor, and Studio Log Files](#)

Installation and Uninstallation Logs

TDV creates installation and uninstallation log files. The log files are created in the first available location listed in the following table for each OS.

Platform	Log File Locations
On Windows	%HOMEDRIVE% %TEMP% %USERPROFILE%
On UNIX	/ /tmp \$HOME

TDV provides the following logs for information about installation activities.

File Name	Description
cisIA.log	Internal installer log file, created at the root level of the local disk on the machine hosting the server. Contains information about the TDV software installation process. Information is logged by the installer.
cisInstall.log	Main TDV installation log file, created at the root level of the local disk on the machine hosting the server. Contains information about the actual process of installation. Information is logged by the TDV Server.
cis<version>_InstallLog.log	Internal installer log file, created at the root level of the installation directory. Contains information about the software components installed, such as registry entry, location of the file, and status of the installation attempt.
cisIA_Uninstall.log	Internal uninstallation log file, created at the root level of the local disk on the machine hosting the server. Contains process information logged by the TDV Server.
cisUninstall.log	Main TDV uninstallation log file, created at the root level of the local disk on the machine hosting the server. Information is logged by the TDV Server.

Server, Monitor, and Studio Log Files

TDV provides log and output files for TDV Server, Business Directory Server, and Studio.

- The Studio main log file is located in <Studio_install_dir>/logs (“Studio” in the table).
- Most of the other files are generated for both TDV Server and Business Directory Server.
- TDV Server log and output files are stored in <TDV_install_dir>/logs (“TDV” in the table).
- Business Directory Server log and output files are stored in <BD_install_dir>/logs (“BD” in the table).
- In a clustered environment, logs are not shared between instances, so be sure to retrieve log files from the appropriate server. For this purpose, always connect through physical addresses, not virtual addresses.

Note: For information about the installer log files, refer to [Installation and Uninstallation Logs](#).

File Name	Log Directory	Description
cs_csmonitor_daemon.log cs_bd_csmonitor_daemon.log	TDV, BD	Tracks the TDV or Business Directory Monitor if it is running as a daemon.
cs_bundles.log cs_bd_bundles.log	TDV, BD	Tracks the activity of TDV and Business Directory when they are installed as a bundle.
cs_cluster.log cs_bd_cluster.log	TDV, BD	Records all Active Cluster log messages. This file resides in the cluster directory under the logs directory. For usage in a TDV or BD cluster environment, refer to the <i>TDV Active Cluster Guide</i> . The Cluster Logging Detail Level and Cluster Event configuration parameters determine what to include.

File Name	Log Directory	Description
cs_data_cache- <day>.log	TDV	Data cache logs, each with a 3-letter day of the week (Mon, Tue, Wed, and so on) in its name. These reside in the cs_data_cache directory under the logs directory.
cs_monitor_ events.log cs_bd_monitor_ events.log	TDV, BD	Monitor Daemon events log. Records the categories of events selected through configuration parameters. See Configuring and Enabling Event Logging .
cs_monitor.log cs_bd_monitor.log	TDV, BD	Monitor Daemon main log. If the TDV Server does not start or stops responding, this log and cs_server.log are the files to check for errors.
cs_monitor.out cs_bd_monitor.out	TDV, BD	Combines stdout and stderr for the Monitor Daemon (MonitorBoot) and Server Boot process (ServerBoot). Any thread dumps of the Monitor Daemon and the Server Boot processes are written to this file. Note: Additional log files are created with a numeric suffix that increments by 1 (for example cs_monitor.out1, cs_monitor.out2, etc), once a defined limit is reached. This limit is specified in the log4j2.properties file in the directories <TDV_Install_Dir>/conf/monitor/ and <TDV_Install_Dir>/bd/conf/monitor/
cs_repository- <day>.log	TDV, BD	Repository logs, each with a 3-letter day of the week (Mon, Tue, Wed, and so on) in its name. These repo logs reside in the logs directory. These files record repository database events and status for TDV and BD. Note: If you use “composite.sh monitor stop” from the command-line, ServerBoot posts the message, “LOG: could not receive data from client: No

File Name	Log Directory	Description
		connection could be made because the target machine actively refused it." Because ServerBoot is a child process, it cannot be prevented from posting this message. However, if this message is not logged twice in a row, No error has actually occurred.
cs_server_client.log cs_bd_server_client.log	TDV, BD	Log of activities of TDV Server or BD Server clients. Use this in combination with cs_studio.log to resolve connection issues between Studio and TDV.
cs_server_dsrc.log cs_bd_server_dsrc.log	TDV, BD	Log of data source functionality.
cs_server_events.log cs_bd_server_events.log	TDV, BD	Server events log. Records the categories of events selected through configuration parameters. See Configuring and Enabling Event Logging .
cs_monitor_lifecycle.out cs_bd_monitor_lifecycle.out	TDV, BD	cs_monitor_lifecycle.out contains information about when, who (user/group), and which command is used with composite.bat. The file also contains lifecycle information about "TDV Server", "TDV Repository" and "TDV Database Cache" services. cs_bd_monitor_lifecycle.out contains information about when, who (user/group), and which command is used with bd.bat. The file also contains lifecycle information about "BD Server", and "BD Repository" services.
cs_server_file_cache.log cs_bd_server_file_cache.log	TDV, BD	Tracks TDV and Business Directory Monitor file-cache activities.

File Name	Log Directory	Description
cs_server_metadata.log cs_bd_server_metadata.log	TDV, BD	Records what objects are being written to the repository, or changes to it. For usage in a TDV or Business Directory cluster environment, refer to the <i>TDV Active Cluster Guide</i> . The Cluster Logging Detail Level and Cluster Event configuration parameters controls the categories of logging to include.
cs_server_status.log cs_bd_server_status.log	TDV, BD	Server Status log files can be used to determine software license conformance and help with corporate asset management. This log keeps data from each server session that is initiated.
cs_server_task.log cs_bd_server_task.log	TDV, BD	Supplements cs_server.log with exceptions that occur outside of the main execution thread (for example, in background threads).
cs_server.log cs_bd_server.log	TDV, BD	Main log. Nearly every error that occurs is logged here. A notable exception is unexpected Server crashes. If the TDV Server does not start or stops responding, this log and cs_monitor.log are the files to check for errors. This file also includes data source type and version information.
cs_server.out cs_bd_server.out	TDV, BD	Standard output and error log for TDV (ServerBoot) and BD Server (BDServerBoot) processes. Any thread dumps of a ServerBoot process are written to this file. Note: Additional log files are created with a numeric suffix that increments by 1 (for example cs_server.out1, cs_server.out2, etc), once a defined limit is reached. This limit is specified in the log4j2.properties file in the directories <TDV_Install_Dir>/conf/server/ and <TDV_Install_Dir>/bd/conf/server/

File Name	Log Directory	Description
cs_tools.log cs_bd_tools.log	TDV, BD	Tracks the errors that occur from the command line utilities in TDV Server and BD (for eg. server_util.sh, encryption_util.sh, bd_encryption_util.sh, bd_server.sh, etc.).
cs_studio.log	Studio	The Studio main log file. Use this in combination with cs_server_client.log to resolve connection issues between Studio and TDV.

Configuring Email Alerts for TDV Events or Actions

You can trigger email notifications for TDV actions or events.

Tip from an expert: The following configuration parameters are left over from functionality that does not send email alerts: Email Addresses for CC, Enable Email Events, Email Addresses.

To enable email alerts

1. Open and log in to Studio.
2. From the Administration menu, choose Configuration.
3. Navigate to Server > Configuration > E-Mail.
4. Set values for the following:

Configuration Parameter	Description of Value	Example
From Address	Email address that you want to appear in the From line for alerts.	meg@queenbeesknees.net

Configuration Parameter	Description of Value	Example
SMTP Host Name	Name of the email server host.	javamail.queenbeesknees.com
Maximum number of rows included in email attachment.	If set to 0, there is no restriction on the size of the email attachment. If set to a value greater than 0, the value is used as the maximum number of rows allowed for the attachment.	0

5. Save and exit the Configuration window.
6. From the Studio resource tree, right-click and select New Trigger.
7. Name and enable it.
8. Set the type of event that you want to trigger the alert. For example, a system event such as a cache refresh or data source going down. For information on how to set the different types of triggers, see the *TDV User Guide*. Choose System Event to collect information for typical TDV events including, Metrics collection, caching actions, and request spikes.

Typical Event Areas	System Event Name
Metrics Alerts	MetricsPersistentFailure
	MetricsTruncationFailure
	MetricsBackupFailure
	MetricsRestoreFailure
	StatisticsGatheringFailure
Caching	CacheRefreshFailure
	CacheRefreshSuccess

Typical Event Areas	System Event Name
Cluster Management	ClusterServerJoined
	ClusterServerConnected
	ClusterServerDisconnected
	ClusterServerShunned
Data Source Management	DataSourceDown
	DataSourceUp
Request Management	RequestFailure
	RequestInactive
	RequestRunForTooLong
	RequestsSpike
	TransactionFailure
Resource Management	ResourceLock
	ResourceUnlock
Errors and Login Management	ErrorsSpike
	FailedLoginSpike
Server Management	ServerStart
	ServerStop
Trigger Management	TriggerStart
	TriggerEnd
	TriggerFail

9. Select the Action Type of Send E-mail.

10. Specify a Resource path. For example, /shared/examples/ds_orders/tutorial/customers.

11. Type the email addresses for which to send the email alerts. For example, meg@queenbeesknees.net.
12. Type a meaningful Message Subject.
13. Type a meaningful Message Body.
14. Save the trigger.

Configuring and Enabling Event Logging

TDV has a number of mechanisms to help you control the logging of system events. The Event Generation configuration parameters specify where events of each type are recorded. The event types are grouped as follows:

- Cache events
- Cluster events
- Data source events
- Request events
- Resource events
- Security events
- Session events
- Storage events
- System overview events
- Transaction events
- Trigger events

Each group of events has its own Enable <group_name> Events configuration parameter, which you must set to True for that group of events to be recorded. This is in addition to setting the overall Enable System Events configuration parameter to True.

For each event type, you can specify one or more places where the event is recorded or handled. Currently supported event filters include:

- DB—Event sent to database only.
- LOG—Event sent to event log file only.

- SNMP—Event sent to SNMP processor only.
- CUSTOM—Event sent to custom event handler only. See [Events that Can Be Sent to Custom Event Handlers](#).
- ALL—Event sent to database, event log, SNMP processor, and custom event handler.
- NONE—Event ignored.

Use commas to separate multiple choices. Multiple choices cannot include ALL or NONE.

This section focuses on the configuration parameters you can use to control logging. Other logging-related features and where they are documented are listed in the table below.

Logging Feature	Purpose	Where Documented
Studio logging option	Enable additional logging of Studio activities to the <TDV_install_dir>/logs/cs_studio.log file for debugging purposes.	TDV User Guide
Monitoring system events	In both Manager and Studio Manager, you can view and monitor system events.	System Management with Manager
Send diagnostic log files to Support	This tool facilitates sending log files to Customer Support.	Log File Collection for Support

Documentation about event logging is included throughout the TDV documentation set. For example:

- Cache events—*TDV User Guide*
- System events and triggers—*TDV User Guide*
- Cluster events—*TDV Active Cluster Guide*

See these sections for more information about logging-related configuration parameters:

- [Enabling Logging of System Events](#)
- [Enabling SNMP Traps in TDV](#)

- [Configuring an SNMP Trap Receiver](#)
- [Events that Can Be Sent to Custom Event Handlers](#)
- [Enabling Recording of Data Source Usage in the Events Log](#)
- [Customizing Audit Log File Behavior](#)
- [Adjusting Time Limits for Request Events](#)

Enabling Logging of System Events

Enable System Events is required to activate any logging. This TDV configuration parameter must be enabled to enable a variety of logging and trigger functionality.

The instructions in this section are required to activate any event logging.

To enable system events

1. Start the TDV Server.
2. Open and log in to Studio as an admin user.
3. From the Administration menu, choose Configuration.
4. Navigate to Server > Events and Logging.
5. Set Enable System Events to True.
6. Optionally, configure the parameters under the Event Generation node.

Events log file

Enabling the System events, allows all the system events to be logged in the cs_server_events.log file. This log file is a tab-delimited file with the following fields:

Field	Type	Description
Date	TIMESTAMP	Date and time the event was logged.
Severity	VARCHAR	The severity of the event.
Status	VARCHAR	The status of the event.

Field	Type	Description
Description	VARCHAR	The short description of the event.
Type ID	INTEGER	The ID that identifies the type of event that occurred.
User Name	VARCHAR	The user that has logged in.
Domain Name	VARCHAR	The name of the domain.
SNMP ID	BIGINT	The SNMP ID of the event.
Event ID	BIGINT	The unique ID for this event.
Parent ID	BIGINT	The ID for the parent of this event. It will be the same as the EVENT_ID if it has no parent.

Enabling SNMP Traps in TDV

Simple Network Management Protocol (SNMP) traps can be used to capture and publish notifications of significant events. You can enable TDV to send these messages to a SNMP client application by following the steps outlined in this section. A configured trap receiver program is required before enabling SNMP traps within TDV. If you do not have one already configured, you can follow the guideline to set up a trap receiver described in [Configuring an SNMP Trap Receiver](#).

The TDV system supports SNMP v3 traps. TDV Server generates traps for monitoring the events that occur in the server.

For a MIB definition of the SNMP traps supported in TDV, see the MIB file available in the product installation directory at:

```
<TDV_install_dir>\apps\server\CompositeSoftware-MIB.mib
```

The CompositeSoftware-MIB.mib file contains the definitions for each trap and trap variable the user sees when they view the trap.

About the SNMP Server Events in TDV

TDV Server creates SNMP events that are compliant with SNMPv3 protocol. The CompositeSoftware-MIB.mib file contains details of these server events.

You can modify SNMP log settings from Studio by selecting the Administration > Configuration menu option, and navigating to Server > Events and Logging > Logging > SNMP.

To enable SNMP traps in TDV

1. Review the available TDV SNMP traps by opening the following MIB file:

```
<TDV_install_dir>\apps\server\CompositeSoftware-MIB.mib
```

2. Optionally, you can add a new trap definition to CompositeSoftware-MIB.mib. For example:

```
csSecurityRBSAssign TRAP-TYPE
```

```
ENTERPRISE csTrapsV3
```

```
VARIABLES{ trapTime, trapServerHostName, trapServerPort, trapPolicyName,
trapResourceName }
```

```
DESCRIPTION "This trap is generated when a Row Based Security policy has been
assigned."
```

```
::= 22005
```

3. Optionally, for any new trap definitions, define the trap variables. For example, a variable definition in the MIB is:

```
trapResourceName OBJECT-TYPE
```

```
SYNTAX DisplayString (SIZE(1..256))
```

ACCESS read-only

STATUS mandatory

DESCRIPTION "A string that indicates the resource name for the generated trap."

::= { trapVars 140 }

trapPolicyName OBJECT-TYPE

SYNTAX DisplayString (SIZE(1..256))

ACCESS read-only

STATUS mandatory

DESCRIPTION "A string that indicates the policy name for the Row Based Security policy."

::= { trapVars 185 }

4. In Studio, select Administration > Configuration.
5. Navigate to Server > Events and Logging > Logging > SNMP.
6. Specify values for the following configuration parameters.

Configuration Parameter	Description of Value to Set
Enable SNMP Events	True
Trap Community	Community string to which to send SNMP traps. Default: public. Make sure the SNMP client uses the same value as you

Configuration Parameter	Description of Value to Set
	set here.
Trap Host List	A comma-separated list of host names or IPs that will be sent the SNMP v3 trap messages.
Trap Port	Specify the port for the trap receiver.

7. Navigate to Event Generation.

8. Add SNMP to the list for event types you want to send to SNMP.

For example, to monitor Request Start events, add a comma and SNMP to the Value for the Request Start configuration parameter under Event Generation > Request Events. This sends the SNMP trap to the trap receiver that you specified using the Trap Port configuration parameter.

Information to Monitor	Description of Notice	Configuration Steps
TDV Session	Alert if it reaches a certain threshold	Create a query that uses the <code>/services/databases/system/SYS_SESSIONS</code> published view in the SYSTEM repository.
TDV Cache	Monitor required for cache success or failure	Define the following traps under Events and Logging > Event Generation > Cache Events configuration parameters: Cache Refresh Start, Cache Refresh Fail, Cache RefreshEnd You can instead define a query that uses the STATUS column from the <code>/services/databases/system/SYS_CACHES</code> view to capture the information and write it to the log files or SNMP monitor program.
Data Sources	Determine whether the data	Define the following traps under Events and Logging > Event Generation > Data Source

Information to Monitor	Description of Notice	Configuration Steps
	source is up and running or not	Events: Data Source Up, Data Source Down You can instead define a query that uses the STATUS column in /services/databases/system/SYS_DATASOURCES to capture the information and write it to the log files or SNMP monitor program.
Trigger	Alert on failure	Define the following traps under Events and Logging > Event Generation > Trigger Events: Trigger Start, Trigger Fail, Trigger End You can instead define a query that uses the STATUS column in /services/databases/system/SYS_TRIGGERS to capture the information and write it to the log files or SNMP monitor program.
Long-running queries	Alert if it reaches a certain threshold	Define the threshold value for the duration which is considered as long in Events and Logging > Event Generation > Request Events: Request Run Time This value must be in minutes and controls the period of time after which a request is considered to be long-running, resulting in the generation of a RequestRunForTooLong event.

Information to Monitor	Description of Notice	Configuration Steps
Number of waiting requests	Alert if it reaches a certain threshold	1. Add “SNMP” to the value of the Events and Logging > Event Generation > Request Events > Request Wait Queue Threshold Passed configuration parameter. 2. Set the Server > Runtime Processing Information > Wait Queue > Wait Queue Threshold configuration parameter. This is the number of requests in the wait queue at which a wait queue threshold event is triggered. This generates a csRequestWaitQueueThresholdPass trap.

Configuring an SNMP Trap Receiver

TDV does not distribute an SNMP Trap Receiver. The following instructions are provided as a guideline. The actual steps that you need to perform will differ depending on the product that you choose to use.

To view and verify the SNMP traps

1. Download and install an SNMP Trap Receiver. for example, iReasoning MIB Browser Free Personal Edition.
2. Launch the program and load the TDV MIB.
For example in iReasoning, navigate to File > load MIB and select <TDV_install_dir>\apps\server\CompositeSoftware-MIB.mib.
3. View the traps.
For example in iReasoning, navigate to the Tools > Trap Receiver menu and open the Trap Receiver tab at the right.
4. Locate the trap filter in your trap receiver tool and specify the port number to listen to.

This will be the same port that you will specify in TDV under Logging > SNMP > Trap Port so that TDV can send the trap to this port.

For example in iReasoning, in the Trap Filter specify the port you want to listen to receive the traps. Because of permissions issues, you might want to set this to something greater than 1024, like 5000.

5. Start the trap receiver.
6. To test the trap receiver, trigger a captured event in TDV

Events that Can Be Sent to Custom Event Handlers

The following is a list of the values of type that can be output from the following Java call in the `handleEvent` method of a custom event handler:

```
String type = (String)eventInfo.get("type");
```

Each Event Type also appears in an Event Group (Cache Events, Cluster Events, and so on) in the Configuration window under Server > Events and Logging > Event Generation. (These are also listed in the “TDV Configuration Parameters” topic of the TDV Reference Guide.) For descriptions and values, find and highlight them in the Configuration window.

Event Type	Event Group
CS_CACHE_CLEAR	Cache Events
CS_CACHE_DISABLE	Cache Events
CS_CACHE_ENABLE	Cache Events
CS_CACHE_REFRESH_END	Cache Events
CS_CACHE_REFRESH_FAIL	Cache Events
CS_CACHE_REFRESH_START	Cache Events
CS_CLUSTER_SERVER_CONNECTED	Cluster Events

Event Type	Event Group
CS_CLUSTER_SERVER_DISCONNECTED	Cluster Events
CS_CLUSTER_SERVER_JOINED	Cluster Events
CS_CLUSTER_SERVER_SHUNNED	Cluster Events
CS_CONN_CHECKED_IN	Data Source Events
CS_CONN_CHECKED_OUT	Data Source Events
CS_CONN_FAIL	Data Source Events
CS_CONN_INVALID	Data Source Events
CS_CONN_POOL_EXHAUSTED	Data Source Events
CS_CONN_POOL_SIZE_DECREASE	Data Source Events
CS_CONN_POOL_SIZE_INCREASE	Data Source Events
CS_DATA_SOURCE_DOWN	Data Source Events
CS_DATA_SOURCE_INTROSPECT_CANCEL	Data Source Events
CS_DATA_SOURCE_INTROSPECT_END	Data Source Events
CS_DATA_SOURCE_INTROSPECT_FAIL	Data Source Events
CS_DATA_SOURCE_INTROSPECT_START	Data Source Events
CS_DATA_SOURCE_MODIFY	Data Source Events
CS_DATA_SOURCE_OFF	Data Source Events
CS_DATA_SOURCE_ON	Data Source Events

Event Type	Event Group
CS_DATA_SOURCE_STATS_COMPLETE	Data Source Events
CS_DATA_SOURCE_STATS_FAIL	Data Source Events
CS_DATA_SOURCE_STATS_START	Data Source Events
CS_DATA_SOURCE_TEST_FAIL	Data Source Events
CS_DATA_SOURCE_TEST_START	Data Source Events
CS_DATA_SOURCE_TEST_SUCCESS	Data Source Events
CS_DATA_SOURCE_UP	Data Source Events
CS_DOMAIN_CREATE	System Overview Events
CS_DOMAIN_DELETE	System Overview Events
CS_GROUP_CREATE	System Overview Events
CS_GROUP_DELETE	System Overview Events
CS_MONITOR_FAIL	System Overview Events
CS_MONITOR_START	System Overview Events
CS_MONITOR_STOP	System Overview Events
CS_PREPARED_STATEMENT_FAIL	Request Events

Event Type	Event Group
CS_PREPARED_STATEMENT_SUCCESS	Request Events
CS_REPOSITORY_DOWN	System Overview Events
CS_REPOSITORY_UP	System Overview Events
CS_REQUEST_CANCEL	Request Events
CS_REQUEST_END	Request Events
CS_REQUEST_FAIL	Request Events
CS_REQUEST_INACTIVE	Request Events
CS_REQUEST_RUN_FOR_TOO_LONG	Request Events (Request Run Time)
CS_REQUEST_START	Request Events
CS_REQUEST_WAIT	Request Events
CS_REQUEST_WAIT_QUEUE_THRESHOLD_PASS	Request Events
CS_REQUEST_WAIT_QUEUE_THRESHOLD_RESET	Request Events
CS_RESOURCE_CREATE	Resource Events
CS_RESOURCE_DELETE	Resource Events
CS_RESOURCE_LOCK	Resource Events
CS_RESOURCE_STATS_COMPLETE	Resource Events

Event Type	Event Group
CS_RESOURCE_STATS_FAIL	Resource Events
CS_RESOURCE_STATS_START	Resource Events
CS_RESOURCE_UNLOCK	Resource Events
CS_SECURITY_RBS_ASSIGN	Security Events
CS_SECURITY_RBS_CREATE	Security Events
CS_SECURITY_RBS_DELETE	Security Events
CS_SECURITY_RBS_DISABLE	Security Events
CS_SECURITY_RBS_ENABLE	Security Events
CS_SECURITY_RBS_REMOVE	Security Events
CS_SECURITY_RBS_UPDATE	Security Events
CS_SERVER_RESTART	System Overview Events
CS_SERVER_RESTART_FAIL	System Overview Events
CS_SERVER_START	System Overview Events
CS_SERVER_STOP	System Overview Events
CS_SERVER_STOP_PLANNED	System Overview Events
CS_SERVER_STOP_UNPLANNED	System Overview Events

Event Type	Event Group
CS_SESSION_END	Session Events
CS_SESSION_LOGIN_FAIL	Session Events
CS_SESSION_MAX_CONNECTIONS_EXHAUST	Session Events
CS_SESSION_NON_LOCALHOST_REQUEST_FAIL	Session Events
CS_SESSION_RUN_FOR_TOO_LONG	Session Events (Session Open Time)
CS_SESSION_START	Session Events
CS_SESSION_TERMINATE	Session Events
CS_STORAGE_LOW_CRITICAL	Storage Events
CS_STORAGE_LOW_WARNING	Storage Events
CS_TRANSACTION_COMMIT	Transaction Events
CS_TRANSACTION_COMPENSATE	Transaction Events
CS_TRANSACTION_FAIL	Transaction Events
CS_TRANSACTION_ROLLBACK	Transaction Events
CS_TRANSACTION_START	Transaction Events
CS_TRIGGER_END	Trigger Events
CS_TRIGGER_FAIL	Trigger Events
CS_TRIGGER_START	Trigger Events

Event Type	Event Group
CS_USER_ADD_TO_GROUP	System Overview Events
CS_USER_CREATE	System Overview Events
CS_USER_DELETE	System Overview Events
CS_USER_PASSWORD_MODIFY	System Overview Events
CS_USER_REMOVE_FROM_GROUP	System Overview Events

Enabling Recording of Data Source Usage in the Events Log

Several configuration parameters can be set to record detailed data source usage information which can then be found in the events log (<TDV_install_dir>\logs\cs_server_events.log). Each of the configuration parameters below specify that certain values be added to the events log and are used in combination to control what is collected.

Configuration Parameter that You Set to True	Values Added to Events Log
Enable System Events	Request ID, Transaction ID, Session ID, Session Host, Session Client Type, User Name, Domain Name, Internal (True for system-generated events), Bytes In, Bytes Out, Rows Affected (number of rows processed)
Include Data Source Timings	All of the above, plus Time to First Row, which is the time, in milliseconds, from the moment TDV received a request to the moment that TDV has fetched the first row from a

Configuration Parameter that You Set to True	Values Added to Events Log
	data source.
Detailed Profiling Enabled	All of the above, plus Data Source Time which is the time, in milliseconds, spent in the data source—not including any time spent in TDV.

The values listed in the table are added to the events log only if the listed combinations of parameters are set to True.

To enable recording of data source usage in the events log

1. Start the TDV Server.
2. Open and log in to Studio as an admin user.
3. From the Administration menu, choose Configuration.
4. Navigate to Server > Events and Logging.
5. Make sure that Enable System Events is set to True.
6. Optionally locate Include Data Source Timings, and set it to True.
7. Optionally locate Detailed Profiling Enabled, and set it to True.

Setting Detailed Profiling Enabled to True can have significant negative impact on performance.

Here is an example of the data source usage information (the last 13 fields of the message) recorded in `cs_server_events.log` with the Enable System Events, Include Data Source Timings, and Detailed Profiling Enabled configuration parameters all set to True:

```
...16338 40803072 318257 x.y.com JDBC test yz false 578 315 2 73 2
```

These fields are interpreted as follows:

Req ID	Transaction ID	Session ID	Host	Client Type	User	Domain	Internal	Bytes In	Bytes Out	Rows Affected	Time to First Row	Data Source Time
	16338	40803072	318257	x.y.z.com			t		e		5	
												2

Log files are discussed in [About TDV Log Files](#).

Customizing Audit Log File Behavior

You can customize the TDV event handler to create an audit log file that captures TDV requests. Depending on your audit file needs, you might want to write a custom event handler and configure your TDV events with different values.

The TDV event handler API is not limited to audit file log events. You can create a custom event handler that captures the specific information that you need.

To customize audit log behavior using a custom event handler

1. Create a customEventHandler.jar event handler that extends `com.compositesw.extension.events.EventHandler`. The JAR file must be named customEventHandler.jar.

The following sample code captures events and saves them to a log file under /tmp on UNIX:

```
import com.compositesw.extension.events.EventHandler;
```

```
import java.io.File;
```

```
import java.io.FileWriter;
```

```
import java.io.PrintWriter;
```

```
import java.util.Map;
```

```
import java.util.Set;
```

```
import java.util.Date;
```

```
import java.util.Locale;
```

```
import java.text.DateFormat;
```

```
import java.text.SimpleDateFormat;
```

```
public class SampleEventHandler
```

```
    implements EventHandler
```

```
{
```

```
    private DateFormat df = new SimpleDateFormat("yyyy-MM-dd'T'HH:mm:ss.SSSZ",Locale.US);
```

```
    public void handleEvent(Map eventInfo)
```

```
    throws Exception
```

```
{  
  
    File f = new File("/tmp/events.log");  
  
    PrintWriter log = new PrintWriter(new FileWriter(f, true));  
  
    log.println("Event:");  
  
    long time = Long.parseLong((String)eventInfo.get("time"));  
  
    String dateString = df.format(new Date(time));  
  
    log.println("\tdate-->" + df.format(new Date(time)));  
  
    Set keys = eventInfo.keySet();  
  
    for (Object key : keys) {  
  
        log.println("\t" + key.toString() + "-->" + eventInfo.get(key));  
  
    }  
  
    log.flush();  
  
}  
  
}
```

2. Stop the TDV Server.
3. Copy customEventHandler.jar and save it in the <TDV_install_dir>/apps/extension/lib directory.

4. Start the TDV Server.
5. Open and log in to Studio as an admin user.
6. From the Administration menu, choose Configuration.
7. Navigate to Server > Events and Logging.
8. Make sure that Enable System Events is set to True.
9. Navigate to Server > Events and Logging > Logging > Custom Logger.
10. For Custom Jar Location, type the fully qualified directory of the customEventHandler.jar that you created in Step 1.
11. For Enable Custom Logging, select True.
12. Enable custom audit requests under Server > Events and Logging > Event Generation > Request Events. There are many types of events that you can schedule.

For this example set the value of the following parameters as indicated in the following table:

Configuration Parameter	Value for the Example
Request Start	DB, LOG, CUSTOM
Request End	DB, LOG, CUSTOM
Request Fail	DB, LOG, CUSTOM

13. Optionally, to capture the web services name in the log file, the Enable Events for Internal Request property must be set to true.
14. Restart the TDV Server.
15. To see the TDV request events in the example log file:
 - Log into Studio.
 - In the resource tree, go to Desktop > Data Services > Databases > system > LOG_EVENTS.
 - Open the LOG_EVENTS table.
 - Execute the table query so that events are sent to the log file.

Outside of Studio, open the log file and see the TDV event details for the Request Start and Request End events.

Adjusting Time Limits for Request Events

You can adjust the time limits set for certain request events. For example, if you need to know that a request has been inactive for 20 minutes or if the request has been running for more than 15 minutes, there are TDV configuration parameters that you can use.

To adjust time limits for request events

1. Start the TDV Server.
2. Open and log in to Studio as an admin user.
3. From the Administration menu, choose Configuration.
4. Navigate to Server > Events and Logging.
5. Make sure that Enable System Events is set to True.
6. Navigate to Server > Events and Logging > Event Generation > Request Events.
7. Adjust the values of the following parameters.

Configuration Parameter	Description
Request Run Time	Controls the number of minutes after which a request is considered to be running too long, and so a RequestRunForTooLong event is generated.
Request Inactive Time	Controls the number of minutes after which a request is to be checked for inactivity. If found inactive, it is considered stale, and a corresponding event is generated.

8. Restart the TDV Server.

Determining Data Source Type and Version Information

It can be helpful to know all the data source types and their versions that are being used as part of your TDV installation. You can locate this information in the `cs_server_status.log` file. For example:

Datasource Info

/shared/myCaches/db-lab-9
size=1, in=1, out=0, total created=1, total destroyed=0, init=0, min=10, max=100, idle=30
JDBC Datasource : Oracle Oracle Database 11g Enterprise Edition Release 11.2.0.2.0 - 64bit Production With the Partitioning, OLAP, Data Mining and Real Application Testing options
JDBC Driver : Oracle JDBC driver 11.2.0.2.0
/shared/examples/ds_orders
size=1, in=1, out=0, total created=1, total destroyed=0, init=0, min=10, max=100, idle=30
Revision: \${svn.Revision})

To determine the data source types and versions

1. Locate the installed TDV servers within your corporate environment.

2. Using your preferred file management tool, navigate to:

```
<TDV_install_dir>/logs
```

3. Open the cs_server_status.log file.

The log file keeps data from each server session that is initiated. Information for each new session is added to the end of the file.

4. Locate the portion of the file that has been added most recently.
5. Review the data in the Datasource Info section of the text file. For example:

```
JDBC Datasource : Oracle Oracle Database 11g Enterprise Edition Release 11.2.0.2.0 -  
64bit Production With the Partitioning, OLAP, Data Mining and Real Application Testing  
options
```

```
JDBC Driver : Oracle JDBC driver 11.2.0.2.0
```

```
/shared/examples/ds_orders
```

```
size=1, in=1, out=0, total created=1, total destroyed=0, init=0, min=10, max=100,  
idle=30
```

Logging Query Execution Statistics

Determining the cause of slow running queries can help you tune the performance of your TDV environment.

To log query execution statistics

1. Enable the logging of purged request statistics from the Studio Configuration window. For example, Administration > Configuration and navigate to SQL Engine > Logging > Query Statistics Logging

Configuration Parameter	Description of Value
Maximum Number of Logged Purged Requests	Use to refine the number of query execution records retained in the log files.
Log Purged Request Query Statistics	Use to turn the logging of query execution statistics on or off.

2. Open the Studio Manager.
3. Navigate to the Requests Panel.
4. At the bottom of the Requests panel, select the Include Logged Requests check box.
5. Review the generated log files.
6. Determine what if anything can be adjusted to improve the performance of the queries that appear to be running slow.

Log File Collection for Support

TDV provides a mechanism, sometimes referred to as the Collector Tool, that helps TDV users collect useful information from both TDV and the system that it is running on to help Support and the user diagnose support cases. You can select various system information and statistics files to collect. You can download the diagnostics zip file to your local machine and upload it directly to Support.

- [About the System Information Files](#)
- [Security with Log File Collection](#)

How Log File Collection Works

When you request to send or collect log files, Studio updates all existing TDV logs and copies all log and configuration files into a temporary folder. Studio runs a set of platform-specific commands to get an overview of the system that is running the TDV Server and saves that information to the tmp folder. See [Saving Log Files to Support](#) .

When collection is complete, the folder is compressed into a zip file, and you have the option to upload it straight to Support or save it locally.

The files that are collected include:

File Location	Description
<TDV_install_dir>/logs/*	The logs directory in the zip file.
<TDV_install_dir>/conf/*	The conf directory in the zip file.
apps/dlm/*/conf/	The apps directory in the zip file.
Results of system commands	Put into the tmp directory in the zip file.

About the System Information Files

To collect the system information, Studio runs the following system commands:

For Windows:

- Disk info: fsutil volume diskfree C:
- IP configuration: ipconfig /all
- Network statistics: netstat -na
- Task list: tasklist /fo table
- System information: systeminfo

For UNIX:

- Kernel information: dmesg
- Disk information: df -h
- IP configuration: ifconfig -a
- Network statistics: netstat -na
- Task list: ps -ef
- System information: uname -a

These files are generated by the system commands.

File Name	Contents
diskinfo.info	The total number of bytes, free bytes, and available free bytes.
ipconfig.info	Windows IP configuration information. Ethernet adapter information.
netstat.info	Active connections and their state.
systeminfo.info	Specifics about the computer running TDV like operating system information, memory information, processors installed, hotfixes applied, network cards installed,
tasklist.info	List of tasks running and memory used at the time of diagnostics collection.

Security with Log File Collection

You must have these access privileges to use log file collection:

- ACCESS_TOOLS
- READ_ALL_CONFIG
- READ_ALL_STATUS

If any resource definitions are returned, you also need:

- READ_ALL_RESOURCES

See [Managing Security for TDV Resources](#).

Saving Log Files to Support

You can collect the files in one of three ways:

- Studio
- Manager
- From the command line

The diagnostics collected are exactly the same regardless of which method you use. The processes for each method are described in the following sections.

- [Generating Log Files in Studio or Manager](#)
- [Generating Log Files Using the Command Line](#)

Generating Log Files in Studio or Manager

Using Studio, you can select which log files to generate. This also creates a zip file from the log files.

To generate log files in Studio or Manager

1. Run Studio, from the Administration menu, choose Save Logs for Support.
2. Click OK.
3. Click the check boxes next to the files you want to collect.
4. Click Save.
5. Name and save the zip files to a location of your choice.
6. Follow the prompts on the screen and collect your zip files as necessary.

Generating Log Files Using the Command Line

You can collect the log files from the command line using `server_util`. See [TDV Command-Line Utilities](#) for more information about `server_util`.

To generate log files from the command line

7. Open a command prompt window.
8. Navigate to the `<TDV_install_dir>/bin`.
9. Enter the `server_util` command with the `-saveLogs` option that specifies which log files to collect for a Support case number.

The command syntax for `server_util` using the `-saveLogs` option is as follows:

```
server_util -server <hostname> [-port <port>] [-encrypt]
```

```
-user <username> -password <password> [-domain <domain>]
```



```
-saveLogs [-port <port>] [-folder <filepath>] [-exclude <File Group>]
```

where <File Group> can be any combination of "logs, conf, sysinfo".

For example, a basic command to upload all files would be:

```
./server_util.sh -server localhost -user admin -password admin -saveLogs
```

Using TDV Log Files to Track Resource Privilege Changes

Occasionally it can be helpful to be able to determine when, if, and how privileges have been changed for a given Studio resource. This information is tracked in the `cs_server_metadata.log` file. By default, this feature is not enabled because it can cause the log file to grow very large, very fast. You must enable privilege logging using Studio configuration parameters before this information will be captured in the log files.

To track resource privilege changes

1. Log into Studio as the admin user.
2. From the Administration menu, choose Configuration.
3. Navigate to Server > Configuration > Security.
4. Set the value of Enable Privilege Logging to True.
5. Click Apply.
6. Click OK.
7. Restart the TDV Server to implement your changes.
8. Wait for TDV Studio resources to undergo privilege changes.
9. Using your preferred file management tool, navigate to:

```
<TDV_install_dir>/logs
```

10. Open the `cs_server_metadata.log` file.
11. Locate and review the data in PRIVILEGE sections of the text file.

For example:

```
UPDATED DATA_SOURCE /shared/security/ReqSignEncRepSignEnc (17663)
```

```
PRIVILEGE composite/all (2) READ WRITE EXECUTE SELECT UPDATE INSERT  
DELETE GRANT
```

```
PRIVILEGE dynamic/all (3) READ WRITE EXECUTE SELECT UPDATE INSERT  
DELETE GRANT
```

```
2013-06-26 11:51:46.006 composite/admin (-1973) saved following changes:
```

```
802 804
```

```
UPDATED DATA_SOURCE /shared/examples/ds_orders (10390)
```

```
PRIVILEGE composite/all (2) READ EXECUTE SELECT
```

Validating TDV Software License Compliance and Asset Management

Diligently following application licensing compliance can save your organization from legal or standards infringement problems. You can gather information from the TDV log files to determine your compliance for auditing and renewal purposes.

This section contains:

- [Determining Your TDV Software License Conformance](#)
- [Tips for Configuring the Number of TDV Processors](#)

Determining Your TDV Software License Conformance

To determine your TDV software license conformance

1. Locate your TDV enterprise license agreement (ELA).
2. From the ELA, determine the values for the following:
 - Version of TDV
 - Number of Cores Licensed OR Number of Processors Licensed
3. Locate all the installed TDV servers within your corporate environment.
4. Using your preferred file management tool, navigate to:

```
<TDV_install_dir>/logs
```

5. Open the cs_server_status.log file.

The log file keeps data from each server session that is initiated. Information for each new session is added to the end of the file.

6. Locate the portion of the file that has been added most recently.
7. Review the data in the Server Stats and License Info sections of the text file.

For example:

```
-----
```

```
| Server Stats |
```

```
-----
```

```
Server Name:      7smith-l8:9408
```

```
Operating System:  Windows 8
```

```
Number of processors:  8
```

Total Memory Used:	6% (71MB of 1058MB)
Total Sessions:	36
...	

License Info	

...	
Product cserver:	
Version = 7.0	
Creation Date = 20150822175953099	
Activation Date = 20150905	
Duration = 728	
Expiration Date = 20180904	
Type = 0	
Owner = development	

8. Open the cs_server.log file.

The log file keeps data from each server session that is initiated. Information for each new session is added to the end of the file.

9. Locate the line with the Number of Processors text. For example:

```
INFO [main] 2013-03-14 18:00:52.656 -0700 LicenseManager - Number of Processors in
the system : 8
```

10. Determine if the Number of processors (cores) and product Version from the log files are consistent with your TDV enterprise license agreement.
11. If you are not in compliance, determine the next step that is appropriate to take to remedy the situation.

Tips for Configuring the Number of TDV Processors

Configuring the number of TDV processors can help you take control of compliance to your TDV license terms. Because the environments at different companies varies so widely, you will need to research and perform testing to determine the best method for your particular environment.

Tips for configuring the number of TDV processors

1. Review documents and instructions for how to set CPU affinity.
For example, navigate to and review:
 - <http://www.cyberciti.biz/tips/setting-processor-affinity-certain-task-or-process.html>
 - <http://pundiramit.blogspot.com/2010/07/how-to-disable-cpu-cores-in-multicore.html>
 - <http://stackoverflow.com/questions/628057/how-to-set-processor-affinity-on-an-executable-in-windows-xp>
 - http://www.experts-exchange.com/OS/UNIX/AIX/Q_27263123.html
 - <http://linux.die.net/man/1/taskset>
2. Determine your number of available CPUs and their unique identifications.
3. Determine the names of the TDV processes that need to be associated with the specific CPUs.

4. Determine if one of the following commands can help you configure your number of TDV processors. Some key commands to help you configure processors depending on operating system are:

Platform	Command	Description
AIX	<code>bindprocessor 1234 1</code>	Bind the kernel threads to the process of a processor.
Windows	<code>start java startServer /affinity:1,2,3,4</code>	Modifies the startup script to provide affinity to 4 CPUs.
Windows	<code>imagecfg -a 0x3 <xxx>.exe</code>	Limits the executable to CPU0 and CPU1
UNIX / LINUX	<code>taskset [options] mask command [arg]...</code> <code>taskset [options] -p [mask] pid</code>	taskset can be used to set the affinity of a running process or to launch a process with a certain affinity.

5. Test your configuration changes and determine if further changes are needed.

Logging Tips from an Expert

- [Using Detailed Logging](#)
- [Controlling the Size of Files Specified in log4j.properties](#)

Using Detailed Logging

It is possible to turn detailed logging on and off without performing a TDV Server restart.

To turn on detailed logging

1. Open the Configuration window from Studio.
2. Locate the Debug Output Enabled parameter and set its value to true.
3. After the process for which you want detailed logging information runs, turn the Debug Output Enabled parameter back to false.

4. Review the log files to determine a potential course of action.

Controlling the Size of Files Specified in log4j.properties

It is possible to configure the size of the files (cs_server.out, *.log) specified in your log4j.properties file using to parameters. MaxFileSize controls the maximum file size. After the value specified is reached the data rolls over to a new file. MaxBackupIndex controls the maximum number of rollover files.

To manipulate the settings in your log4j log file

1. Navigate to `<TDV_install_dir>/conf/server/log4j.properties`.
2. Locate the file for which you want to control the size.
3. Edit the following lines:

```
log4j.appender.MONITOR_STDOUT.MaxFileSize=10000KB
```

```
log4j.appender.MONITOR_STDOUT.MaxBackupIndex=100
```

4. Save the file.
5. Restart the TDV Server.

Hiding Enabling Error Details

Errors encountered in TDV include a Log Id which can be interpreted with the corresponding information from Stack trace. While this is helpful for debugging, it is advisable to turn off the stack trace output for security.

You can enable/disable the stack trace by setting the configuration parameter "Enable Stack Traces for Debugging" to True or False.

Setting this value to "true" will show stack traces for troubleshooting purposes. Otherwise, the stack traces would not be printed.

Configuring TDV Data Connections

This topic describes how to install and configure connection interface adapters. It also covers configuration changes to make when connecting to the AIX platform.

The following topics are included:

- [Installing and Using Preconfigured JDBC Drivers](#)
- [Using the ODBC Driver on Windows](#)
- [Using the ODBC Driver on UNIX](#)
- [Using the ODBC Driver on UNIX](#)
- [Configuring TDV for Using a JMS Broker](#)
- [Configuring TDV for AIX Platforms](#)

Make sure that you have completed the steps in the “Installing the TDV Client Drivers that are Distributed with TDV” topic that is in the *TDV Installation and Upgrade Guide*.

For information on how to configure custom connection adapters, see “Working with Data Sources” in the *TDV User Guide*.

Installing and Using Preconfigured JDBC Drivers

JDBC drivers provide API calls for Java programs to communicate with databases. A single JDBC driver can be used to connect to any number of the same type of data sources. After uploading, the JDBC driver can be used with other JDBC data sources, such as Oracle, SQL Server, or PostgreSQL.

You need to install the necessary drivers in the locations that are documented so that the TDV Server and the data source can interact. If you plan on using more than one of the drivers described, we recommend obtaining and placing them all in the necessary directory locations before restarting the TDV Server.

TDV ships with a JDBC interface and provides adapters to connect to relational data sources. You can customize these adapters to connect to new or custom data sources.

To override the default JDBC location (for example, `TDV_INSTALL_DIR/apps/dlm/cis_ds_postgresql/lib` for the Postgresql data source), place the JDBC driver jar file in the location `TDV_INSTALL_DIR/conf/adapters/system` and restart the TDV Server. Note that the server does not make any accommodations for JDBC drivers that do not supply correct metadata about the data source.

Refer to the corresponding Adapter guides for a description on which preconfigured JDBC drivers are required for use of specific data sources and where those drivers should be placed so that they can be used when connecting to specific data sources. These data source drivers must be installed separately from the TDV Software installation because of third-party licensing restrictions. After the required data source driver is installed, you can create a JDBC data source using Studio as described in the *TDV User Guide*.

Using the ODBC Driver on Windows

ODBC is an API for programs that use SQL statements to access data. The ODBC drivers can provide access to more than relational databases. ODBC defines the client side of database connectivity but not the server side. ODBC drivers typically rely on the presence of a proprietary driver. ODBC drivers transform ODBC calls into access requests and responses. You must install and configure an ODBC driver on each client and install the vendor-specific proprietary driver on the server side.

This section covers the following topics:

- [Supported ODBC Data Types](#)
- [Adding ODBC Data Sources on Windows](#)

Supported ODBC Data Types

The TDV ODBC driver supports and maps the following data types.

Supported Data Type	Data Type is Mapped To
CHAR, VARCHAR	VARCHAR
BIT, TINYINT, SMALLINT	SMALLINT
BIGINT, INT	INT
DECIMAL, REAL, FLOAT, NUMERIC	FLOAT
SHORT, LONG, DOUBLE, TIME, DATE,	VARCHAR

Supported Data Type

TIMESTAMP

Data Type is
Mapped To

Data Types of Unknown Length

The TDV ODBC driver supports retrieval of parameters with values longer than 255 characters if the client provides adequate memory for the task.


If the parameter is a wildcard of unknown length, the parameter is defined with a data type of VARCHAR (255). SQL parameters of a given length submitted in a prepared statement are assigned a data type of CHAR with the actual character length of the parameter submitted.

Adding ODBC Data Sources on Windows

TDV supports the native Windows driver managers. If you need to install the ODBC drivers, see the *TDV Installation and Upgrade Guide*.

To add an ODBC data source to a Windows machine

1. Select Windows Control Panel > Administrative Tools > Data Sources (ODBC) to open the ODBC Data Source Administrator. Or, when using the 32-bit driver on 64-bit machines, navigate to ...\\Windows\\SysWOW64\\ and run odbcad32.exe.



2. Select 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 Add.
4. In the Create New Data Source screen, select the TDV <version> driver, and click Finish.
5. In the Driver Configuration window, enter the following information that is required for configuring the driver:

Field	Description
DSN Name	Name of the data source that the clients refer to. After a DSN is created, its name cannot be changed.
Host	Server name (or IP address) on which TDV is running.
Port	TCP port used to communicate with TDV Server, which must match the port that the server is listening on. With default installation settings TDV listens on port 9401, but that setting should be verified by checking the port setting in the Configuration window accessible through Studio Administration menu: TDV Server > Client Drivers > Communications > Port
Integrated Authentication	Method for authenticating the ODBC connection: disabled (default), Kerberos, or NTLM.
Kerberos SPN	SPN for Kerberos to use to authenticate the ODBC connection. Ungrayed if Kerberos is selected as Integrated Authentication.
User Name, Password, and Domain	Must be valid within TDV Server. The password is nullable. NOTE: The ODBC manager may truncate the password at 14 characters.
Datasource	Name of the TDV data source that the ODBC connection accesses. This entry sets the default scope of client queries to a particular datasource. Querying outside the scope of this data source requires super-qualified tables or stored procedures.

Field	Description
Catalog	Connects with a default data source catalog.

6. Click Refresh to retrieve the catalogs available to this user on the server.
7. Click Test 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 machine.

Using the ODBC Driver on UNIX

TDV ODBC drivers are available for 32-bit or 64-bit UNIX operating systems. You must install the correct version for your environment.

Creating a DSN is done through the configuration utility. The interactive utility is driverConfig. Use driverConfig to reconfigure the driver files (when the driver file location has changed), and create, edit, list, or delete DSN entries.

The following describes the tasks for using the ODBC driver on UNIX:

- [Setting the ODBC Environment Variables on UNIX](#)
- [Creating a DSN with driverConfig on UNIX](#)
- [Configuring TDV for Using a JMS Broker](#)

Setting the ODBC Environment Variables on UNIX

For examples and instructions for how to set UNIX environment variables, refer to your favorite UNIX guidelines. A typical command might be:

```
setenv PATH "/bin:/usr/bin:/usr/sbin:ucb/bin"
```

To set the ODBC environment variables

1. Log into the installation machine as the same user that installed the TDV software.

2. Set the following environment variables:

Variable	Description and Values
COMPOSITE_HOME	This optional variable allows you to specify the location where the TDV ODBC driver is installed. This is the full path to the <TDV_ODBC_install_dir> for the TDV ODBC driver. If this configuration is not set, you can run driverConfig with an absolute or relative path, for example: <pre>./home/release/apps/odbc/linux64/bin/driverConfig</pre> <pre>./odbc/linux64/bin/driverConfig</pre> <pre>./bin/driverConfig</pre> If this variable is set to /home/release, then when you create a DSN, it goes to /home/release/apps/odbc/linux64/lib to find the so files.
ODBCINI	Full path to the configuration file odbc.ini. It is generated during creation of DSN configuration with driverConfig. The ODBCINI and ODBCINSTINI files do not exist yet and will be created during DSN creation in the next step. It should be: <TDV_install_dir>/odbc.ini
ODBCINSTINI	Full path to the ODBC drivers configuration file odbcinst.ini. It is generated during DSN configuration with driverConfig. The ODBCINI and ODBCINSTINI files do not exist yet and will be created during DSN creation in the next step. It should be: <TDV_install_dir>/odbcinst.ini
LD_LIBRARY_PATH	This is specific to Linux. This path refers to the location of the iODBC driver manager files. The default location is: <pre><TDV_install_dir>\apps\odbc\<platformType>\lib</pre> If you already have this variable, add the additional path to the existing path definition.
LIBPATH	This is specific to AIX. This path refers to the location of the iODBC

Variable	Description and Values
	driver manager files. The default location is: <TDV_install_dir>\apps\odbc\<platformType>\lib
SHLIB_PATH	This is specific to HP-UX. This path refers to the location of the iODBC driver manager files. The default location is: <TDV_install_dir>\apps\odbc\<platformType>\lib

3. Add these variables to the users profile that will be accessing the ODBC driver.

Creating a DSN with driverConfig on UNIX

A DSN is the logical name that is used by ODBC to access data. You can use an ODBC DSN entry to store the connection string values externally, to minimize the complexity of the connection string that you must define in your program.

You can create a DSN using the configuration utility driverConfig. This configuration utility helps you to reconfigure the driver files and create, edit, list, or delete DSN entries. You can use it when the driver file location has changed or is to be changed after installation.

On UNIX platforms, SysV semaphores are used to synchronize the read and write operations, and they are never deleted by ODBC drivers. The ODBC driver might run into an error if it is unable to create a new one because the maximum SysV count has been reached.

You can clean up semaphores using the ipcrm command.

To create a DSN using driverConfig

1. Make sure that you have Read and Write permissions on the following files:

```
odbc.ini
```

```
odbcinst.ini
```

2. Locate driverConfig.

3. Run the utility using the following command:

```
driverConfig
```

For example:

```
./home/release/apps/odbc/linux64/bin/driverConfig
```

```
./odbc/linux64/bin/driverConfig
```

```
./bin/driverConfig
```

4. Supply driverConfig with responses to set configurations in the odbc.ini and odbcinst.ini.

The Usage instructions of driverConfig is given below:

```
Usage: driverConfig [options]
```

where options are:

```
[-help]
```

```
[-view]
```

```
[-deleteDSN <name>]
```

```
[-uninstallDriver]
```

```
[-installDriver <driver path>]
```

```
[-configDSN <DSN attributes>]
```

The table below describes the different options:

Option	Description
help	Show this information and exit
view	View configuration and DSNs on this system and exit
deleteDSN <name>	Delete the named DSN
uninstallDriver	Uninstall ODBC 3.5 64-bit Driver
installDriver <driver path>	Install ODBC 3.5 64-bit Driver
configDSN <DSN attributes>	Create a new ODBC 3.5 DSN, where DSN attributes are of the form: "DSN=test;host=localhost;port=9401;uid=userId;pwd=password;domain = composite;datasource=ds;catalog=cat"

Note:

- If this utility is executed without any options then it will run in interactive mode using Terminal I/O
- All option names are case-insensitive.
- If -help or -view is one of the options used, then no other options that are included, will be executed.
- The utility accepts more than one option.
- The order of actions that will be executed is:
 - deletedsn
 - uninstallDriver
 - installDriver
 - configDSN
- If the environment variable COMPOSITE_HOME is defined, then it is used for evaluating the default ODBC driver path. COMPOSITE_HOME is used for evaluation only during driver installation in interactive mode.

Connecting SAS System to TDV ODBC

The SAS system must be installed and functional, for example in the following location:

```
/opt/sas
```

For more information, see the installation instructions for the SAS System for Unix.

To configure the connection between the TDV server through ODBC from SAS

- 1. Make sure that the SAS System is installed and has write access to SAS environment sasenv_local configuration script.
- 2. Adjust the SAS environment script sasenv_local by adding the necessary environment variables. The script is typically located at:

```
/opt/sas/SASFoundation/9.3/bin/sasenv_local
```

Make sure to set the following environment variables:

Variable	Description
COMPOSITE_HOME	COMPOSITE_HOME=/usr/local/composite/
	export COMPOSITE_HOME
ODBCHOME	ODBCHOME=/usr/local/unixODBC
	export ODBCHOME
ODBCINI	ODBCINI=/usr/local/unixODBC/etc/odbc.ini
	export ODBCINI

Variable	Description
ODBCINSTINI	ODBCINSTINI=/usr/local/unixODBC/etc/odbcinst.ini export ODBCINSTINI

3. Test the SAS/Access to ODBC connectivity using the Composite data source by running the following command:

```
libname comptest odbc datasrc=COMPDEV;
```

Your results should be similar to:

NOTE: Libref COMPTTEST was successfully assigned as follows:

Engine: ODBC

Physical Name: COMPDEV

Configuring TDV for Using a JMS Broker

Java Message Service (JMS) provides a way to publish message-based Web services. By default the installation of TDV supports Sonic and TIBCO JMS brokers, but a few drivers must be copied from the JMS broker installation to the installed directory of TDV to connect the two servers. TDV can also work with other message queues through its open API.

- [Configure Communications between TDV and the JMS Broker](#)
- [Adding JMS Connectors to the TDV Server](#)

Configure Communications between TDV and the JMS Broker

To enable communications between TDV and the JMS broker, several JAR files must be obtained. TDV supports connection to JMS through Java Naming and Directory Interface (JNDI). TDV ports using JMS can only be configured with a queue destination type, but procedures and triggers can use topic connections factories.

To configure communication with JMS brokers

1. Find and copy the following files.

JMS Type	File to Copy	From
Sonic MQ	mfcontext.jar sonic_<x>.jar	Sonic installation directory
TIBCO MQ	tibjms.jar	The TIBCO installation

2. Paste those files into the directory:

```
<TDV_install_dir>\apps\server\lib
```

3. Restart the TDV Server.
4. Configure your JMS broker according to manufacturer instructions.
5. Make sure that the following JMS provider objects are created:
 - A suitable QueueConnectionFactory (QCF)
 - A suitable Queue
6. Register the QCF and the Queue with the JNDI

You can now add your JMS connectors to the TDV Server.

Adding JMS Connectors to the TDV Server

JMS connectors must be configured so that the TDV Server can publish data services using JMS through JNDI after the JMS Broker has been configured.

To add JMS connectors to the TDV Server

1. Launch Manager.

- From a URL, locally or remotely, type:

```
http://localhost:9400/manager/#home
```

```
http://[Host]:[BasePort]/manager/#home
```


- From the Studio, select Administration > Launch Manger (Web).

2. Log in to the Manager as administrator.

3. Choose Connectors from the CONFIGURATION menu.

4. On the CONNECTOR MANAGEMENT page, click Add Connector.

5. The Add a JMS through JNDI Connector window is displayed. Publishing directly to JMS is not supported.



6. Enter values in the fields displayed on the Info tab.

Field	Description
Connector Name	Enter a name to call the connector.
Group Name	Connectors that share a group name use a common connection pool, with the added advantage of failover; that is, if a connector instance fails, other connectors in the group can send and receive messages using the same connection pool.

Field	Description
Annotation	(Optional) Adds notes for the JNDI connector that are visible on the Connector Management page.

7. Enter values in the fields displayed on the JMS through JNDI tab.

Field	Description
Initial Context Factory	Typically, the JNDI initial context factory is the class name. Type <code>c</code> to see the following default string values: - For Sonic—<code>com.sonicsw.jndi.mfcontext.MFContextFactory</code> - For TIBCO—<code>com.tibco.tibjms.naming.TibjmsInitialContextFactory</code> - For OpenMQ—<code>com.sun.jndi.fscontext.RefFSContextFactory</code>
JNDI Provider URL	URL for connection with the JNDI. A TCP protocol is generally used. The TIBCO default port is 7222, and the Sonic default port is 2506. Make sure that the port in the firewall is open 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 TDV connections with the JMS broker.
Connection Factory	Queue or topic connection factory name. For multiple queue or topic connection factories, create additional connectors.

8. Enter a name-value pair on the JNDI Properties tab.

9. Click the plus button to add more name-value pairs.

Sonic requires you to specify a domain name; TIBCO does not.

10. Enter values on the Pool tab.

The Pool tab lets you specify connection thread timeout and pool size parameters. The default values are typical for development needs.

Name	Value description
Pool Timeout	Maximum waiting time (in seconds) for a new connection. If a connection is not provided within this period, the service checks for an available connection through a valid user and uses it. If no connection is available, the least recently used connection for another user is dropped and a new connection for the required user is opened.
Minimum Pool Size	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 through JNDI, the pool creates one connection based on the information provided in the Info panel. To improve response, connections remain available even when there is no activity. After a period of JMS connection inactivity, the pool size begins to shrink.
Maximum Pool Size	The number of connections (active and idle) available on the data source. When the limit is reached, new incoming requests must wait for the next available connection. Connectors with identical group names use the same pool of connections.

11. Click OK.

For more information on publishing to JMS queues, see the *TDV User Guide*.

Configuring TDV for AIX Platforms

TDV works with AIX platforms, but some additional configuration is required. See the sections below if you experience these issues:

- [Improving Studio Response Times for AIX Connections](#)

Improving Studio Response Times for AIX Connections

If you experience slow response times and cannot log in to Studio, follow the steps below.

To improve Studio connection times on AIX

1. Add the following attribute to the <TDV_install_dir>/conf/server/server_values.xml file:

```
<common:attribute>
```

```
<common:name>/server/config/net/useBlockingIOConnectors</common:name>
```

```
<common:type>BOOLEAN</common:type>
```

```
<common:value>>true</common:value>
```

```
</common:attribute>
```

2. Restart TDV Server to make the change effective.

System Monitoring with Studio Manager

TDV system monitoring entails tracking an analysis of system activities, system status, events, and system data. TDV provides two interfaces for system monitoring:

- Studio Manager—in Studio as documented in this topic
- Manager—a Web browser interface (see [System Management with Manager](#) for more information).


Studio Manager displays summary views of TDV status, server information, cached resources, data sources, requests, sessions, transactions, triggers, and event logging. System event and log monitoring is covered in [System Event and Log Monitoring](#)

The following topics are covered in this chapter:

- [Studio Manager Window and Toolbar Overview](#)
- [Using Studio Manager](#)
- [Studio Manager UI Reference](#)

Studio Manager Window and Toolbar Overview

The following graphic identifies the main components of the Manager window.



- **Manager Menus**—The menu options are specific to the Studio Manager.
- **Manager Toolbar**—The menu options are specific to the Studio Manager.
- **Manager Console List**—A list of available consoles for you to monitor and manage TDV activities. When you select a console, Manager displays the relevant information in the right pane.
- **Manager Console**—Displays the summary statistics and details about the selected Manager console.
- **Console Summary**—Displays summary statistics for the selected Manager console.
- **Console Details**—For Server Overview, you see the status console for all consoles. For the other consoles, you see detailed real-time status information.
- **Status Bar**—Provides current TDV status information.

Studio Manager Toolbar

The Studio Manager toolbar contains buttons that generally can be used for all of the manager panels. You can hover your cursor over the buttons to view tooltips that explain what the button does. The Studio Manager Toolbar toolbar is displayed just below the Studio menu bar.

The table below describes the use of each button. The buttons are listed as they appear on the toolbar from left to right.

Menu Option	Use to...
Full Server Backup	Open the Full Server Backup dialog to back up all resources. See the “Resource Management Basics” in the <i>TDV User Guide</i> for more information.
Configuration	Open the Configuration dialog to access the TDV configuration parameters. See the “TDV Configuration Parameters” in the <i>TDV User Guide</i> or the <i>TDV Reference Manual</i> for more information.
Studio Log	Opens the Studio log which is a log of all activities. See TDV Logging Information .
Refresh Rate	Choose the frequency, in seconds, that you want to refresh the Studio Manager panels.
Refresh Now	Click to immediately refresh the Studio Manager panels.

Using Studio Manager

This section describes how to use many of the basic features that appear in Studio Manager. Many of the panels in Studio Manager include features to help you tailor the current display:

- [Launching Studio Manager](#)
- [Selecting Columns for Display](#)
- [Viewing Table Row Details](#)
- [Sorting Rows](#)

- [Customizing Filters for Studio Manager](#)
- [Configuring the Columns on the Cached Resources Panel](#)
- [Enabling and Disabling Caches in Studio Manager](#)
- [Modifying the Cache Schedule in Studio Manager](#)
- [Refreshing a Cache in Studio Manager](#)
- [How to Troubleshoot Cache Refresh](#)
- [Configuring Time for Requests to Stay Active on the Studio Manager Request Panel](#)
- [Scheduling Data Source Connection Testing](#)


Launching Studio Manager

Studio Manager is one of the three main tabs in Studio. It gives you access to a number of individual panels that you can use to understand, analyze, and manage TDV performance and system utilization.

To launch Studio Manager

1. Start Studio.
2. Click Manager on the left edge of Studio.

Studio displays the Manager page:



Selecting Columns for Display

Your choices about which columns to show or hide are saved and returned when you restart the Studio.

To specify the columns you want to have displayed

1. Open Studio Manager.
2. On the Server Overview page, in the table, right-click the header, Console or Status.
3. Select the columns to be displayed from these options:

Option	Description
Show Default Columns	Displays those columns that are considered default in the system.
Only Show This Column	Lets you specify any one column to be shown. You can choose to show/hide any column you think is relevant for your needs.
Show All Columns	Displays all the columns.

4. Use the Show Default Columns option to reset the column choices to their original settings.

Viewing Table Row Details

To view the details of a table row

1. Open Studio Manager.
2. Double-click the row in the table for which you want details.

Or select a row and click Show Row Detail.

Sorting Rows


Any sorting you specify is automatically saved and reused when you restart the Studio.

To sort rows using the advanced sort

1. Open Studio Manager.
2. Select any option with a Details table such as Cached Resources, Data Sources, Events, Requests, Sessions, Transactions, or Triggers.
3. Click the column header by which you want to sort the rows.
For example, click Time to sort the rows by the Time column.
4. Locate and click the Sort icon on the page.
The Advanced Sort Dialog opens.
5. Use the Add icon to add sort criteria.
6. Make choices for the following options:.

Option	Description and Action
Sort By	Lists the columns displayed in the table view
Direction	Direction, ascending or descending, by which to sort the table

Option	Description and Action
	entries
Move Rule	Direction, upward or downward. Changes the order in which sort rules are applied. Click up or down in the Move Rule column to further filter the sort order.




- Click OK after you have made all the specifications.

Customizing Filters for Studio Manager

Filter definitions can be created for in Studio Manager for these options: Cached Resources, Data Sources, Events, Requests, Sessions, Transactions, and Triggers. Filter definitions in addition to the filter that is currently being used on a console are saved when you exit the Studio. The filter will be available when you restart the Studio and Studio will automatically reuse the same active filter for the specific console.

To filter data for displaying in the table view

- Open Studio Manager.
- Select any option with a Details table such as Cached Resources, Data Sources, Events, Requests, Sessions, Transactions, or Triggers.
- From the Filter field, select Edit Filters.



4. Click the green plus-sign to add a filter.
5. Name your new filter.
6. Make choices for the following options.

Option	Description and Action
Filter By	Lists the columns displayed in the table view. You can select a column in this list to choose the column on which to apply the filter.
Operator and Condition	Work together as the two sides of an equation with the column. Operator lists a set of conditions for your selection and Condition lets you specify the value for the Operator. Specify your condition in the Operator and Condition columns.
Match Any Rule	Find data that matches any of the rules defined. This option typically returns more results and might take more time.
Match All Rules	Find data that matches all of my defined rules. This option typically returns a smaller set of data.

7. Click Use Selected Filter or Do Not Filter.
All filters are saved.

Configuring the Columns on the Cached Resources Panel

You can show and hide columns that are displayed on the Cached Resources page in Studio Manager. For example, if you want the End column to display, you can use this

procedure to add that column to the display.

To configure columns

1. Open Studio Manager.
2. Select Cached Resources.
3. Right-click on any of the column headings.
4. Select or clear the columns from the list of values.

Enabling and Disabling Caches in Studio Manager

To change cache enabling in Studio Manager

1. Open Studio Manager.
2. Select Cached Resources.
3. In the table view of the console, select the event to be scheduled for caching.
To select multiple views, hold down the Shift key or Ctrl key and select the views. The Shift key lets you select adjacent rows, and the Ctrl key lets you select any row.
4. Use the Change Enabling button to disable or enable caching for that view.
The UP status changes to DISABLED when you click the Change Enabling button. The DISABLED status changes to UP when you click the Change Enabling button.

Modifying the Cache Schedule in Studio Manager

To modify a caching schedule in Studio Manager

1. Open Studio Manager.
2. Select Cached Resources.
3. In the table view for Cached Resources, select the event to be scheduled for caching.
4. Click Schedule.

The Cache Schedule window opens.

5. In the Status section, select the Enable check box, if it is not selected.
6. For more information on scheduling and caching options, see “TDV Caching” in the *TDV User Guide*.
7. Click OK.

Refreshing a Cache in Studio Manager

To refresh a cache in Studio Manager

1. Open Studio Manager.
2. Select Cached Resources.
3. In the table view of the console, select the event to be scheduled for caching.
To select multiple views to refresh their respective cache, hold down the Shift key or Ctrl key and select the views. The Shift key lets you select adjacent rows, and the Ctrl key lets you select any row.
4. Click Refresh Cache.

Cancel Cache Refresh

In-progress cache refreshes may be canceled manually. Click on the Cancel button that can be found above the result set. You can also call a REST API from the command line to cancel in-progress cache refreshes. Cancellation is allowed only by the users that initiated the refresh or member of the admin group.

How to Troubleshoot Cache Refresh

Caching is often needed in TDV implementations, often with scheduled cache refreshes. These refresh activities can fail or hang for a number of reasons. In tightly controlled environments such as user acceptance testing or production, failures can be difficult to diagnose.

This topic assumes a typical production logging level, and also assumes that it is impractical to change the debug setting if it requires a TDV restart to take effect.

When executing a cache refresh, TDV performs several steps:

- **Invokes the Cache Process**—How is the cache refresh initiated, by schedule or by event?
- **Reads the Source Data**—What data sources are supplying the data to cache?
- **Computes the Result Set**—What calculations does TDV perform to generate the result set? Is a federated join algorithm used?
- **Writes to the Cache Target**—How does TDV write the result set to the cache database? Does it use SQL INSERT or some advanced mechanism?
- **Completes the Refresh Process**—Once the cache table has been fully updated, TDV cuts over to the new data. Has transactional integrity been maintained?

If you suspect a cache refresh is failing, try to locate the problem among these steps. The two main log files that contain cache-related errors are `cs_server.log` and `cs_server_task.log`.

Invokes the Cache Process

To see information about the most recent cache refreshes, you can do the following:

- Click the Studio's Manager tab on the left, and select **Cached Resources** from the list. This window lists the names of all recent refreshes and when they were launched.
- On the same page, select a cache refresh from the **Details** list to see more information for that refresh.
- Click the Studio's Modeler tab on the left, navigate to **Composite Data Services > Databases > System** in the resource tree, and locate the **SYS_CACHES** system table.

For a list of the fields in the **SYS_CACHES** table, refer to the **TDV System Tables** chapter of the *TDV Reference Manual*.

- In the **SYS_CACHES** table, **INITIAL_TIME** indicates when the refresh was kicked off, so you can tell whether it was triggered as expected. The table also tracks **NUM_SUCCESS** and **NUM_FAIL**.
- Check the **cache_tracking** and **cache_status** tables.

For more information about these tables, refer to the **TDV Caching** chapter of the *TDV User Guide*.

- If your TDV implementation has enhanced logging capability, such as the open source KPI Module (https://github.com/TIBCOSoftware/ASAssets_KPI), you can log the cache refresh history. This provides a record of each cache's refresh processing time.

Reads the Source Data

TDV issues one or more SQL statements to fetch data from data sources. Under normal logging settings, these fetches are not logged. To see them, open the view's execution plan and look for FETCH nodes. If one of the data sources does not seem to be responding normally, you can copy a FETCH node's SQL and execute it in a client native to the database (such as Oracle SQL Developer).

Note: This technique is not reliable for clustered data sources, because system changes may have been made since the problem occurred (such as a bad node being taken out of service), meaning that you cannot reliably reconstruct what TDV encountered.

On the rare occasions when a data source connection issue involves the `cache_status` table, the cache refresh can hang. For a remedy, read the Managing Cache Status Table Probe Row Conflicts section of the TDV Caching chapter of the *TDV User Guide*.

Computes the Result Set

After the source data has been retrieved, TDV computes the requested result set using SQL statements, procedural logic (even for cached views), and so on. A prolonged CPU spike can indicate computation problems, such as a runaway thread. Unfortunately, you cannot tie a CPU spike to a particular request or query. However, you can see a given request's *memory* consumption. If a request's memory consumption is high and stays high throughout the request's life cycle, something is probably wrong.

- Click the Studio's Manager tab on the left, and select Requests from the list. This window lists the names of all recent requests and when they were launched. You can use the start time and TRUE in the Cache column to help find the request associated with the cache refresh.
- You can double-click a request and examine its Full SQL to verify that you have found the correct one.

You will not be given the fully qualified path to the cached resources (the best identifier), but you can usually make a confident guess. After you have identified the request, you can see its Memory and Max Memory consumption.

Writes to the Cache Target

When you troubleshoot a cache refresh, always check the writing step.

With the result set computed, TDV begins writing to the cache database. By design, TDV does not always wait for request completion. Instead, TDV may start streaming partial results to the cache database as soon as they are available.

You can determine whether the first batch of records has been inserted into the cache table. Use the `cache_tracking` and `cache_status` tables to determine the correct target table. (For multitable caching, pick the appropriate bucket.) When querying this table, set transaction isolation to `READ UNCOMMITTED`, either from TDV Studio or from a native database client.

- If the target table appears to be fully populated, but the cache refresh hangs, TDV may have been unable to cut over from the old cache data to the new cache data. See [Completes the Refresh Process](#).
- Always check for errors related to table spaces that have reached their maximum—a common occurrence. When this happens, TDV hangs, because it is unable to commit a transaction.
- If multitable caching is configured, TDV drops indexes before inserts and recreates them afterward.

Completes the Refresh Process

After the cache table is fully populated, TDV performs a cutover. It updates the `cache_tracking` and `cache_status` tables to start using the new cache data. If the old cache data is currently being used to service a request, it continues to do so. Only requests issued after the cutover use the new cache data. When TDV determines that the old cache data is no longer needed, it purges the old cache.

Cancel Cache Refresh

In-progress cache refreshes may be canceled manually. Click on the Cancel button that can be found above the result set. You can also call a REST API from the command line to cancel in-progress cache refreshes. Cancellation is allowed only by the users that initiated the refresh or member of the admin group.

Configuring Time for Requests to Stay Active on the Studio Manager Request Panel

There are several configuration parameters that control how long requests are maintained on the Request panel. By manipulating the configuration parameter values you can configure how long requests are maintained.

If you do not see the completed requests in the Requests panel in Studio Manager, typically it means they were purged after the Request Purge Period expired.

To configure the length of time that requests stay on the Request panel

1. Open Administration > Configuration.
2. Search for and check the value of the following settings:

Parameter	Description
Maximum Requests Tracked	This is the maximum number of requests tracked. For example: 10000. Changing this value will have no effect until the next server restart.
Request Purge Period	Controls how often the server cleans out completed requests that are older than the purge period. For example: 5 minutes.
Maximum Sessions Tracked	The maximum number of simultaneous sessions. For example: 10000. Use '0' for no limit.
Session Purge Period	Controls how often the server cleans out closed sessions that are older than the purge period. For example: 5 minutes.
Session Idle Purge Period	Controls how often the server's session reaper attempts to end idle sessions.

If you used the example settings described in this task, after the 5 minute purge period, ALL completed requests that ended more than 5 minutes ago and ALL closed sessions that ended more than 5 minutes ago are purged.


Scheduling Data Source Connection Testing

If you manage many data sources using Studio, it can be good to validate the connections to those sources on a regular or scheduled basis. If a data source connection returns invalid results, it might mean that the data has gone dormant or has been moved to a new location which could invalidate any views that you have that depend on fresh data from that source. The scheduled tests can be configured to send notification emails.

If you do not see the completed requests in the Requests panel in Studio Manager, typically it means they were purged after the Request Purge Period expired.

To schedule data source connection tests with email notifications

1. Open Administration > Configuration.
2. Search for SMTP in the Configuration search field.



3. Depending on the SMTP requirements for your environment, configure valid values for the necessary parameters.
 - SMTP Authentication Required
 - SMTP Authentication User Name
 - SMTP Authentication Password
 - SMTP Host Name
 - SMTP Port
4. In the Configuration window, locate From Address.
5. In the Value field, type the address that you want to appear on notification email messages sent from the TDV products.
6. Open Manager.

7. Select Data Sources.
8. Click Schedule Test.
9. Set the frequency that you would like the tests run.
10. Specify the email address for the location that notifications should be sent to.

Studio Manager UI Reference

- [Server Overview Panel](#)
- [Cached Resources Panel](#)
- [Data Sources Panel](#)
- [Events Panel](#)
- [I/O Panel](#)
- [Memory Panel](#)
- [Requests Panel](#)
- [Sessions Panel](#)
- [Storage Panel](#)
- [Transactions Panel](#)
- [Triggers Panel](#)

Server Overview Panel

The Studio Manager Server Overview window displays this summary information for all current and recent sessions:

- **Server Name:port#**—The name of the TDV server and its port number.
- **Total Memory Used**—Percentage of total available Java Heap Memory (RAM) currently in use. Total memory is further divided into Managed and Reserved memory with a built in margin to prevent OOM errors.
- **Total Sessions**—Total number of sessions started since the server started.
- **Active Sessions**—Number of currently active sessions.

- **Total Server Requests**—Total number of requests made to the TDV server since the server started.
- **Active Server Requests**—Total number of currently active requests made to the TDV server.
- **Total Data Source Requests**—Total number of requests made to the data sources since the server started.
- **Active Data Source Requests**—Total number of currently active requests made to the data
- **Maximum Viewable Events**—Maximum number of events that can be viewed from the Events console in Studio Manager. In Studio, you can set this number using the Configuration dialog setting for Maximum Viewable Entries.
- **Maximum Event Entries**—Maximum number of events to be stored in the TDV repository. When the number of events reaches this threshold the oldest events are discarded in FIFO (first in, first out) order. The log files are generally configured to retain a more expansive archive of event entries. In Studio, you can change the Maximum Event Entries using the Configuration window setting for Maximum Log Entries
- **Privilege Cache Access**—The percentage of the privilege cache hits of the total number of accesses.
- **Privilege Cache Capacity**—The percentage of the privilege cache that is being used. In Studio, you can change privilege cache capacity setting in the Configuration window at Privilege Cache Size (On Server Restart).
- **User Cache Access**—The percentage of the user cache hits of the total number of accesses.
- **User Cache Capacity**—The percentage of the user cache that is being used. In Studio, you can change the user cache capacity setting in the Configuration window at User Cache Size (On Server Restart).
- **Repository Cache Access**—The percentage of the repository cache hits of the total number of accesses.
- **Repository Cache Capacity**—The percentage of the user cache that is being used. In Studio, you can change the user cache capacity setting in the Configuration window at User Cache Size (On Server Restart). In Studio, you can change the repository cache capacity setting in the Configuration window at Metadata Cache Size (On Server Restart).

- Clear Repository Cache button—Immediately empties the repository cache. Clearing the repository cache requires the Modify All Status right. The button is visible, but grayed out for any other users who can have access to the Manager.

The Server Overview Details table displays the status of all of the system resources and activities.

You can also get server overview information in Manager. See [SERVER OVERVIEW Page](#) for more information.

Cached Resources Panel

The Studio Manager Cached Resources panel displays:

- The status summary for existing caches at the top of the panel.
- Details about each cache in the middle of the panel.
- The status of caches refreshes at the bottom of the panel.
- Total Cache Refresh Failures—Total number of cache refresh failures during this session.
- Storage Used—Amount of storage currently in use by cached resources.

You can also monitor cached resources in Manager. See [CACHED RESOURCES Page](#) for more information.

Data Sources Panel

The Studio Manager Data Sources panel displays:

- The status of the existing data source activity at the top of the panel.
- Details about the status of each of the data sources at the bottom of the panel.

Studio Manager displays this information about data sources:

- Total Requests—Total number of requests made to the data sources since the server started.
- Active Requests—Total number of currently active requests made to the data sources.

- Bytes From Data Source (Estimated)—An estimate of the total number of bytes of data sent to the data sources since the server started.
- Bytes Into Data Source (Estimated)—An estimate of the total number of bytes of data received from all data sources since the server started.
- Test All Now button—Enables manual test of the availability of all resources. Use requires the Modify All Status right.
- Schedule Test button—Enables configuration of a time interval that should occur between automated attempts to test all of the data sources.

When you press the Schedule Test button, the Test All Data Sources Schedule window opens, where you can set the time and the interval at which the testing should occur. If the Do Not Execute Automatically button is selected in the Test All Data Sources Schedule window, then no automated testing of all of the data sources will occur. If an individual data source is tested at a given time (using the Test button), that test time will override this setting. Disabled data sources will fail this test.

- Next Test Time—Time when the next test is scheduled to run.

See [Data Source Details in Studio Manager](#) for more information.

You can also monitor data sources in Manager. See [DATA SOURCES Page](#) for more information.

Data Source Details in Studio Manager

Above the table in the Data Sources panel, you can use the buttons and controls to manage data sources:

- Info button—Opens the Data Source Information panel and enables basic and advanced configuration of the individual data source. Data source connection information can be edited directly from this panel, and tables can be added or removed from the definition.
- Change Enabling button—Enables or disables the data source. Enabling makes the data source accessible through TDV definitions and configurations. Disabled takes the data source offline and makes it inaccessible to TDV defined channels.
- Clear Pool button—Clears currently allocated pool connections, dropping the current connection pool to allow current processes to restart connections when necessary.
- Test button—Verifies connection status of the data source.

Events Panel

The Studio Manager Events panel displays the status of recent events.

- **Maximum Viewable Events**—Maximum number of events that can be viewed from the Events console in Studio Manager. In Studio, you can set this number using the Configuration dialog setting for Maximum Viewable Entries under Memory.
- **Maximum Storable Events**—Maximum number of events that can be stored from the Events console in Studio Manager. In Studio, you can set this number using the Configuration dialog setting for Maximum Log Entries under Database Logger.

See [Event Details in Studio Manager](#) for more information.

You can also monitor events in Manager. See [About Events](#) for more information.

Event Details in Studio Manager

Every individual event has additional detailed information available. Studio displays these event details:

- **Time**—The date and time the event occurred.
- **Description**—A description of the event, such as the request id.
- **Type** —The type of event that occurred which can be anything in the event lifecycle.
- **User**—User who generated this event.
- **Domain**—Domain to which the owner of the resource that triggered the event belongs.
- **Attributes**—An arbitrary list of additional properties specific to the type of event.
- **SNMP ID**—Unique SNMP ID that describes the context of the event.
- **ID**—Unique event identifier.
- **Parent ID**—Unique identifier of the parent event.

I/O Panel

The Studio Manager I/O panel displays a graph showing input and output of data between data sources and clients and TDV over time.

You can adjust the information displayed in the I/O panel by selecting or clearing the check boxes at the bottom of the panel:

- Total Input/Output—Aggregate data input/output from the server. This is the sum total of the TDV data service activity and the data source activity.
- Data Service Input/Output—Total requests made between the TDV Server and the clients.
- Data Source Input/Output—Total requests made between the data sources and the clients. These are the numbers displayed in the Bytes Into Data Source (Estimated) and Bytes From Data Sources (Estimated) fields at the top of the Data Sources panel.

You can also monitor input and output in Manager. See [I/O Log](#) for more information.

Memory Panel

The Studio Manager Memory panel displays a graphical log of TDV memory usage.

You can adjust the information displayed in the I/O panel by selecting or clearing the check boxes at the bottom of the panel:

- Total Memory check box—The actual TDV usage levels.
 - Maximum—Available computational memory.
 - Throttle—Displays the maximum throttle memory available. TDV has a runtime configuration setting for a wait queue minimum memory threshold which is visible as the “Throttle” on the Memory console. If actual memory usage crosses the minimum threshold, all new queries are queued until more memory is available. This threshold is meant to avoid potentially fatal out of memory errors on the server.
- Managed Memory check box—Managed memory is a configurable percentage of that total made available to the Java Virtual Machine (JVM).
 - Maximum—Displays the maximum total memory available as set by the Available Managed Memory configuration parameter.

You can use the Free Unused Memory button at the bottom of the panel to free unused memory. This action starts the Java VM garbage collection cycle that is started automatically when the maximum managed memory level is exceeded. You must have the `Modify_All_Status` right to use this button.

You can also monitor memory usage in Manager. See [Memory Log](#) for more information.

Requests Panel

The Studio Manager Requests panel displays:

- Summary information about requests activity at top of the panel.
- Details about the status of each request at the bottom of the panel.
- 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.
- Server Requests (Total and Active)—The total number of requests made to the server since the server was started and the number of currently active requests.
- Data Source Requests (Total and Active)—The total number of requests made to the data sources since the server was started and the number of currently active requests.
- Internal Requests (Total and Active)—The total number of internal requests made to the server since the server was started and the number of currently active requests.
- Total Bytes From Data Source (Estimated)—An estimate of the total number of bytes of data sent to the data sources since the server started.
- Total Bytes Into Data Source (Estimated)—An estimate of the total number of bytes of data received from all data sources since the server started.
- 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 can be used for notification.
- Memory Utilization—Amount of memory currently in use by TDV.
- Clear Plan Caches button—Clears all query plan caches, in the event that current statistics gathering can 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 button**—Immediately removes all completed and failed requests. This is useful to reset the view prior to testing a set of requests.

Request Details in Studio Manager

In the Request Details pane, you can select the row and then perform these actions:

- **Cancel button**—Clears the request.
- **Show Query Plan button**—View the query plan for a request. This action displays the full query plan of the request with statistics. You can use the plan to diagnose any type of query.
- Click the **Info button** to view the row details.

By default, the requests in the Details table are sorted in the order they are received (oldest first). See [Request Details in Studio Manager](#) for more information.

You can also monitor requests in Manager. See [REQUESTS page](#) for more information.

Sessions Panel

The Studio Manager Sessions window displays this summary information for all current and recent sessions:

- **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.
- **Total Sessions**—Total number of sessions started since the server started.
- **Active Sessions**—Number of currently active sessions.
- **Studio Session Timeout**—Timeout limit (in minutes) for JDBC/ODBC clients.

You can also monitor sessions in Manager. See [SESSIONS page](#) for more information.

Storage Panel

The Storage Graph provides a Status icon at the top of the log:

Storage Status	Description
OK	The storage used is within the set thresholds.
EXCEED	Maximum threshold was exceeded.
RESET	Maximum threshold was reset.
FAIL	Out of disk space.

The lines on the graph are:

- Used Disk Space—The total usage.
- Maximum—Total amount of available disk space.
- Low Disk Warning Threshold—Threshold at which a warning will be issued if the storage becomes lower than the threshold.
- Low Disk Critical Threshold—Threshold at which a the storage becomes critically low.

You can also monitor storage in Manager. See [Storage Log](#) for more information.

Transactions Panel

The Studio Manager Transactions panel displays:

- Summary information about transactions activity at top of the panel.
- Details about the status of each transaction at the bottom of the panel.

Note: Studio Manager only shows current transactions in the Details table. Manager displays the same information about transactions as Studio Manager but includes recent transactions in addition to current transactions. See [TRANSACTIONS Page](#).

Triggers Panel

The Studio Manager Triggers panel displays:

- Summary information about triggers activity at top of the panel.
- Details about the status of each trigger at the bottom of the panel.

You can also monitor transactions in Manager. See [TRIGGERS page](#) for more information.

The Studio Manager Triggers window displays this summary information for all current and recent sessions:

- **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**—The total number of triggers processed since the server was started.
- **Total Failed Runs**—The total number of triggers that failed since the server was started.
- **Resource**—The name of the trigger.
- **Path**—Fully qualified path to the resource that has been scheduled for execution.
- **Type**—Type of trigger. Timer Event, System Event, or User-Defined Event.
- **Total Runs**—Total number of times the this trigger was invoked.

System Management with Manager

Manager provides a thin client Web-page based interface for managing TDV. Manager displays summary and detail views of TDV status, server information, cached resources, data sources, requests, sessions, transactions, triggers, and event logging. These monitoring capabilities are also provided in Studio Manager which is an integral part of Studio. See [System Monitoring with Studio Manager](#) for more information.

The following topics are covered:

- [Using Manager](#)
- [Manager UI Reference](#)
- Domain, group, and user management including configuring LDAP servers and selected groups for use with TDV. Domain, group, and user management are described in:
 - [Composite Domain Administration](#)
 - [LDAP Domain Administration](#)
 - [Dynamic Domain Administration](#)
- Resource Management privilege settings, dependency privilege analysis, and modifying privileges on resources and their dependencies are described in [Managing Security for TDV Resources](#).
- Cluster creation, configuration, and monitoring. Cluster management is described in the *TDV Active Cluster Guide*.
- SSL keystore management. See [Configuring the Java Keystore File](#) .
- How to configure the JMS connector required to use the JMS Broker is described in [Adding JMS Connectors to the TDV Server](#).
- Row-based security configuration is described in [Managing Security for TDV Resources](#).

Manager is available only to users with administrative rights (Access Tools right is the minimum right required to view the Manager).

Using Manager

Many Manager pages have shared 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.

- [Launching Manager](#)
- [Refresh the Current Page](#)
- [Sort with Manager](#)
- [Filter Table Data](#)
- [Creating a New Table Filter](#)
- [Copying an Existing Filter](#)

Launching Manager

Manager enables users with appropriate TDV rights to view, monitor, and update selected TDV summary views and status. Additionally, authorized users can perform some server management tasks, establish and maintain active clustering, and manage domains, groups, and users and their associated TDV rights, and others. The Access Tools right is the minimum right required to view Manager.

To launch Manager

1. Launch Manager from Studio:

Administration > Launch Manager (Web)

Or direct a Web browser to the Manager using one of these URLs:

```
http://localhost:9400/manager (when TDV is locally installed)
```

```
http://[TDV_HostName]:[PortNumber-Default9400]/manager
```

2. If you are using Internet Explorer, turn Compatibility View mode on.
Manager launches and displays a login dialog.
3. Login with a user name and password with administrative rights.

If an OAuth 2.0 domain has been set up in Web Manager and access tools privileges given, then you can login to TDV Manager by choosing the domain name from the drop down or by entering the name in the Domain field. You will not need to provide username and password while using the OAuth 2.0 domain.

Notes:

- You need to be actively logged into your IdP account, before attempting to login to any of the TDV client tools using the OAuth2 domain. If not, you will be redirected to your IdP login page. Once logged in, click on the Oauth2 domain in TDV Studio, Web Manager or Web UI to successfully login to the TDV client.
- In order to allow access to the Web Manager, the TDV Web Manager url must be added to the Redirect URIs list in your Identity Provider. Using the following hostname based url is the preferred method:


```
http://xyz.com/manager  
https://xyz.com/manager
```

You may also add an IP based or localhost based url in the Redirect URI list as given below:

```
https://1.1.1.1:9402/manager  
http://localhost:9400/manager  
https://localhost:9402/manager  
http://1.1.1.1:9400/manager
```

However, it is important to note that this is not a preferred approach as:

- localhost references are not host specific.
 - IP addresses will not work in deployment environments where server IP addresses are not statically assigned (For example, in the cloud deployments).
- After successful login, the MANAGER HOME page is displayed.



- Logging out of TDV Manager invalidates the token that was used to authorize the user. Some providers do not support access token revocation. Refer to your IdP documentation for the specific help regarding renewal/revocation of access and refresh tokens.

Refresh the Current Page

Most Manager 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 refresh the page at any time by clicking the green arrow “Refresh Now” icon.

Sort with Manager

Sortable table column headers display a small white arrow 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.

You can change the table row display sort order by clicking the Sort... icon: . Manager displays the Advanced Sort dialog for you to define the sorting rules.



Filter Table Data

You can use the Filter setting above most tables to filter the data displayed based on the column data. To apply a filter to the table, choose one from the Filter drop-down list.

For tables, only the Show All and <Edit Filters...> options are offered.

When working with table filters, these conventions apply:


- Filter definitions are defined and saved for a specific table.
- Manager remembers the currently applied filter when you exit a page and applies that filter on redisplay of the table.
- Only you can see filters that you have created and only on the computer on which you created them.

Creating a New Table Filter

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

To create a new filter

1. Above any table in Manager, choose <Edit Filters...> from the Filter drop-down menu. The Advanced Filter dialog is displayed:



2. Click Add Filter above the list of existing filters.
3. In the Name field, type a unique name for your filter.
4. Specify a rule for your filter using these fields:

Field	Value to specify
Filter By	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	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	the value or condition the data in the column must match.

5. To specify another rule, click Add Rule.
6. To remove any rule, click Remove Rule to the left of the rule definition.
7. Select 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.

Copying an Existing Filter

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

To copy an existing filter

1. Above any table in Manager, choose <Edit Filters...> from the Filter drop-down menu.
2. Select the filter in the list box on the left.
3. Click Copy Filter above the list of filters.
4. Edit the name of the filter.
5. Make changes to the filter as you want.
6. Click OK to save the filter.

Manager UI Reference

- [MANAGER HOME Page](#)
- [SERVER OVERVIEW Page](#)
- [CACHED RESOURCES Page](#)
- [DATA SOURCES Page](#)
- [REQUESTS page](#)
- [SESSIONS page](#)
- [TRANSACTIONS Page](#)
- [TRIGGERS page](#)

MANAGER HOME Page

The Manager HOME page is the first page displayed when you log in to Manager. MANAGER HOME provides a quick summary of the current TDV status. The MANAGER HOME page is shown in [Launching Manager](#).

This topic includes:

- [SERVER INFO Panel](#)
- [SERVER STATUS Panel](#)
- [QUICK LINKS Panel](#)

SERVER INFO Panel

The SERVER INFO panel shows summary information and links to the SERVER OVERVIEW page.

- Server Name - the HTTP base port is displayed. All other TDV ports are derived from the HTTP base port as follows:
 - base port +1 = JDBC, ODBC, and ADO.NET
 - base port +2 = HTTP SSL
 - base port +3 = JDBC SSL, ODBC SSL, and ADO.NET SSL
 - base port +4 = Reserved
 - base port +5 = Reserved
 - base port +6 = TDV Process Monitor
 - base port +7 = Active Cluster / JGroup
 - base port +8 = Default for Repository
 - base port +9 = Monitor

In Studio, you can view and change the HTTP base port setting when required on the Configuration panel at: *Port*

Changing the base port changes all other ports on server restart, so system impact must be carefully considered before changes are made.

- Total Memory Used

Percentage of total available Java Heap Memory (RAM) currently in use. The value is controlled by TDV configuration parameters that requires a TDV Server 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

The SERVER STATUS panel displays 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), where # is the number of warnings or failures for the module listed. 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 these pages:

- EVENT LOG (see [EVENT LOG Summary Information](#))
- CLUSTER MANAGEMENT (see “Working with Active Cluster” in *TDV Active Cluster Guide*)
- USER MANAGEMENT (see [User Management](#))

Switching the Language of Web Manager UI

After logging into Web Manager, you can change the language of your user interface. To do that,

1. Choose the language of your choice in the Language drop down found in the top right corner of the Manager HOME page.
2. You will be prompted to login again. Click OK.
3. Once you log back in, the UI will be displayed in the language of your choice.

SERVER OVERVIEW Page

In Manager, you access TDV server overview information by choosing Server Overview from the MONITORING menu. The SERVER OVERVIEW page is displayed, providing a consolidated overview of the system, and the overall status of all of the other system components.

The SERVER OVERVIEW page displays 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.

- [Server Status Information](#)
- [Session and Request Information](#)
- [Privilege, User, and Repository Caches](#)
- [Server Status Indicators](#)
- [Work with the SERVER OVERVIEW Page](#)

Server Status Information

The upper left section of the SERVER OVERVIEW page displays server status information. Each item listed is described in this section. Where appropriate, the related configuration parameter in Studio is also described.

- **Status**
Status reports the presence and count of errors and warnings from all pages under the MONITORING tab.
In Studio Manager, more status information is available: Server is running, Server is stopping, Server stopped, Server is starting, Server failed, or Unknown server status (which is shown if the Monitor is not running.) See [Server Overview Panel](#) for more information.
- **Server Name**
Server Name with the HTTP base port is displayed.
- **Total Memory Used**
Percentage of total available Java Heap Memory (RAM) currently in use. Java Heap Memory is a TDV configuration setting that requires restart to change. 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 can be viewed from the Events console in the Manager.

In Studio, you can set this number using the Configuration dialog setting for Maximum Viewable Entries under Memory:

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

- **Maximum Event Entries**

Maximum number of events to be stored in the TDV repository. When the number of events reaches this threshold the oldest events are discarded in FIFO (first in, first out) order. The log files are generally configured to retain a more expansive archive of event entries.

In Studio, you can change the Maximum Event Entries using the Configuration window setting for Maximum Log Entries in Database Logger:

Session and Request Information

The SERVER OVERVIEW page also displays summary information about sessions and requests.

- **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 metadata number of resources loaded on the server. includes all requests completed or otherwise.

Privilege, User, and Repository Caches

Summary cache information on the SERVER OVERVIEW page displays usage of TDV 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.

For more information on caching, see the *TDV User Guide*.

The system caches store data such as user session values of privileges, recently introspected data source metadata, and execution plans.

Access (Hits/ Accesses), the first column in the sub-section, 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.

A high percentage of hits to total access attempts is one indicator of enhanced performance. It would indicate 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. In Studio, you can change privilege cache capacity setting in the Configuration window at Privilege Cache Size (On Server Restart).

- User Cache

Current user cache data stored in the repository. In Studio, you can change the user cache capacity setting in the Configuration window at User Cache Size (On Server Restart).

- Repository Cache

Repository cache is a resource metadata store enabling quick use of configured resources. In Studio, you can change the repository cache capacity setting in the Configuration window at Metadata Cache Size (On Server Restart).

Server Status Indicators

The SERVER STATUS summary box at right of the SERVER OVERVIEW page displays a summary status for other Manager pages, which are all available for display from the MONITORING and LOGGING tabs. The red, green, and yellow indicate the status of key system components. You can click on any of the links to go to the related page for detailed information.

Work with the SERVER OVERVIEW Page

You can start and stop a server or clear the repository cache for the server by clicking these buttons:

- **Stop**
Stops TDV after acknowledgment of a verification prompt. Actually stopping TDV requires the Modify All Status right. The button is visible, but grayed out and inactive for other users who can have access to the Manager.
- **Clear Repository Cache**
Immediately empties the repository cache. Clearing the repository cache requires the Modify All Status right. The button is visible, but grayed out for any other users who can have access to the Manager.

CACHED RESOURCES Page

In Manager, you access cache resources information by choosing Cached Resources from the MONITORING menu. The CACHED RESOURCES page is displayed, providing information about the cached views and procedures, both enabled and disabled, that are configured for use in TDV. Summary information is displayed at the top of the CACHED RESOURCES page, and information about each individual cache is displayed in the table.

- [Work with the CACHED RESOURCES Page](#)
- [The CACHED RESOURCES Table](#)
- [The IN PROGRESS REFRESHES Table](#)
- [Cached Resource Details](#)

Work 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 caches, you can click these buttons:

- **Change Enabling**—toggles the enabled/disabled status of selected caches. Changing of the Enabled status requires user to have the Modify All Status right.
- **Refresh Cache(s)**—refreshes the selected caches. Cache refresh requires that the user have the Write privilege on the selected resource.
- **Clear Cache(s)**—refreshes the selected caches. Cache refresh requires that the user have the Write privilege on the selected resource.

The CACHED RESOURCES Table

The CACHED RESOURCES table displays summary information with cache details available.

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:

Status	Event
NOT LOADED	Cache is configured, but it has not been loaded.
UP	Cache has been loaded successfully.
FAILED	Cache is not loaded and the most recent refresh failed.
STALE	Cache is loaded with valid data, but the most recent refresh failed. Reads against the cache can succeed. If the status is STALE with an "All buckets are in use ..." message, you can clear the cache to get out of that status.
DISABLED	Cache has been disabled.
CONFIG	Cache cannot operate due to a configuration error.

Status	Event
ERROR	
SETTINGS MISMATCH	This indicates a difference between the data source and TDV for the case sensitivity or trailing spaces settings. It is informational only and can be ignored. It indicates that the cache is available.

Type - Values can 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 metadata.

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.

Message - More details about the cache cancel operations (For example, Cancelled, Reason for failing, etc.)

The IN PROGRESS REFRESHES Table

The IN PROGRESS REFRESHES table displays any cache resources that are in the process of being refreshed.

Cached Resource Details

Each row also has a Show Row Details button which you can use to display the fully qualified path and other details about cached resources.

Path - TDV 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 TDV 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 TDV restart.

Total Failures - Count of failed refreshes since last TDV restart.

Message - Error message returned from cache refresh failure.

DATA SOURCES Page

The Manager DATA SOURCES page provides information about all data sources added to the TDV Server, including:

- 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 TDV since the last restart.
- Estimations of the total volume of data passed from TDV to all data sources and back to TDV since the last TDV restart.

This section includes:

- [DATA SOURCES Summary Information](#)
- [Work with the DATA SOURCES Page](#)
- [The DATA SOURCES Table](#)

DATA SOURCES Summary Information

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), where # is the number of warnings or failures for the module listed. 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.

Work with the DATA SOURCES Page

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

- Enable or disable the data source - the Change Enabling button toggles the status of the data source. Enabling makes the data source accessible through TDV definitions and configurations. Disabled takes the data source offline and makes it inaccessible to TDV 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 Administrative user with the Modify All Status right can use the Test All button.

The DATA SOURCES Table

The following columns are shown for each data source in TDV.

Name - User-defined name for the datasource.

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 - Overall status for the data source, which can be one of the following:

Status	Indicates
DISABLED	The data source is disabled; represented by a gray 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 TDV Server; represented by a green circle.

Type - Native data source type, some of the more common supported data sources categories include: DB2, Composite, Custom Java Procedure, Infirmity, FileCache, LDAP, Microsoft Access, Microsoft Excel, Microsoft SQL Server, PostgreSQL, Netezza, Oracle, Sybase, Teradata, WSDL, XML, and XmlHttp

Category - Can be File, LDAP, Relational, WSDL, XML/HTTP.

Owner - Resource owner.

Active Requests - Current count of all outstanding data source requests

Total Requests - Cumulative count of all requests made on data sources (through TDV) since start of the TDV.

Pool Size (In Use) - Current connection-pool size for relational data sources

Allocated Pool Size - Current number of connections allocated for TDV 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 - Usage of pool represented as a percentage where allocated connections is divided maximum connections for a relational data source.

Data Source 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 - Can be File, LDAP, Relational, WSDL, XML/HTTP.

Type - Kind of 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 TDV 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 - Usage 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 page

The Manager REQUESTS page provides information about all current requests for service including:

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

Summary information is displayed at the top of the REQUESTS page, and information about individual request is displayed in the table. 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.

This section includes:

- [REQUESTS Summary Information](#)
- [Work with the REQUESTS Page](#)
- [The REQUESTS Table](#)

Some of the information displayed on the REQUESTS page is controlled by the following Studio configuration parameters:

- Request Events
- Requests

Requests are removed from the table periodically, based on the Studio Manager configuration setting: Request Purge Period

The default setting purges requests every 5 minutes.

REQUESTS Summary Information

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 can be used for notification.
- **Server Requests (Active and Total)** - The number of currently active requests, and the total number of requests made to the server since the server was started.
- **Data Source Requests (Active and Total)** - The number of currently active requests, and the total number of requests made to the data sources since the server was started.

Work with the REQUESTS Page

Select one or more requests in the REQUESTS table with the respective check boxes and then perform the following actions on those requests:

- **Clear Plan Caches button** - Clears all query plan caches, in the event that current statistics gathering can 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 button** - Immediately removes all completed and failed requests. This is useful to reset the view prior to testing a set of requests.
- **Cancel Requests button** - Clears all requests.

The REQUESTS Table

The REQUESTS table displays these columns for each request:

ID - Unique request identifier.

Status - Can be any one of the following values:

Status	Indicates that the request...
STARTED	Was created but not invoked or executed.
READY	Data is ready for a client, but no client has read the data.
RUNNING	Is currently executing. If a client were blocked on the request, then the status would be running.
WAITING	Exists in a wait queue.
COMPLETED	Execution successfully completed but is not yet closed.
CLOSING	Is in the process of closing.
SUCCESS	Successfully executed and has closed.
FAILED	Execution failed.
TERMINATED	Was closed by canceling.
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 in Studio using this property: <i>Number of Top Requests Tracked</i> 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 - User ID of the user who submitted the request.

Parent ID - Unique ID for the requests parent process.

Session ID - STUDIO, HTTP (Web service), INTERNAL, or client procedure.

Session Name - Name of the component that initiated this request.

Start Time - Date and time the request started to execute.

End Time - Date and 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 Used Memory - Maximum memory used 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. To view the read-only details, click the Show Row Details button for the row.

In addition to the information presented in the REQUESTS table (and described in [The REQUESTS Table](#)), 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 requests 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 page

In Manager, you access sessions information by choosing Sessions from the MONITORING menu. The SESSIONS page is displayed, providing 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.

Some of the information displayed on the SESSIONS page is controlled by the following Studio configuration parameters:

- *Session Events*
- *Sessions*

This section includes:

- [SESSIONS Summary Information](#)
- [The SESSIONS Table](#)
- [Working with the SESSIONS Page](#)

Studio Manager provides a Sessions panel. See [Sessions Panel](#) for more information.

Sessions are removed from the table periodically, based on the Studio Manager configuration setting: Session Purge Period

The default setting purges sessions every 30 minutes. You can also use the Purge Completed Sessions button to immediately remove all sessions. This is useful to reset the view prior to testing a set of sessions.

SESSIONS Summary Information

The SUMMARY page provides the following summary 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 (Active and Total) - 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 by clicking 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

The SESSIONS table displays these columns for each session:

ID - Unique session identifier.

Status - Can be any one of the following values:

Status	Indicates that the session...
ACTIVE	The session is currently active.
TIME_OUT	The session has timed out. The Session Timeout configuration setting in Studio determines how long the session can remain idle before it times out. See TDV 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 available. To view the read-only details, click the Show Row Details button for the row.

In addition to the information presented in the SESSIONS table ([The SESSIONS Table](#)), 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 Page

In Manager, you access sessions information by choosing Transactions from the MONITORING menu. The TRANSACTIONS page is displayed, providing information about the current and recently active transactions. Summary information and information about each individual transaction is displayed.

This section includes:

- [TRANSACTIONS Summary Information](#)
- [Work with the TRANSACTIONS Page](#)
- [The TRANSACTIONS Table](#)

Studio Manager provides a Transactions panel. See [Transactions Panel](#) for more information.

You might need to scroll the display to the right to see all of the information provided.

By default, the TRANSACTIONS page displays transactions that occurred within the last 5 minutes. Transactions created by Manager itself are not displayed.

Some of the information displayed on the TRANSACTIONS page is controlled by the following Studio configuration parameters:

- Transaction Events
- Transactions

Transactions are removed from the table periodically, based on the Studio Manager configuration setting: Transaction Purge Period

The default setting purges transactions every 5 minutes.

TRANSACTIONS Summary Information

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.

Work with the TRANSACTIONS Page

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

The TRANSACTIONS table displays these columns for each transaction:

ID - Unique transaction identifier.

Status - Can be any one of the following values:

Status	Indicates that the transaction...
START	Was started.
COMMITTED	Has been committed.
FAIL	Failed.
ROLLBACK	Was rolled back.
COMPENSATE	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. For example, Studio.

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 transaction has additional detailed information available. To view the read-only details, click the Show Row Details button for the row.

In addition to the information presented in the TRANSACTIONS table (and described in [The TRANSACTIONS Table](#)), these details are provided:

- Owner domain - Name of the domain to which the session owner belongs.
- Session Type - STUDIO, HTTP (Web service), INTERNAL, or client procedure.
- Active Requests - The total number of currently active transactions.
- Total Requests - The total number of transactions since the server started.

TRIGGERS page

In Manager, you access trigger information by choosing Triggers from the MONITORING menu. The TRIGGERS page is displayed, providing information about the current and recently active triggers. Summary information and information about each trigger is displayed.

This section includes:

- [TRIGGER Summary Information](#)
- [Work with the TRIGGERS Page](#)
- [The TRIGGERS Table](#)

Studio Manager provides a Triggers panel. See [Triggers Panel](#) for more information.

By default, the Manager TRIGGERS page displays triggers that occurred within the last 5 minutes. Triggers created by Manager itself are not displayed.

The TestAllDataSources trigger is provided by default. Configure the TestAllDataSources trigger from the Schedule Test button on the Studio Manager DATA SOURCES page.

Some of the information displayed on the TRIGGERS page is controlled by the following Studio configuration parameters:

- *Trigger Events*
- *Triggers*

For more information about triggers, see Triggers in the *TDV User Guide*.

TRIGGER Summary Information

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.

Work with the TRIGGERS Page

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 TRIGGERS table displays these columns for each trigger:

Name - The name of the trigger.

Status - Can be any one of the following values:

Status	Indicates that the trigger...
ACTIVE	Is currently processing.
DISABLED	Is disabled.
CONFIG ERROR	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.

Trigger Details

Each trigger has additional detailed information available. To view the read-only details, click the Show Row Details button for the row.

In addition to the information presented in the TRIGGERS table (and described in [The TRIGGERS Table](#)), these details are provided:

- Path - Fully qualified path to the resource that has been scheduled for execution.

- Parent Type - Type of parent for this trigger.
- Owner Domain - The domain to which the Owner belongs.
- Initial Time - The first time the trigger execution occurred.
- Last Fail - Last time the trigger execution failed.
- Total Successes - Total number of times the trigger execution successfully occurred since the server started.
- Total Failures - Total number of times the trigger execution failed since the server started.
- Message - Displays an error message if the trigger caused an error.

TDV Configuration Options

Studio provides many configuration parameters that allow modification of TDV values and behavior. Use of the Studio Configuration parameter window to set these parameters is described in the *TDV User Guide*. This topic describes some of the administration-specific parameters.

- [Fine Tuning Memory](#)
- [Fine Tuning Performance Using Connection Pools](#)
- [Using Pass-through Optimization with Oracle Data Source Clients](#)
- [Using Pass-through Optimization with SQL Server Data Source Clients](#)
- [Using Pass-Through Introspection with Vertica Data Source Clients](#)
- [Enabling Studio Locking](#)
- [Configuring Case Sensitivity and Trailing Spaces Settings](#)
- [Function Overrides](#)
- [Management of Data Source Customization](#)
- [Password Storage Options](#)
- [Customizing the Login Screen Default Domain Value](#)

Fine Tuning Memory

You can fine tune caching and memory settings after taking into account factors including how paging works and the memory usage statistics in Studio Manager.

- [About Paging](#)
- [Configuring the Caching and Data Processing Directory](#)
- [Viewing Usage and Cleaning Up Memory](#)
- [Changing Default Memory Settings](#)

About 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 to keep the applications running when memory limits are exceeded. When a paged-out memory location is accessed, the operating system restores that area of memory from disk and then, to make room, moves some other part of memory to disk. Consequently, performance can suffer. Some amount of paging is fine on the client side. But on the server, where many queries must be performed simultaneously, paging must be minimized.

Procedures that are known to require large amounts of memory and exceed the allotted managed memory should be cached and scheduled for refresh at non-peak processing hours (provided the business requirement allows for that). That query should be forced to process on disk to allow other queries to run and execute efficiently for more optimized performance. See the *TDV Reference Guide* for the use of the `FORCE_DISK` query engine option.

Configuring the Caching and Data Processing Directory

The disk space must be adequately sized to allow for processing. It is recommended that the Temp directory used by TDV Server be at least 10 GB in size. The default Temp directory is created on the same partition as the TDV Server installation directory and it expands as needed up to the specified limit.

To change the location used for caching and data processing

1. Log into Studio as the admin user.
2. From the Administration menu, choose Configuration.
3. In the tree pane, navigate to TDV Server > Configuration > Files.
4. Modify the Temp Directory (On Server Restart) parameter.
5. Click Apply.
6. Click OK.
7. Restart the TDV Server.

Viewing Usage and Cleaning Up Memory

Studio displays real-time memory usage in the lower right corner of the Studio interface. The Studio Manager displays a more detailed view of the TDV Server memory usage and it also has a Free Unused Memory button that can invoke Garbage Collection.

The Resource Monitor displays how much Kernel Memory is being used and the degree to which Page Filing is occurring. If that number is bigger than the amount of physical memory, it means that your server is paging and performance can suffer for it. Investigate the queries being run and adjust your server memory settings accordingly.

If memory usage is high, you can release memory using the garbage collection feature. Typically, garbage collection is an automatic process.

To view usage and clean up unused memory

1. Open Studio.
2. Select the Manager tab on the left side of the screen.
3. Select Memory.
4. Determine what might be impacting your memory usage.
5. Click Free Unused Memory.

Changing Default Memory Settings

You must balance the following to determine optimal memory configuration:

- Queries run faster with more memory. So, giving the server as much memory as possible is highly recommended.
- Giving the server too much memory can cause excessive paging and significantly degrade performance.

To change the default memory setting

1. Log into Studio as the admin user.
2. From the Administration menu, choose Configuration.
3. In the tree pane, navigate to TDV Server > Memory.
4. You can modify the parameter settings for the following categories:

- Managed Memory
 - Sampling
5. You can also locate and modify the following configuration parameters to optimize memory settings for your system:

Configuration Parameter	Description
Total System Memory	The total physical memory calculated on the target installation machine.
Upper Limit of Total Available Memory	The upper limit for the total amount of Java heap (memory) available for the server's use. The default value is the total physical memory calculated on the target installation machine.
Piped Buffer Size	This is the Piped Stream buffer size for SOAP or REST message encoder. The default value is 1024 and the maximum value is 2048. It is important to note that setting a very high value for Piped Buffer Size will get allocated from the Java Heap Space which means TDV will reduce the memory from total available memory.

6. Click Apply.
7. Click OK.
8. Restart the TDV Server to implement your changes.

Fine Tuning Performance Using Connection Pools

A connection pool is a cache of database connections for reuse when future requests to the database need them. These open and available connections make sending requests to the database faster. The connection pool can reduce or eliminate the need to open a new connection for each request.

To fine tune connection pool performance

1. Log into Studio as the admin user.
2. From the Administration menu, choose Configuration.
3. In the tree pane, navigate to Configuration > Data Sources > Common to Multiple Source Types.
4. You can manipulate the following parameters to tune connection pool performance. You might need to iterate through several values until you find the values that provide the best performance.

Configuration Parameter	Descriptions
Default Execution Timeout	Sets the number of seconds that the data source waits for an execution to occur. If the limit is exceeded, a timeout occurs. If this value is zero, it means there is no limit. If it is greater than 0, the execution waits at most the given number of seconds.
Delayed Connection Commit/Rollback Timeout	When the TDV Server is asked to rollback or commit a JDBC connection on a data source, it waits until all the result sets are closed. This timeout controls how long the server waits. After the timeout, if there are still open result sets, the connection is forcibly closed. If the timeout value is 0, the server does not timeout. The default is 300 seconds.

5. Click Apply.
6. Click OK.
7. Restart the TDV Server to implement your changes.

Using Pass-through Optimization with Oracle Data Source Clients

Pass-through optimization affects DATE and TIMESTAMP data types in Oracle data sources. The default value of Pass Thru Optimizations is true. To take advantage of this optimization for 6.1.0.1 JDBC clients, leave it set to true. Set it to false for pre-6.1.0.1 JDBC, ODBC and ADO.NET clients.

To set pass-through optimization for Oracle data source clients

1. Log into Studio as the admin user.
2. From the Administration menu, choose Configuration.
3. In the tree pane, navigate to Configuration > Data Sources > Oracle Sources > Performance.
4. If you are using a 6.1.0.1 JDBC client for this data source, leave Enable Pass Thru Optimizations set to true; otherwise, set it to false.
5. Click Apply.
6. Click OK.
7. Restart the TDV Server to implement any changes you have made.

Using Pass-through Optimization with SQL Server Data Source Clients

Pass-through optimization affects DATE, DATETIME, DATETIME2, MONEY, SMALLDATETIME, SMALLMONEY, and TIME data types in Microsoft SQL Server data sources. The default value of Pass Thru Optimizations is true. To take advantage of this optimization for 6.1.0.1 JDBC clients, leave it set to true. For pre-6.1.0.1 JDBC, ODBC and ADO.NET clients, set it to false.

To set pass-through optimization for SQL Server data source clients

1. Log into Studio as the admin user.
2. From the Administration menu, choose Configuration.
3. In the tree pane, navigate to Configuration > Data Sources > MS SQLServer Sources > Performance.
4. If you are using a 6.1.0.1 JDBC client for this data source, leave Enable Pass Thru Optimizations set to true; otherwise, set it to false.
5. Click Apply.
6. Click OK.
7. Restart the TDV Server to implement any changes you have made.

Using Pass-Through Introspection with Vertica Data Source Clients

When Introspect as Pass-Through Following DDL is enabled on a Vertica data source, tables created in the data source using the DDL feature are created using the same pass-through user that submitted the DDL. Enabling this will make introspection ignore the saved credential if the DDL was issued by a pass-through user.

To set pass-through optimization for Vertica data source clients

1. Log into Studio as the admin user.
2. From the Administration menu, choose Configuration.
3. In the tree pane, navigate to Configuration > Data Sources > Vertica Sources.
4. Select true, to enable introspection as a pass-through user.
5. Click Apply.
6. Click OK.
7. Restart the TDV Server to implement any changes you have made.

Enabling Studio Locking

When the Studio locking is enabled, Studio forces users to acquire a lock prior to changing a resource. However, the TDV Server Web services API does not honor this Studio Configuration setting.

The requirement for Studio resource locking prior to changing and saving resources can be disabled for the entire server by any administrator with the Modify All Resources right.

Existing locks persist regardless of whether Studio requires locks prior to modification. When Studio locking is disabled, users can still optionally use resource locks to prevent simultaneous, conflicting changes to a resource.

To enable locking

1. Make sure you have both Modify All Config and Access Tools rights.

2. Log into Studio as the admin user.
3. From the Administration menu, choose Configuration.
4. In the tree pane, navigate to Studio > Locking > Enabled.
5. For Enabled, select True.
Disable Studio Locking by setting Enabled to False.
6. Click Apply.
7. Click OK.

This Studio configuration change is not immediately propagated to other open instances of Studio connected with this server, but any attempt to save resources forces a check to see whether the Studio Lock is enabled.

8. Restart the TDV Server.

Configuring Case Sensitivity and Trailing Spaces Settings

Case sensitivity and trailing space mismatches are common when working with different data sources. By default TDV is set to be case insensitive and to ignore trailing spaces. Changing the policy might change query results or performance. This section describes configuration tasks you can perform to control TDV behavior for case sensitivity and trailing spaces.

For example, the test ('ABC' = 'abc') returns False for a case-sensitive comparison and True for a non-case-sensitive comparison. The test ('ABC ' = 'ABC') returns False when trailing spaces are considered and True when trailing spaces are ignored.

With TDV, case sensitivity and trailing spaces mismatches occur only under the following conditions:

- A mismatch between TDV and the underlying data source's case sensitivity or trailing spaces settings.
- A WHERE clause that contains a CHAR or VARCHAR.

The following topics are covered in this section:

- [Determine Whether Case or Trailing Space Settings Affect Query Performance](#)

- [Setting Server-wide Case and Trailing Space Behavior Using Configuration Parameters](#)
- [Setting Session-wide Case and Trailing Space Behavior Using Connection Properties](#)
- [Configuring Case and Trailing Space Behavior for Built-in Procedures](#)
- [Configuring Case and Trailing Space Behavior for Queries](#)
- [Mismatch Effects on String Comparisons](#)

Determine Whether Case or Trailing Space Settings Affect Query Performance

To determine if the TDV settings are affecting query performance, you can evaluate any filter nodes or the SQL underlying each FETCH node in the execution plan in Studio. Focus primarily on the WHERE clause or filter nodes.

For example, under certain conditions and with certain configuration parameter settings, TDV might apply RTRIM or UPPER functions to string comparisons in a WHERE clause. But this prevents the underlying system from using an index on that column, which affects query latency.

As another example, when a filter is applied at the TDV level, all rows must be returned from the underlying table, which could impact performance for large tables.

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

TDV Setting	Data Source Setting	TDV Query Behavior
case_sensitivity=true	case_sensitivity=true	No special action.
case_sensitivity=true	case_sensitivity=false	Pushes WHERE clause string comparison to data source, but applies the case sensitivity filter to the results returned.
case_	case_sensitivity=true	Adds UPPER to both sides; mismatch

TDV Setting	Data Source Setting	TDV Query Behavior
sensitivity=false		is not necessarily a performance issue.
case_ sensitivity=false	case_sensitivity=false	No special action.
ignore_trailing_ spaces=true	ignore_trailing_ spaces=true	No special action.
ignore_trailing_ spaces=true	ignore_trailing_ spaces=false	Adds RTRIM to both sides; mismatch is not necessarily a performance issue.
ignore_trailing_ spaces=false	ignore_trailing_ spaces=true	Pushes WHERE clause string comparison to data source, but applies the trailing spaces filter to the results returned.
ignore_trailing_ spaces=false	ignore_trailing_ spaces=false	No special action.

TDV reports settings matches and mismatches in the Resource Capabilities section display in the opened data source configuration window. Similar reports appear on the reintrospection and cache configuration displays.

Whenever possible, set TDV case and trailing space behavior to match the data sources. If this is not possible, the following effects might occur.

Case Sensitivity	Trailing Spaces
If TDV Server is set to be case sensitive and the data source is not case sensitive, the query is pushed to the data source, and then the case sensitivity filter is applied to the result set returned.	If TDV Server is not set to ignore trailing spaces and the data source <i>is</i> set to ignore them, the query is pushed to the data source, and then TDV Server applies the trailing spaces filter to the result set returned.
If you do not want the filter to be applied to the returned results, set Push Even if Case Sensitivity Mismatch to True.	If you do not want the filter to be applied to the returned results, set Push Even if Trailing Spaces Mismatch to True.

Case Sensitivity	Trailing Spaces
If TDV Server is <i>not</i> set to be case sensitive and the database <i>is</i> case sensitive, TDV Server adds an UPPER function to both values to ensure the data source performs a non-case-sensitive comparison. This should have little impact on performance. To keep the UPPER function from being added to such queries, set Disable Case Sensitivity Correction to True.	If TDV Server is set to ignore trailing spaces and the database is not, you can force comparisons to show a match (so they can be pushed) by wrapping values in TRIM or RTRIM functions. You can have TDV push unchanged comparison syntax to the data source by setting Push Even if Trailing Spaces Mismatch to True.

Setting Server-wide Case and Trailing Space Behavior Using Configuration Parameters

TDV uses the configuration parameter values for case sensitivity and trailing spaces that have been set in the Studio Configuration window. This is useful if the data sources have consistent case and trailing space behavior among themselves. However, if you change configuration parameter settings, TDV also needs to re-evaluate other query plans against the new settings.

Note: Configuration parameters settings are server-wide. They are overridden at the session level by client interface property settings. Parameter and property settings are in turn overridden by query options.

To configure case and trailing space behavior using configuration parameters

1. Log into Studio as the admin user.
2. From the Administration menu, choose Configuration.
3. In the tree pane, navigate to TDV Server > SQL Engine > SQL Language.
4. Determine the best settings for the following parameters.

Parameter	Description
Case Sensitivity	Controls the default case sensitivity of queries. The default value of False ignores case in string comparisons. If this setting does not match the case-sensitivity setting of a data source, performance is degraded when querying that source. Changing this has no effect on currently running queries.
Ignore Trailing Spaces	Controls whether to ignore trailing spaces during string comparisons in queries. The default value of True ignores trailing spaces during string comparisons. If this setting does not match the trailing spaces setting of a data source, performance is degraded when querying that source. Changing this has no effect on currently running queries.

5. In the tree pane, navigate to TDV Server > SQL Engine > Overrides.

6. Determine the best settings for the following parameters.

Parameter	Description
Disable Case Sensitivity Correction	Determines whether the server uses UPPER functions to normalize the SQL when case sensitivity settings do not match. The default value is False.
Disable Ignore Trailing Spaces Correction	Determines whether the server uses TRIM functions to normalize the SQL when “ignore trailing spaces” settings do not match. The default value is False.
Character Functions Conform to Case Sensitivity Setting	Determines whether the server execution of character functions conforms to the case sensitivity setting of the query execution environment. If not, functions TRIM, BTRIM, LTRIM, RTRIM, INSTR, POSITION, STRPOS and TRANSLATE are case sensitive when executed within TDV. The default value is True.
Push Even if Case Sensitivity Mismatch	Determines whether the server ignores case sensitivity setting differences between the server and the data source. The default value is False.

Parameter	Description
Push Even if Trailing Spaces Mismatch	Determines whether the server ignores trailing space setting differences between the server and the data source. The default value is False.

7. Click Apply.
8. Click OK.
9. Restart the TDV Server.

Setting Session-wide Case and Trailing Space Behavior Using Connection Properties

You can set up session-wide case sensitivity and treatment of trailing spaces using the following two JDBC- and ODBC-driver connection URL properties:

- `caseSensitive`
- `ignoreTrailingSpaces`

For details, see the *TDV Client Interfaces Guide*.

Note: These properties override TDV configuration parameter settings.

Configuring Case and Trailing Space Behavior for Built-in Procedures

You can set up procedure-wide case sensitivity and treatment of trailing spaces using the following two built-in-procedure environment variables:

- `System.CASE_SENSITIVE_IN_COMPARISONS`
- `System.IGNORE_TRAILING_SPACES_IN_COMPARISONS`

Procedures (`getEnvironment`, `SetEnvironment`, `SetEnvironmentFromNodeValues`, `SetNodeValuesFromEnvironment`) are available for viewing and manipulating these environment variables. For details, see the *TDV Application Programming Interfaces Guide*.

Note: These variables override TDV configuration parameter settings for procedure execution.

Configuring Case and Trailing Space Behavior for Queries

Query options allow you to control the case and trailing space behavior for an individual query, but can produce unpredictable results when numerous data sources are used with varying case-sensitivity and trailing-space settings.

The query options are `CASE_SENSITIVE` and `IGNORE_TRAILING_SPACES`, which can be used with `SELECT`, `INSERT`, `UPDATE`, and `DELETE`. For descriptions, see the “TDV Query Engine Options” topic of the *TDV Reference Guide*.

Note: These query options override TDV configuration parameter settings.

For example, you submit the following SQL statement:

```
SELECT v1.balance FROM accounts v1
```

```
WHERE v1.account_name = 'bob'
```

If TDV submits this syntax to a case-sensitive database, TDV expects to get only accounts with lowercase 'bob' as the name. If TDV submits this syntax to a case-insensitive database, TDV expects to get accounts with 'bob' in any combination of uppercase and lowercase letters.

If you know the database is case sensitive and you want a case-insensitive comparison, submit:

```
SELECT v1.balance FROM accounts v1
```

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

The same is true of TDV. However, if TDV is *not* case sensitive and the underlying database *is* case sensitive, TDV adds the `UPPER` function to the SQL sent to the underlying database.

Note: Doing this invalidates any existing index and therefore requires a new table scan, which might affect performance.

To configure case and space using SQL query options

1. Determine what query is having the most effect on performance.
2. Use SQL query options to override the configuration settings for that query.

Using STRICT to Control Case and Trailing Space Behavior for Queries

If the TDV and data source settings do not match for case sensitivity or trailing spaces, use the STRICT option in any query that includes one of these operators:

- DISTINCT
- UNION
- INTERSECT
- EXCEPT

Mismatch Effects on String Comparisons

The case-sensitive and trailing-spaces policies affect string comparison. The policies do not affect the actual value of strings. Affected functions and operators include the following.

Case Mismatch	Trailing Spaces
Comparison operators in WHERE and JOIN ON: =, <, <=, >=, >, and <>	Comparison operators in WHERE and JOIN ON: =, <, <=, >=, >, and <>
REPLACE (src,pattern,escape) The pattern is matched according to the policy.	LENGTH (column) The string length returned does not include trailing spaces.

Case Mismatch	Trailing Spaces
MIN (column) The strings 'ABC' and 'abc' are considered the same, so either can be chosen by this function.	MIN (column) The strings 'abc ' and 'abc' are considered the same, so either can be chosen by this function.
MAX (column) The strings 'ABC' and 'abc' are considered the same, so either can be chosen by this function.	MAX (column) The strings 'abc ' and 'abc' are considered the same, so either can be chosen by this function.
GROUP BY The strings 'ABC' and 'abc' are considered the same, so the group include sboth sets of values.	GROUP BY The strings 'abc ' and 'abc' are considered the same, so the group includes both sets of values.
ORDER BY The strings 'ABC' and 'abc' are considered the same, so they sort together and can be intermixed.	ORDER BY The strings 'abc ' and 'abc' are considered the same, so they sort together and can be intermixed.

Function Overrides

Setting the overrides options enables users to override the default behavior of the TDV server.

LPAD and RPAD return length

The LPAD and RPAD functions truncate strings or pads them with spaces (or specified characters), to make all returned values the same specified length.

The length of the result string can be tuned by altering the configuration parameter "LPAD and RPAD return length". From Studio, select Administration > Configuration > Server > SQL Engine > Overrides > LPAD and RPAD return length.

When this override option is set to TRUE, the functions LPAD and RPAD will result in a VARCHAR(4000). When set to FALSE the result will be a VARCHAR(2147483647). Default value for this option is FALSE.

Management of Data Source Customization

Capabilities files are a mechanism for TDV to determine the behavior of data sources and their connections. TDV uses an abstraction called “data source adapter” that lets customers change data source capabilities and have their changes persist across TDV patches and upgrades.

TDV determines data source capabilities by reading configuration files for that data source in a specific order. The value of any capability in each file overrides the value of the same capability in a previously read file. An Oracle 11g thin driver data source is used as an example.

- <TDV_install_dir>\apps\dml\cis_ds_oracle\conf\oracle.capabilities—A text file that defines the basic capabilities for an Oracle data source
- <TDV_install_dir>\apps\dml\cis_ds_oracle\conf\oracle_11g_thin_driver.oracle.capabilities—A text file that defines additions and overrides specific to an Oracle 11g thin driver data source
- <TDV_install_dir>\conf\adapters\system\oracle_11g_thin_driver\oracle_11g_thin_driver_values.xml—An XML file that overrides some capabilities with values specific to the 11g thin driver version of Oracle
- <TDV_install_dir>\conf\adapters\custom\MyOracleAdapter\myoracleadapter_values.xml—An XML file, present only if a custom adapter has been built off the Oracle adapter, that overrides some capabilities with custom adapter values

The final two files, when present, are included in CAR exports so that they are later among the import files on the target server instance or version. In this way, any adapter customizations are carried over.

If an extension adapter has been created, a fifth layer of configuration files is present. Some adapters, such as Greenplum, are less complex and have fewer levels of configuration files.

Customization and capabilities for extension adapters and their data sources are handled differently from what this section describes. For details, see the *TDV Data Source Toolkit Guide*.

Password Storage Options

By default users defined in the composite domain are encrypted and stored in the TDV repository. The encryption converts the password to a 100-byte random salt that is used to create and SHA-512 hash key, which is what is needed if you require WSSE Username Token authentication for your SOAP documents.

If storing the passwords in the repository is not a concern, you can set this parameter to TRUE.

To store unencrypted passwords in the TDV repository

1. Make sure you have both Modify All Config and Access Tools rights.
2. Log into Studio as the admin user.
3. From the Administration menu, choose Configuration.
4. In the tree pane, navigate to Server > Configuration > Security > Store User Passwords.
5. For Store User Passwords in an unencrypted format, select True.
Encrypt passwords by setting to False.
6. Click Apply.
7. Click OK.

This Studio configuration change is not immediately propagated to other open instances of Studio connected with this server.

8. Restart the TDV server.

Customizing the Login Screen Default Domain Value

By default the login screen displays composite as the domain value. You can customize the default domain value by setting the value of a Studio configuration parameter.

To customize the default login screen domain

1. Make sure you have both Modify All Config and Access Tools rights.

2. Log into Studio as the admin user.
3. From the Administration menu, choose Configuration.
4. Locate the Default Domain Name parameter, type the default value that you want login screens to display when opening the product.

If you want the value to be the composite domain, no change is necessary. If you want the value to be one of your other LDAP defined domains, type the value of the LDAP domain.

5. Click Apply.
6. Click OK.

This configuration change is not immediately propagated to other open instances of Studio connected with this server.

7. Restart the TDV server.

Custom Server Settings For External Tools

The following server configuration options can be tuned to use with external tools.

Property	Description
Custom Server Properties	This configuration allows users to set arbitrary properties on the TDV server that can be used by external tools.
Custom Cluster Properties	This configuration allows users to set arbitrary properties on the TDV server that can be used by external tools.

Using the key-value pair value for the above properties, it is recommended that users construct keys for the map using a dotted notation (for example, tibco.psg.xxxx or tibco.<tool>.xxxx) to help avoid conflicts.

Composite Domain Administration

TDV supports the composite, dynamic, and LDAP domains, each of which controls a particular set of users and groups that can access TDV. This topic describes the composite domain and how to create and manage its users and groups.

- [About the Composite Domain](#)
- [About Domain Management](#)
- [Group Management](#)
- [User Management](#)
- [Auditing User Access to TDV Defined Resources](#)
- [Changing Passwords for Other Composite Domain Users](#)
- [Changing Ownership of Resources](#)
- [Manage User and Group Privileges](#)

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

- [LDAP Domain Administration](#)
- [Dynamic Domain Administration](#)

About the Composite Domain

The Composite domain includes users and groups defined within TDV to access TDV. TDV has predefined specific users and groups in the Composite domain which you can use and modify as appropriate. You can create additional users and groups within the Composite domain to meet your specific needs.

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

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

- From Studio, choose Administration > Launch Manager.


- From a Web browser (when TDV is locally installed), use this URL:
http://localhost:9400/manager
- From a Web browser (when TDV is not locally installed), use this URL:
http://[TDV_host_name]:[port_number]/manager

Note: Internet Explorer 8 has compatibility issues when viewing Manager pages. If you have difficulties displaying Manager, engage the Compatibility View option available from the IE8 Tools menu.

About Domain Management

Domain management includes adding and removing domains and the users and groups assigned to a domain. Two domains—composite and dynamic—are already defined for use when TDV is installed. The Manager DOMAIN MANAGEMENT page lists the defined domains and provides links to view the groups and users within those respective domains.

Note: Use of @ and : characters is not fully supported in domain names in TDV.



The DOMAIN MANAGEMENT page is used primarily for the specification of LDAP domains and external groups that will have rights and privileges to view and use TDV defined resources.

Domain management for LDAP domain configurations are described in [LDAP Domain Administration](#). For more information on dynamic domain administration see, [Dynamic Domain Administration](#).

Group Management

Administrators with the Access Tools and Read All Users rights can create groups of users who share the same rights to perform administrative tasks on the server, and groups who need access to TDV tools to create, view, access, and change objects defined with Studio. Developers, operations personnel, and administrators should each have their own groups to access Manager and other TDV tools and options.

Group rights templates enable quick assignment of rights based on an expected level of interaction with TDV, Studio, and other TDV tools. Group rights templates exist for: Administrators, Developers, Operations, Backup, Restore, Backup & Restore, and End Users. See [Understanding TDV User Templates and Rights](#).

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. You can 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.

Groups are also useful for assigning and managing resource privileges for sets of users. Assign resource privileges to groups so members can access and use data sources, views, and procedures. Administrators with the Modify All Resources right, resource owners, and users with a grant privilege on that resource can assign privileges to groups and users.

- [Built-in Groups](#)
- [Adding Groups to the Composite Domain](#)
- [Removing Groups](#)

Built-in Groups

There are three built-in groups that are created by the system and cannot be deleted:

- **admin (composite)**—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 (composite)**—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 Composite Domain

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

For more information on how to customize the rights as required, see [Understanding TDV User Templates and Rights](#).

To add a group to the composite domain

1. Launch Manager from Studio or direct a Web browser to the Manager using one of these URLs.

- When TDV is locally installed:

```
http://localhost:9400/manager
```

- When TDV is not locally installed:

```
http://[TDV_host_name]:[port_number]/manager
```

After login, the MANAGER HOME page is displayed.

2. From the SECURITY tab, choose Group Management.
3. Click Add Group.
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.
6. In the "Group Restrictions" section,
 - Select or deselect "Direct DS Data Access" restriction for the selected group. This access control option is given to prevent direct data access to the data source asset. The restricted group is subject to direct data source data access restriction and cannot access the data in the TDV tools.
 - Select or deselect "Direct DS Invocation" restriction for the selected group. This access control option is given to prevent stored procedure calls to the data source asset. The restricted group is subject to direct data source invocation restriction and cannot make stored procedure calls to the data source from any TDV tools.

7. 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 administration.
8. Click OK.

The group is added to the GROUP MANAGEMENT page.

Removing Groups

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

TDV users who were members of a deleted group 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.

Note: Deletion of a composite domain group does not remove its member users from the TDV.

Removing a group from the composite domain

1. In Manager, choose SECURITY > Group Management.
2. Select one or more groups using the check box.
3. Click Remove Group(s) to delete the group.
4. Accept the confirmation prompt and the group is deleted.

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 TDV system.


Removing an externally defined LDAP group

5. In Manager, choose SECURITY > Domain Management.
6. Select the LDAP domain using the radio button
7. Click Edit External Groups.

User Management

User management means much more than adding and removing users from the Composite domain. Administrators have an active role in mediating access to both data resources and to design tools that enable creation and modification of resources. From the User Management page an administrator with the Read All Users right can review all users, their rights, and their group membership. An administrator who also has the Modify All Users and Modify All Resources rights can add, edit, and modify users, change group membership, and edit resource privileges. The Edit Resource Privileges function is also useful as a user-based security audit showing all the resources for which a selected user has privileges.

Note: Use of the “:” character is not fully supported in user names in TDV.



User management is facilitated by group management of both privileges and rights. See [Group Management](#) for more information on setting up those groups.

- [Built-in Users and Their Privileges](#)
- [Adding Users to the Composite Domain](#)
- [Removing Users from the Composite Domain](#)
- [Auditing User Access to TDV Defined Resources](#)
- [Managing Group Membership](#)

- [Viewing Group Membership](#)
- [Editing Group Membership](#)

Built-in Users and Their Privileges

The composite domain has the following permanent users that are automatically created. These users cannot be removed:

- **admin**—This user has privileges to access and use any resource in the system; admin can also grant and 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 TDV 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 be removed. It owns items that even the users with administrative privileges cannot modify. The SYSTEM account is used to control TDV communication with the repository. The SYSTEM account cannot be used to login to a TDV instance.
- **Monitor**—This user is for TDV to communicate with the monitor.

The all group includes all composite users and all dynamic users, but not the user named nobody. All members of this group have READ privileges for all folders created with the installation, but not newly created folders and resources. Privileges must be assigned by the creator or owner of the resource, or by an administrator or user explicitly given the GRANT right on that object.

All semi-editable folders (for example, /shared, /services/databases, /services/Webservices) have no privileges, but they are editable.

All precreated tables and procedures have SELECT and EXECUTE privileges for the all groups in the composite and dynamic domains, and the anonymous user in the composite domain. For example:

```
/services/databases/system
```

```
/services/webservices/system
```



```
/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, and grant the READ privilege to the data service, service, and port that you want the anonymous user to be able to access and use. The global option `anonymousOptionsRequest` controls whether to allow an HTTP anonymous login request even if the server configuration for anonymous login is disabled.

Anonymous users cannot connect to the server using JDBC, because no TDV data service of the type database is automatically available. To enable them to connect, grant READ privileges to services/databases, the data service, and any catalogs or schemas that you want to make available.

Resources in the Data Services area point to resources in the work area. To access a resource in the Data Services area, the anonymous user needs permission to read all the folders above that item, and have appropriate permission on the item to which the resource points.

To expose a resource to Web services or JDBC clients, 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, repeat the process with those resources as well.

This is similar to what you would do for any other user, except that for those folders that have the READ privilege by default for the all group, you might need to override privileges on those folders.

The anonymous user is denied access to the /users folder; admin cannot change this. All published resources you want anonymous to be able to use must reside in the /shared folder.

Adding Users to the Composite Domain

TDV administrators with the Modify All Users and Modify All Resources rights can add users to the composite domain.

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

To add a user to the composite domain

1. Launch Manager from Studio or direct a Web browser to the Manager using one of these URLs.

- When TDV is locally installed:

```
http://localhost:9400/manager
```

- When TDV is not locally installed:

```
http://[TDV_host_name]:[port_number]/manager
```

After login, the MANAGER HOME page is displayed.

2. From the SECURITY tab, choose User Management.
3. Click Add User.
4. Enter the new user name and password with a confirmation entry for the password.
New passwords must be between six and 64 characters long. The password can contain numeric and uppercase alphabetic characters, and selected symbols.
5. Select a base template to begin rights assignment.
6. Select or deselect rights as appropriate for the local security policy and the expected level of user interaction with the TDV Server and underlying data sources.

7. User Restrictions

In the "User Restrictions" section,

- Select or deselect "Direct DS Data Access" restriction for the selected user. This access control option is given to prevent direct data access to the data source asset. The restricted user is subject to direct data source data access restriction and cannot access the data in the TDV tools.
 - Select or deselect "Direct DS Invocation" restriction for the selected user. This access control option is given to prevent stored procedure calls to the data source asset. The restricted user is subject to direct data source invocation restriction and cannot make stored procedure calls to the data source from any TDV tools.
8. Enter comments in the Annotation field to give administrators an indication of the user's role in the system or organization.

9. Click OK.

The newly added user name is added to the composite domain.

Removing Users from the Composite Domain

Removing a user from the composite domain removes the user from TDV.

Removing a user who is derived from an LDAP domain/group does not prohibit the user from logging into the system again. See [Remove LDAP Users from TDV](#) for more information.

To remove a user from the composite domain

1. Launch Manager from Studio or direct a Web browser to the Manager using one of these URLs:
 - When TDV is locally installed:

```
http://localhost:9400/manager
```

- When TDV is not locally installed:

```
http://[TDV_host_name]:[port_number]/manager
```

After login, the MANAGER HOME page is displayed.

2. From the SECURITY tab, choose User Management.
3. Select check box to the left of the each user to be removed from the domain
4. Click Remove User(s).
5. Confirm that you want to delete the users and the user or users are removed from TDV Server.

Auditing User Access to TDV Defined Resources

Administrators with the Read All Users and Modify All Resources rights can use the User Management page in Manager to get a single page view of all resource privileges held by any selected user or set of users.


Whether the privilege is explicit or implicit does not matter to the actual functionality provided to the user or group of users with that privilege.

To view all the resource privileges held by a user or a set of users

1. Open the User Management page from the Users tab selection in Manager.
2. Select a user or a group of users.
3. Click Edit Resource Privileges to display a list of all resource for which the selected user has more than one privilege.

All privileges assigned to the user (or selected group of users) for any TDV defined resource are displayed in the Edit Resource Privileges window.

The Edit Resource Privileges window allows for direct modification of user privileges. An administrator with the Modify All Resources right can add or remove explicitly assigned privileges for the user or set of users selected.



When more than one user is selected, different privilege settings for the same resource privilege will be represented by one of the following:

- Users all have an explicitly assigned privilege which can be removed directly.
- Some users have an explicitly assigned privilege while others do not have the privilege. Clicking this icon once will add an explicit privilege for all users/groups selected. Clicking it again removes all explicitly assigned privileges, and clicking the box a third time leaves the assignments unchanged.
- Users have an implicitly derived privilege from either group membership or admin right. Implicit privileges can only be removed by either removing the

user from the group (or groups) that grant that privilege, or by removal of the administrative right that grants that privilege.

- Some users have an implicitly derived privilege while at least one other does not. Clicking this box once will assign explicit privileges to all users.
 - All users have both an explicit and implicit privilege. This kind of redundant assignment is harmless, unless of course they should not have this privilege at all.
 - Some users have an explicit privilege and others have an implicit privilege. Clicking the privilege check box once will assign the privilege explicitly to all users, clicking it a second time will remove all explicitly assigned privileges, and clicking that box a third time will leave the mixed privilege setting as it was originally.
4. Click any of these icons once to add an explicit privilege for all the users selected.
 5. Click OK.

Managing Group Membership

A group must exist in the composite domain before you can try to add a user to that group. See [Adding Groups to the Composite Domain](#).

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 TDV. To use the new privileges as an administrator, the user has to log out and re-log into the Studio.

To add or remove a user to or from a group in the Composite domain

1. Launch Manager from Studio or direct a Web browser to the Manager using one of these URLs:
 - When TDV is locally installed:

`http://localhost:9400/manager`
 - When TDV is not locally installed:

```
http://[TDV_host_name]:[port_number]/manager
```

After login, the MANAGER HOME page is displayed.

2. From the SECURITY tab, choose User Management.
3. Select the link in the # Groups column for the user.

The Edit the User's Group Membership window is displayed.

4. Select or clear the groups in which the user will be a member.
5. Click OK.

Viewing Group Membership

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

To view a user's group membership in the composite domain

1. Using Manager, choose User Management from the SECURITY tab.
If the user belongs to a single group, it is displayed in the Groups column listing for that user.
2. Expand the Groups column listing with a click on the expander icon and the list of groups is shown.

To view a group's membership


3. Using Manager, choose Group Management from the SECURITY tab.
4. Select the link in the # Users column for the group of interest. 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 as shown to select the group of interest.

Editing Group Membership

To edit group membership


1. Using Manager, choose Group Management from the SECURITY tab.
2. Select a single group using the check box in the left column.
3. Click Edit Users.
4. Add users who should belong to the group.



5. Click OK.

Edit group membership for a single user

6. Using Manager, choose User Management from the SECURITY tab.
7. Select a single user.
8. Click Edit Group Membership.
9. Select those groups to which the user should belong.



10. Click OK.

Changing Passwords for Other Composite Domain Users

TDV administrators with the Modify All Users and Modify All Resources rights can change any composite user password.

Note: TDV administrative rights do not permit management of LDAP domains, group membership, or passwords. Neither TDV administrators nor LDAP users logged into Studio can change LDAP profiles or passwords through normal TDV interfaces.

Changes made to the user rights profile take effect nearly immediately because TDV checks for appropriate rights every time feature access is attempted.

To change passwords for other Composite domain users

1. Launch Manager from Studio or direct a Web browser to the Manager using one of these URLs:
 - When TDV is locally installed:

```
http://localhost:9400/manager
```

- When TDV is not locally installed:

```
http://[TDV_host_name]:[port_number]/manager
```

After login, the MANAGER HOME page is displayed.

2. Go to the SECURITY > User Management page.
3. Select any user name link.
4. Click Edit User.
5. Reset the user's password and optionally change rights.
6. Click OK.

Changing Ownership of Resources

The administrator 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 reassigned to the nobody user. The nobody user cannot log in or be removed. The administrator can use the change ownership feature available in Studio to change the ownership from nobody to a valid user. Also, if there becomes a need to change the ownership of resources from one regular user to another, the change ownership feature is useful.

To change the ownership of a resource

1. In the resource tree, select the resource to change its ownership
2. Select Resource > Change Owner of <resource>, or right-click on the resource name and select 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.

Note: Ownership cannot be changed for system-owned resources or home folders.

3. Select a new owner's user name from the User drop-down list. Note that any restrictions that have been set on the selected new user will affect the user's ability to access the data source or execute stored procedures. Refer to the section [User Restrictions](#) for more information on the User Restrictions and how it impacts access to data.
4. Optionally, check the Apply the change recursively check box. The check box is selected 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 check box is enabled when the resource is not in your home folder.

5. To selectively change the ownership, select the Change if the current owner is box.

Resources within a container can have more than one owner, so this specification avoids unintentional transfers.

6. Click OK to see a list of resources ready to be transferred to the new owner.
7. View the list of resources.
8. Click Commit if the list is acceptable, or click Cancel to return to the previous window and make a different selection.

Clicking Commit changes resource ownership as specified.

Manage User and Group Privileges

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

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.

To facilitate decentralized management of specific resource privileges, 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 [Managing Security for TDV Resources](#) for more details on access privileges and for description of the Manager Resources pages.

LDAP Domain Administration

TDV supports the following domain types: composite, LDAP, Azure, and dynamic. This topic focuses on how to configure and administer LDAP domains for use with TDV.

- [LDAP Domain - Active Directory 2003 Limitation](#)
- [Configure the LDAP Properties File](#)
- [LDAP Domain Administration](#)
- [LDAP User Management](#)
- [Configuring LDAP for Use with Certificate Authentication](#)
- [Configuring LDAP for Use with Nested Groups](#)

About the LDAP Domain

TDV can leverage enterprise LDAP implementations of Active Directory, eDirectory domains, groups, and users to authorize views and use, creation and management of TDV-defined resources. Currently supported LDAP authentication servers are listed in the *TDV Installation and Upgrade Guide*.

To configure and manage LDAP domains, groups, users, and rights, the administrator needs two rights: Read and Modify All Users, and Access Tools.

Manager also lets the administrator specify LDAP authentication. LDAP configurations and usage are described in [Dynamic Domain Administration](#).

LDAP Domain - Active Directory 2003 Limitation

There is a known problem with Microsoft related to JRE that results in disabling of the 3DES_EDE_CBC transport layer security algorithm. If you encounter this problem, you can enable 3DES_EDE_CBC in <TDV_install_dir>/jdk/conf/security/java.security.

To re-enable 3DES_EDE_CBC:

1. Navigate to the <TDV_install_dir>/jdk/conf/security/java.security.
2. Open the file and remove 3DES_EDE_CBC from the jdk.tls.disabledAlgorithms setting.
3. Restart the TDV Server.

Configure the LDAP Properties File

Query searches for retrieving user and group information are controlled by a properties file. This properties file is in the following directory:

<TDV_install_dir>/conf/server/ldap.properties

The ldap.properties file contains relevant query parameters for the supported LDAP directory servers.

If you add LDAP domains to TDV Server, you should configure the ldap.properties file after installation and prior to adding and configuring the LDAP domain on the Studio Manager Domain Management page. You should also use the properties file to indicate whether you want permissions granted to nested groups.

Note: For TDV Active Cluster, custom configurations of the ldap.properties file are not copied to other TDV instances in a clustered environment. The ldap.properties file is **not** automatically synchronized with other machines in the cluster. Each server is considered LDAP independent unless you copy these files to all members of the cluster.

This section describes:

- [Structure of the LDAP Properties File](#)
- [Example of an ldap.properties File](#)
- [LDAP Properties File Symbols and Attributes](#)
- [Query Examples](#)

Structure of the LDAP Properties File

The ldap.properties file uses these conventions:

- PREFIX must be replaced with the values of the domain name or type of your directory services that are in use.

If you have multiple LDAP directories, you can replace PREFIX with the domain name given to each specific LDAP directory.

For example for single LDAP directories, use “activedirectory” or “restaurantOwners” where “restaurantOwners” comes from the domain name given to that LDAP directory.

- Property file variables are designated with a capital letter enclosed by angled brackets: <A>, , ... <X>.
- The ldap.properties file does not support mix and match domain names with domain subtypes.
- The MaxPageSize property controls the maximum number of entries that are returned in a single search result.

Used for Querying all Users

LDAP Property and Value	Description
<PREFIX>.all.users.search.context=<A>	Search-context used to find all users.
<PREFIX>.all.users.filter=	Filter to pass to a query for finding all users.
<PREFIX>.all.users.username.attribute=<C>	Username attribute to retrieve the name of user found from a query.
<PREFIX>.all.users.search.timeout=<D>	Search timeout value to limit the time for infinite search of all users. The timeout is in milliseconds and should be greater than 0. A value of 0 represents infinite timeout.

Used for Querying all Groups

LDAP Property and Value	Description
<PREFIX>.all.groups.search.context=<A>	Search-context used to find all groups.

LDAP Property and Value	Description
<PREFIX>.all.groups.filter=	Filter to pass to a query for finding all groups.
<PREFIX>.all.groups.groupname.attribute=<C>	Group name attribute to retrieve the name of a group found from a query.
<PREFIX>.all.groups.search.timeout=<D>	Search timeout value to limit the time for infinite search of all groups. The timeout is in milliseconds and should be greater than 0. A value of 0 represents infinite timeout.

Used for Authenticating LDAP Users

LDAP Property and Value	Description
<PREFIX>.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 can be set to False.
<PREFIX>.user.search.context=	Search-context used to find the user attempting authentication.
<PREFIX>.user.filter=<C>	Filter used to authenticate user in LDAP directory server. The USERNAME keyword will be replaced at runtime with the appropriate username.
<PREFIX>.user.username.attribute=<D>	User name attribute to retrieve the name of the user attempting authentication from a query.
<PREFIX>.user.search.timeout=<E>	Search timeout value to limit the time for infinite search of the user

LDAP Property and Value	Description
	attempting authentication. The timeout is in milliseconds and should be greater than 0. A value of 0 represents infinite timeout.

Used for Querying all Groups for a User

LDAP Property and Value	Description
<PREFIX>.user.groups.search.context=<A>	Search-context used to find all the groups for a user.
<PREFIX>.user.groups.filter=	Filter to pass to a query for finding the members of a group. The USERDN keyword is replaced at run time with the appropriate user distinguished name.
<PREFIX>.user.groups.groupname.attribute=<C>	Group name attribute for finding the name of a group to which a user belongs.
<PREFIX>.user.groups.search.timeout=<D>	Search timeout value to limit the time for infinite search of all the groups for a user. The timeout is in milliseconds and should be greater than 0. A value of 0 represents infinite timeout.

Example of an ldap.properties File

This section presents an example of the operational lines in the default ldap.properties file. [LDAP Properties File Symbols and Attributes](#) explains the symbols and attributes that can be used in the file.

```
iplanet.max.page.size=1000
```

```
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.max.page.size=1000
```

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

```
activedirectory.all.users.filter=(&(objectCategory=person)(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)(objectCategory=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)  
(objectCategory=person))
```

```
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)
(objectCategory=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, `|`, can be used to separate multiple search context property values. This can be interpreted as a disjunction (or).

LDAP Search Filter Symbols

Symbol	Name	Description
&	Conjunction	(and) All items in the list must be true.
	Disjunction	(or) One or more alternatives must be true.
!	Negation	(not) Item being negated must not be true.

Symbol	Name	Description
=	Equality	Items must be equal according to the matching rule of the attribute.
~=	Approximate equality	Items must be approximately equal according to the matching rule of the attribute.
>=	Greater than	First item must be greater than or equal to the second item according to the matching rule of the attribute.
<=	Less than	First item must be less than or equal to the second item according to the matching rule of the attribute.
=*	Presence	The entry must have the attribute. Returns the attribute value.
*	Wildcard	Searches for zero or more characters in the position of the attribute. A wildcard cannot be used for the placeholders USERNAME and USERDN (name and distinguished name of the current TDV user attempting LDAP authentication). In the following example, USERNAME is a placeholder: <pre>activedirectory.user.filter=(& (samaccountname=USERNAME) (objectclass=user))</pre> You cannot replace USERNAME with a wildcard to become: <pre>activedirectory.user.filter=(&(samaccountname=*) (objectclass=user))</pre>
\	Escape	Searches for the character following the backslash (asterisk, open parenthesis, or closed parenthesis) inside of an attribute value, rather than interpreting the character as part of search syntax.

LDAP Attribute Key

Symbol	Description
o	
ou	
cn	
dn	
dc	

Query Examples

This section shows example Directory LDAP server query examples. The Active Directory LDAP server configurations are similar except where the object class values, search contexts, and user or group attribute values can be different where:

<PREFIX> = { activedirectory }

- [Search for Specific Groups with a Group Filter](#)
- [Specify Multiple Locations to Find Users or Groups](#)
- [Disable Case Sensitivity for LDAP Authentication](#)
- [Get All Users](#)
- [Get All Users Under Container ou=people](#)
- [Get All Groups](#)
- [Get All Groups Under Container ou=groups](#)

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:

```
<PREFIX>.all.groups.filter=(&(objectclass=Y)(Z=cs_*))
```

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

Specify Multiple Locations to Find Users or Groups

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

```
<PREFIX>.*.search.context=CONTEXT_1|CONTEXT_2|...|CONTEXT_N
```

Example

```
<PREFIX>.all.groups.search.context=cn=users|cn=users2
```

This example is for groups under cn=users and cn=users2 search contexts.

Disable Case Sensitivity for LDAP Authentication

By default the TDV Server is case sensitive when used with either a directory domain, but that can be changed with ldap.properties.

Example

Enable case insensitive user names for LDAP authentication. How can the default case sensitive mode used for LDAP authentication be disabled?

Example solution:

```
<PREFIX>.user.username.comparison.is.case.sensitive=false
```

When the LDAP user name comparison is not case sensitive, the user “cn=sam,ou=users,dc=domain,dc=com” can log in to a TDV LDAP domain with user name sam or SAM. All variations of the user name used to log in to TDV tools 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. You can differentiate users by search context. For instance, in Active Directory, the

samaccountname attribute for a user object is globally unique in the LDAP server, but cn (common name) is not.

Get All Users

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

```
<PREFIX>.all.users.search.context=
```

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

```
<PREFIX>.all.users.filter=(&(objectclass=person))
```

To retrieve user names from the user object name attribute:

```
<PREFIX>.all.users.username.attribute=uid
```

To perform a search without a timeout:

```
<PREFIX>.all.users.search.timeout=0
```

Get All Users Under Container ou=people

This search context finds only groups under container ou=people:

```
<PREFIX>.all.groups.search.context=ou=people
```

This search finds only groups that match the objectclass filter:

```
<PREFIX>.all.groups.filter=(&(objectclass=person))
```

This search retrieves group names from this group object name attribute:

```
<PREFIX>.all.groups.groupname.attribute=cn
```

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

```
<PREFIX>.all.groups.search.timeout=0
```

Get All Groups

Using a null value (blank) starts searching from the root node and retrieves all groups:

```
<PREFIX>.all.groups.search.context=
```

To find only those groups that match the objectclass filter:

```
<PREFIX>.all.groups.filter=(&(objectclass=groupofuniquenames))
```

To retrieve group names within this group object name attribute:

```
<PREFIX>.all.groups.groupname.attribute=cn
```

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

```
<PREFIX>.all.groups.search.timeout=0
```

Get All Groups Under Container ou=groups

This search context finds only groups under the container ou=groups:

```
<PREFIX>.all.groups.search.context=ou=groups
```

To find only groups that match the objectclass filter:

```
<PREFIX>.all.groups.filter=(&(objectclass=groupofuniquenames))
```

To retrieve group names from this group object name attribute:

```
<PREFIX>.all.groups.groupname.attribute=cn
```

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

```
<PREFIX>.all.groups.search.timeout=0
```

Directory User Authentication

TDV LDAP user authentication dependent on directory servers requires configuration prior to successful user authentication through a TDV interface.

- The LDAP server must be configured for use.
- The LDAP domain must be configured for use in the Manager console.
- Specific Directory groups within the specified domain must be authorized to use TDV defined resources.

Note: All members of TDV authorized LDAP groups have the basic set of privileges granted to the all group. Other resource privileges and TDV 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 can authenticate properly using TDV resources.

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

```
<PREFIX>.user.filter=(&(uid=USERNAME)(objectclass=person))
```

```
<PREFIX>.user.username.attribute=uid
```

LDAP Domain Administration

LDAP domain administration involves the following tasks:

- [About Kerberos Configuration Files and LDAP Login Credentials](#)
- [Adding an LDAP Domain](#)
- [Working with Groups from an LDAP domain](#)
- [Editing LDAP Domain Connection Parameters](#)
- [Removing an LDAP Domain](#)

About Kerberos Configuration Files and LDAP Login Credentials

Kerberos configuration files often contain definitions for multiple Kerberos realms in the realms section of the file and a default realm specified in the libdefaults section.

Depending on what realm a user belongs to as specified in the libdefaults section of the Kerberos configuration file, their user name might need to be specified differently during login:

Realm Type	User Name Syntax	Example
Non-Default	<user>@<non-default_realm_name>	mmhennington@2K8.HLP.NET
Default	<user>	mmhennington

Passwords are treated one of the following ways:

Password	New Tickets Obtained
specified during login	The user principal and password are used to obtain: • A ticket-granting ticket from the Key Distribution Center (KDC) server • A service ticket for the Kerberos enabled LDAP server based on the new ticket-granting ticket
left blank during login	The specified user principal obtains a ticket-granting ticket from the ticket cache or the Local Security Authority.

The kinit command can be used to obtain a list of available tickets that reside in the ticket cache or Local Security Authority for principals.

Examples

To connect to an external LDAP server residing in the 2K8.HLP.NET realm and the Kerberos configuration file contains the realm settings for the 2K8.HLP.NET realm, but

the default realm is SUPPORT.NET, then the user name would have to be specified as <user>@2K8.HLP.NET.

Adding an LDAP Domain

You can add more than one LDAP domain to TDV Server, provided each of those domains has a unique name. The names “dynamic” and “composite” are reserved domain names in the TDV system.

To add an LDAP domain

1. Launch Manager.
2. From the SECURITY tab, choose Domain Management.
3. Click Add Domain.
4. 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).

5. Specify the LDAP directory type.

When using Novell eDirectory or Oracle Directory Server as the authentication source, select Other as the LDAP directory type and make changes in the ldap.properties file.

6. Type the path to the LDAP server in the Server URL field using the format:

```
ldap://<hostname:port>/<directory suffix>
```

```
ldaps://<hostname:port>/<directory suffix> (for secured LDAP)
```

Example:

```
<port> = 389 and <directory suffix> is dc=composite,dc=com
```

```
<port> = 686 and <directory suffix> is dc=composite,dc=com
```

Note: To use secured LDAP (LDAPS; default port 686), the TDV Server must have the keystore from the LDAP server placed in the trusted store.

Example for Windows Active Directory:

<directory suffix> is dc=composite,dc=com

7. Enter an administrative LDAP user name and password. The fully-qualified name always works, because it is unambiguous, but you can also use the common name.

Example for Windows Active Directory:

cn=Administrator,cn=Users,dc=composite,dc=com

8. Select Simple, Digest, or Kerberos authentication.

Option	Description
Simple	The client sends the LDAP server its fully qualified domain name and a clear-text password. This authentication mechanism can be used within an encrypted channel such as SSL, if it is supported by the LDAP server.
Digest	Sets the security authentication mechanism to DIGEST-MD5.
Kerberos	Enables authentication against a LDAP service that has Kerberos authentication, such as Microsoft Active Directory, without transmitting passwords, encrypted or otherwise, over the network. The authentication is done by obtaining a cached ticket-granting ticket from the system's underlying Kerberos implementation, and using it to obtain service tickets from the ticket-granting service for the other services in use. Required configuration: TDV JRE installation must be 1.6.0_44 or higher. Update the krb5.conf or krb5.ini file to include details of the Kerberos realm that the LDAP domain with Kerberos authentication belongs to with the following information: – A new realm tag, containing the Key Distribution Center (KDC) hostname, default domain name, KDC admin server hostname, KDC password server hostname, supported encryption types and principal name to user name mappings (if necessary). Other properties might

Option	Description
	be necessary, based on your unique Kerberos realm configuration. Cross-realm authentication is not supported. – A single or multiple entries in the domain_realms section to specify local domain name to Kerberos realm mappings. – Only if necessary, modify the libdefaults section of your configuration file. After this option is enabled, the behavior of TDV is modified in a way that will be unique to your location. It is recommended that you make your users aware that when logging into TDV as a user on an LDAP domain with Kerberos authentication, the password field is non-editable. For some additional information on how TDV user name and passwords are managed, see About Kerberos Configuration Files and LDAP Login Credentials.
Certificate-based	Choose this option if you are using certificate-based SSL Mutual authentication.
Certificate Field Name	This is the field name in the certificate keystore file that will be referenced in the mapping table for additional security (For example, CN, OU, etc.). If left blank, the global server setting value will be used. Refer to SSL Mutual Authentication Using Certificate for instructions on how to set the global setting.
Mappings Table Path	This is the path of the mappings table which contains a mapping of the certificate field name and the principal. If left blank, the global server setting value will be used. Refer to SSL Mutual Authentication Using Certificate for instructions on how to set the global setting.

9. If you want to synchronize the local external domain with the external domain server, then choose the option “Auto” in the Synchronization field. By default the Auto-synchronization is turned off.

Note: If you do not check the “Auto” synchronization option, synchronization of the user groups with LDAP server will happen only at login. You may choose to run the system procedure syncDomainGroups to manually do this operation. Refer to the *Application*

Programming Interface Guide Chapter *Web Services Operations* for a detailed usage of this procedure.

10. Specify the time period in which you want the synchronization to happen using the Periodic field. By default, this is set to 30 minutes.
11. Click OK.
12. Designate the LDAP groups (and users in those groups) who can access TDV resources. The access tools rights must be given in order to access the TDV tools.

Working with Groups from an LDAP domain

As you add groups from an LDAP domain to TDV, you are selecting groups or users from the LDAP server and adding them to the TDV Server. This enables differentiated group and user access, use, creation, and modification of TDV resources as LDAP authenticated users.

LDAP domain users must belong to at least one LDAP group selected to use TDV Server as an authenticated user. This enables you to implicitly assign rights, privileges, and ownership of defined resources.

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

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

A security check on user rights and privileges is made every time a request is made of TDV applications or defined resources. Authentication status with the LDAP domain is checked and maintained with a non-persisting session.

Authenticated users can own and use resources as defined by the rights and privileges assigned to them explicitly as individuals, or implicitly by group membership. Members of a group defined for use within TDV inherit all the rights and privileges defined for that group.

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

- [Adding a Group to an LDAP Domain](#)
- [Removing a Group from an LDAP Domain](#)
- [Viewing Group Membership](#)
- [Adding and Removing LDAP Users from a Group](#)

Adding a Group to an LDAP Domain

Adding external groups from an LDAP domain gives the TDV system a way to support differentiated access, and use of TDV-defined resources for selected groups without including the entire domain.

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

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

LDAP users are given rights and privileges to use TDV resources by explicit addition of the groups to which those members belong. LDAP managers should make sure that appropriate groupings of users are enabled to use TDV resources.

Set appropriate rights and privileges for LDAP groups in the same way that TDV 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 access data through JDBC, ODBC, or Web services clients. Unauthenticated users, anonymous, and dynamic users with pass-through authentication can be given privileges to view, access, and execute procedures on data resources, but they cannot receive rights to change TDV definitions and settings.

Groups of developers, operations users, and administrators should have explicit rights to access tools, and rights to read or modify TDV resources at design time.

After initial TDV use, LDAP domain users can be added directly to specifically defined TDV groups, thereby granting them implicit rights and privileges, or they can be given individual rights and privileges explicitly. Managing rights and privileges by group (role-based access control) makes it easier to control large groups of users.

See [Managing Security for TDV Resources](#) for more information.

To add users to LDAP domains and groups, see [Adding Users to TDV from an LDAP Domain](#), and [Add Users to Groups](#).

To add a group from an LDAP domain

1. In Manager, choose SECURITY > Domain Management and select the LDAP domain by using the row selector at the left of the Domain table.

2. Click Edit External Groups at the bottom of the table.

The Add External Groups window displays all groups in the LDAP domain.

3. Select those groups that you want to grant access to TDV 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.

Initially, no users are shown as members of the selected groups. Users from the groups appear in the TDV system after their first use of any TDV resource.

Removing a Group from an LDAP Domain

Removing a group from an LDAP domain deletes the LDAP group, all of its users, and all implicit rights and privileges on the TDV Server.

Resource definitions for /shared resources owned by users in a deleted group retain access privileges for the remaining LDAP groups to which they belong. 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 reintrospected to ensure that the resources remain accessible.

Group deletion removes all access privileges for the deleted group and its members. Group deletion also clears users' personal work space in the /users node. However, the external LDAP server is unaffected by these TDV definition changes.

To remove a group from an LDAP domain

5. In Manager, choose SECURITY > Domain Management and use the row selector at the left of the Domain table to select the LDAP domain.

6. Click Edit External Groups.

The window displays all groups in the LDAP domain.

7. Select the groups to remove.

Use the navigation arrows and page numbers at the bottom of the window to display additional groups.

8. Click OK.

Viewing Group Membership

The TDV administrator with Read All Users right can review and monitor user group membership from the Manager.

To view a user's group membership in an LDAP domain

9. In Manager, choose SECURITY > User Management.

The table of users can be filtered by domain and group, and sorted on multiple attributes.

10. 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

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

The TDV Server and Manager do not manage LDAP group membership. LDAP users can be added to TDV groups as described above, but LDAP groups are not modifiable from Manager.

To add or remove LDAP users to or from a group

11. In Manager, choose SECURITY > Group Management.

12. In the Users column, select the Edit Users icon for the group.

The Edit Group Membership window is displayed.

13. Add or remove users by checking or clearing the users.

14. Click OK.

Editing LDAP Domain Connection Parameters

You can edit an LDAP domain to change the connection parameters required to connect and read data from an LDAP authentication server. Everything but the domain name display text can be modified.

To edit an LDAP domain

1. In Manager, choose SECURITY > Domain Management
2. Select the Domain Name link for the LDAP domain that you want to edit.
3. Make your changes and click OK.

Removing an LDAP Domain

When you remove an LDAP domain, all users, groups, rights, and privileges associated with that domain are deleted and removed from TDV, and ownership of those users' shared TDV resources is moved to user nobody. LDAP users and groups on the LDAP server are unchanged.

Privileges to use resources owned by nobody stay the same for those groups and users who remain after the LDAP domain is removed.

To remove an LDAP domain

1. In Manager, choose SECURITY > Domain Management.
2. Select the row for the LDAP domain that is to be removed.
3. Click Remove Domain.
4. A verification prompt will ask whether you want to remove the selected domain.
5. Click OK and the domain, groups and users from that domain are no longer configured for use of TDV resources.

LDAP User Management

By default, without additional rights and privileges, all members of LDAP groups selected for use with TDV are able to log into JDBC, ODBC, and Web services clients configured for TDV. Rights to use TDV tools and to view and use other resources must be added to group definitions or assigned explicitly to users.

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

- [Adding Users to TDV from an LDAP Domain](#)

- [Remove LDAP Users from TDV](#)
- [Add Users to Groups](#)

Adding Users to TDV from an LDAP Domain

Typically, LDAP users are added to TDV indirectly by addition of their groups, with group rights and privileges appropriate to their role. To add users to an LDAP domain, the TDV administrator must first add the LDAP domain to TDV Server, and then add groups to that domain.

When adding a user from an LDAP domain, three conditions must be satisfied:

- The user's LDAP name and password are successfully authenticated with the LDAP server.

If LDAP authentication fails or the LDAP user does not belong to any local group definition, the user is not added to the LDAP domain in TDV Server and is not allowed to log into Studio.

- 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 TDV.

This user is added to each local LDAP group as a member where appropriate. The domain sync process adds or removes LDAP users to or from the appropriate local LDAP groups.

To add a user to an LDAP domain

1. Make sure the user belongs to a group in the LDAP server, see [Adding a Group to an LDAP Domain](#).
2. Start Studio.
3. In the login screen, log in with a valid LDAP user name, password, and domain.

Remove LDAP Users from TDV

Removing a user from a domain and group configured for use in TDV only removes the user locally from TDV Server while the user can still exist in the LDAP server and possess implicit rights and privileges given by membership in the LDAP domain and

group. Removing a user who is derived from an LDAP domain or group does not prevent the user from logging into the system again.

To remove an LDAP user and prevent that user from accessing resources defined by TDV, do one of these three things:

- Redefine the LDAP group membership at the source directory to exclude the user.
- Restrict rights and privileges for the entire LDAP group, and then explicitly assign rights and privileges to other members of that LDAP group, or make them members of a TDV group that gives them the needed rights and privileges.
- Remove the entire LDAP group from those included in the TDV external groups list.

TDV services are not normally used as interfaces to manage LDAP users directly. Typically, users and group memberships are managed using Active Directory interfaces. For example, if an individual LDAP Active Directory user needs to be refused TDV access, a management task must be performed directly on the LDAP server to change the column values for `memberOf`.

TDV users can be removed in Manager, but LDAP users selected for removal are only removed temporarily, because LDAP group membership continues to give implicit rights and privileges. Removing an LDAP user resets rights and privileges to those inherited through group membership. The user's Studio workspace is also deleted, but it is recreated when the user next logs into Studio.

Ways to work around this issue include:

- You can delete an LDAP group (see [Working with Groups from an LDAP domain](#)) to remove all group users, rights and privileges.
- You can initially grant no rights and privileges to the group, and then add selected members to other groups with the desired set of rights and privileges.

Add Users to Groups

The LDAP administrator can add LDAP and dynamic users to TDV groups to give them their rights and privileges implicitly. (For recommendations about dynamic users, see [Dynamic Domain Administration](#).)

Add LDAP users to TDV groups using the Group Management page in the Manager.

Configuring LDAP for Use with Certificate Authentication

You have the option to use LDAP with certificate authentication for TDV. If your site requires certificate authentication, you must modify two files.

To configure a TDV LDAP environment for use with certificate authentication

1. Configure TDV for use with LDAP.
2. If your LDAP server is using certificates signed by well known certificate authentication, use the LDAP URL that starts with `ldaps://` and skip to the final step.
3. If your LDAP server is using a certificate that is self-signed or signed by a an untrusted certificate authority:

Import the necessary chain of certificate signers to the `cis_server_truststore.jks` file.

You can use the Java key and certificate management utility (`keytool`) to import the certificates. For example:

```
<TDV_install_dir>/jdk/bin/keytool -import -alias myalias -trustcacerts -file Thawte.crt -  
keystore <TDV_install_dir>/conf/server/security/cis_server_truststore.jks
```

Import the necessary chain of certificate signers to `cacerts`, which is typically found in:

```
<TDV_install_dir>/jdk/conf/security
```

4. Restart the TDV Server.

Configuring LDAP for Use with Nested Groups

If you use LDAP with Active Directory, you have the option to use nested groups with TDV.

Nested groups allow you to define a group as a member of another group, allowing inheritance of permissions.

Note: The privileges on a parent group can be inherited only by up to 2 child level groups.

To configure your TDV LDAP environment for use with nested groups

1. Configure TDV for use with LDAP.
2. Locate the LDAP properties file, which is in the following directory:

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

3. In a text editor, locate the Active Directory section with group context search properties. For example:

```
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
```

4. Add the following two lines below the section:

```
activedirectory.user.parentgroups.filter=(&(distinguishedName=USERDN)  
(objectclass=group)(objectCategory=group))
```

```
activedirectory.user.parentgroups.attribute=memberOf
```

5. Save the file.
6. Restart the TDV Server.

Azure Domain Administration

TDV supports the following domain types: composite, LDAP, Azure, and dynamic. This topic focuses on how to configure and administer Azure domain for use with TDV.

About the Azure Domain

TDV can leverage Azure Active Directory, groups, and users to authorize views and use them in the creation and management of TDV-defined resources. To configure and manage Azure domains, groups, users, and rights, the administrator needs two rights: Read and Modify All Users, and Access Tools.

Currently supported Azure authentication servers are listed in the *TDV Installation and Upgrade Guide*.

Azure Domain Administration

Azure domain administration involves the following tasks:

- [Adding an Azure Domain](#)
- [Azure Group Management](#)
- [Editing Azure Domain Connection Parameters](#)
- [Removing an Azure Domain](#)

Adding an Azure Domain

You can add more than one Azure domain to TDV Server, provided each of those domains has a unique name. The names “dynamic” and “composite” are reserved domain names in the TDV system.

To add an Azure domain

1. Launch Manager.
2. From the SECURITY tab, choose Domain Management.
3. Click Add Domain.
4. Choose Azure.
5. 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).

6. Specify the Client ID. The ClientID parameter is the client/application Id of the application that is registered in the [Azure Active Directory portal](#). Contact your organization's system administrator for details about registering an application in the portal.

Note: When registering the application in the Azure AD portal, you need to choose the option "*Accounts in any organizational directory (Any Azure AD directory Multitenant)*", in order to login as an Azure domain user in TDV Web manager, TDV Web UI or Studio.

7. Enter an Azure Domain User Principal Name (eg: username@hostname.com) and password.
8. Click OK.
9. Designate the groups (and users in those groups) who can access the TDV resources.

Azure Group Management

Group and User Management in an Azure domain is similar to the LDAP Domain. Refer to [Working with Groups from an LDAP domain](#) for instructions about adding and removing groups and to view the group membership. Refer to [Adding Users to TDV from an LDAP Domain](#) for information on User Management.

Editing Azure Domain Connection Parameters

You can edit an Azure domain to change the connection parameters required to connect and read data from an Azure authentication server. Everything but the domain name

display text can be modified.

To edit an Azure domain

1. In Manager, choose SECURITY > Domain Management
2. Select the Domain Name link for the Azure domain that you want to edit.
3. Make your changes and click OK.

Removing an Azure Domain

When you remove an Azure domain, all users, groups, rights, and privileges associated with that domain are deleted and removed from TDV, and ownership of those users' shared TDV resources is moved to user nobody. Azure users and groups on the Azure server are unchanged.

Privileges to use resources owned by nobody stay the same for those groups and users who remain after the Azure domain is removed.

To remove an Azure domain

1. In Manager, choose SECURITY > Domain Management.
2. Select the row for the Azure domain that is to be removed.
3. Click Remove Domain.
4. A verification prompt will ask whether you want to remove the selected domain.
5. Click OK and the domain, groups and users from that domain are no longer configured for use of TDV resources.

Dynamic Domain Administration

TDV supports the following types of domains: composite, LDAP, and dynamic. This topic focuses on how to enable and administer dynamic domains for use with TDV.

- [About Dynamic Domains](#)
- [About Dynamic Domain Administration](#)
- [Enabling the Dynamic Domain](#)
- [About Group Administration for Dynamic Domains](#)
- [Considerations for Granting Privileges to Dynamic Domain Users](#)
- [About User Administration](#)
- [Adding Users to the Dynamic Domain](#)
- [Remove Users from the Dynamic Domain](#)
- [Dynamic Users Group Membership](#)
- [Viewing Dynamic User Group Membership](#)

About Dynamic Domains

Dynamic domains enable users to negotiate “direct” access to a secured data source by way of a TDV Server pass-through login. The TDV 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 (through TDV 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.

Pluggable Authentication Modules (PAM) use dynamic domain sessions for pass-through authentication and authorization. The dynamic domain must be enabled for

PAM security so that the user's dynamic session can be used for login pass-through and for use of ephemeral objects like Kerberos session tokens. Composite PAM wraps a Kerberos login module and uses any session token granted with positive authentication for use with data sources that are configured to use those Kerberos session tokens. See [Pluggable Authentication Modules](#) for more information.

With the dynamic domain, the TDV solution can be made more transparent. End users can use their existing login information for authentication with a data source to gain the same permissions they had in the past, without needing to log into TDV separately.

Note: Only one login is permitted for dynamic domain pass-through authentication. More than one pass-through-enabled data source can be used for federated queries if the data sources are set to authenticate using the same login.

Dynamic domains also accommodate a potentially large user base that does not require a TDV or an LDAP domain structure.

About Dynamic Domain Administration

Aside from enabling the TDV configuration settings to enable the dynamic domain no special user management is required to enable users to access resources given correct 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 Studio and other TDV administrative 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 TDV configuration setting. If it is enabled, dynamic users have Read access to basic resources in Studio. See [Considerations for Granting Privileges to Dynamic Domain Users](#) for the specific resources.

Enabling the Dynamic Domain

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

To enable the dynamic domain

1. In Studio, choose Administration > Configuration.
2. In the Configuration window, expand TDV Server > Configuration > Security.
3. Select Enable Dynamic Domain Login.
4. Set its value to True.
5. Click Apply.
6. Click OK.

About Group Administration for Dynamic Domains

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

The dynamic domain cannot use 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.

Considerations for Granting Privileges to Dynamic Domain Users

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

- 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; such users are not authenticated by the TDV system.

When the dynamic domain is enabled, dynamic users have default Read access to the following basic resources in Studio:

/

/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.
- 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 can use that user name.

Viewing Dynamic User Names

To view the dynamic user names that have been used

1. Open and log in to Manager.
2. Select SECURITY > Group Management.

3. In the # Users column, select the hyperlinked number representing the count of users in the all group row of the dynamic domain.

The USER MANAGEMENT page opens. The number link filters the display of users to show only those users in the all group of the dynamic domain.

About User Administration

Management of dynamic domain users is mostly passive as far as TDV 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 TDV. The new user is assigned an ID and can be treated like a normal user who has been cautioned not to expose sensitive resources. Dynamic domain users do not have a home directory; hence, they cannot create or own resources.

Considerations and Precautions

- 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 be given privileges to access public resources, while data sources enabled with pass-through login can independently authenticate and authorize dynamic users to gain access to secured data.
- Individual users in the dynamic domain can be deleted, but the all group and the dynamic domain cannot be deleted.
- Deleting 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.
- The ODBC manager may truncate the password at 14 characters.

Adding Users to the Dynamic Domain

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"
```

To add a user to the dynamic domain

1. Enable the dynamic domain as described in [Enabling the Dynamic Domain](#).
2. Connect to TDV Server through JDBC or ODBC, supplying the value dynamic for the argument domain in the connection string.

See the *TDV Client Interfaces Guide* for details.

Remove Users from the Dynamic Domain

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

If you were to remove a user from the TDV list of users registered in the dynamic domain, 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 the same user ID for login.

Dynamic Users Group Membership

Dynamic users are created at first login. All dynamic domain users automatically become members of the dynamic all group. See [Adding Users to the Dynamic Domain](#) for details on adding users to the dynamic domain.

It is not recommended to add dynamic users to the composite groups unless the group privileges are an entirely public set of permissions and no TDV rights are granted.

Viewing Dynamic User Group Membership

It is not recommended that dynamic users be given regular group membership because their user names can be used by any individual to gain access to group resources.

To view the group membership of a dynamic domain user

1. Launch Manager.
2. From the SECURITY tab, choose User Management.
3. Filter the list of users by selecting the dynamic domain using the domain pull-down.
4. If a user belongs to more than the dynamic group, an integer is displayed next to the icon in the #Groups column.
5. 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 groups with a check mark selection to show to which groups the user belongs. All users in the dynamic domain belong to the dynamic all group.

OAuth Domain Administration

This topic focuses on how to configure and administer OAuth2 domain for use with TDV.

- [About the OAuth2 Domain](#)
- [OAuth2 Domain Administration](#)
- [OAuth Supported Data Sources](#)
- [Defining Security Policies for Claims](#)
- [Hybrid OAuth2 Domain](#)

About the OAuth2 Domain

OAuth2.0 is the industry-standard protocol for authentication that apps can use to provide client applications with secure delegated access. OAuth works over HTTPS and authorizes devices, APIs, servers, and applications with access tokens rather than credentials.


For a list of TDV Data sources that support OAuth 2, refer to [OAuth Supported Data Sources](#).

OAuth2 Architecture

The high level architecture of the authentication protocol is illustrated below. An access token is generated by an IdP or an authorization server for a client application. It is used to login to the TDV Server. TDV Server uses the Signature in the token to validate credibility and communicates with the authorization server and rejects the token in case of a missing/incorrect information. Once validated, the token is pushed down to the protected resources, to retrieve the dataset based on the claim that is carried in the token.

Okta, Auth0, Azure AD and Google are the certified Identity Providers for TDV. It means that these are the Identity providers that TDV has been tested with.

To configure and manage OAuth2 domains, claims and privileges in TDV manager, the administrator needs two rights: Read and Modify All Users, and Access Tools.



The tokens used by OAuth2 for the purpose of authentication is called “access tokens”. Access tokens are used to prove an identity between consumers and service providers. The access token is an identifier, that is found in the http header of the requests that are made against the endpoints that support OAuth2.

JWT Tokens

The common format used for access tokens is **JWT**, which is represented in a specific format -

```
<header>.<payload>.<signature>.
```

Each of the parts of the token is in a JSON format.

Header

The Header of the access token usually consists of the encoding algorithm used and the type of the token.

Payload

The Payload part of the token contains Claims about the principal (user). Claims are a set of name-value pairs. Some claims are standard (such as the issuer, audience, subject) and some vary for each customer.

Signature

This is the signature of the tokens containing the Client Secret and different ways to sign them based on the chosen algorithm. The signature ensures the authenticity of the header and payload information of the access tokens.

Opaque Tokens

Opaque tokens are un-decodable arbitrary strings. The client application cannot read the contents of the token without interacting with the authorization server by querying an introspection endpoint or checking the claims of the token with the userinfo endpoint. Opaque tokens are widely used by many providers including Google. As the Opaque token is just an arbitrary string value, it is not possible to authorize the token as the JWT tokens (which has a specific format). An ID token is needed in this case, from the provider. The provider issues an ID token based on the "Scope" value passed from TDV. By default Scope is set to "openid". In certain cases "email" can be used as Scope. When there is no access token or ID token then "userinfo" will be used to validate the token.

OAuth Supported Data Sources

The following data sources are supported for OAuth:

- Facebook
- Google Ads
- Google Analytics
- Google BigQuery
- Google Calendar
- Google Contacts

- Google Directory
- Google Drive
- Google Sheets
- Hubspot
- Marketo
- NetSuite
- OData
- Oracle eloqua
- REST
- Salesforce.com
- Salesforce with SSO
- Siebel
- SOAP
- Twitter
- WSDL
- XML File
- XML HTTP

OAuth2 Domain Administration

OAuth2 domain administration involves the following tasks:

- [Adding an OAuth2 Domain](#)
- [Working with Claims from an OAuth2 domain](#)
- [Editing OAuth2 Domain Connection Parameters](#)
- [Removing an OAuth2 Domain](#)

Adding an OAuth2 Domain

You can add more than one OAuth2 domain to TDV Server, provided each of those domains has a unique name. The names “dynamic” and “composite” are reserved domain names in the TDV system.

To add an OAuth2 domain

1. Launch Manager.
2. From the SECURITY tab, choose Domain Management.
3. Click Add Domain.
4. 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).
5. Specify the domain type as OAuth2.
6. The table below describes the different fields in the New Domain window for the OAuth2 Domain. The values for some of the fields are available in the metadata that your IdP provides when you register TDV as an application. Refer to [TDV to IdP Field Mapping](#) for more information about the IdP specific fields.

Field	Description
Domain Name	The name of the domain.
Domain Type	Type of the domain. Choose OAuth2 to create an OAuth2 domain.
Issuer Value	Issuer indicates who issues the OAuth tokens. It typically matches with the value that appears in the oauth discovery url (For example, https://xyz-374535.okta.com/well-known/openid-configuration) It is a unique value in a TDV domain and therefore you cannot create two domains with the same issuer value. This is a required field.
User ID Claim	This is an optional field. However, entering the keyword "preferred_username" in this field will enable the user to login to the

Field	Description
	TDV applications (TDB Web UI, Studio and Manager) using the user name tied to the OAuth account.
Issuer Claim	Indicates the domain binding claim name. If this is not set, “iss” is the default value. This field is used to receive tokens that do not carry an “iss” claim.
Validation	The method used for validating the token. If set to AUTO, the order of validation is JWKS, Public Key and then Secret. JWKS - JSON Web Key Set endpoint contains information about public keys. The public keys are used to verify the JSON Web Token (bearer token) issued by the authorization server. Public Key - This is the authorization server’s Public Key. Public keys are in JSON Web Key (JWK) format and is used to verify the bearer token issued by the authorization server. Secret - The Secret is part of the signature in the bearer token. The signature is a hash generated by a cryptographic algorithm looking at the header and payload. The hash will be used to verify that the token created by the authorization server has not been tampered.
Claim Info	JSON: A set of claims as a name and value pair (with a separator, which is usually a colon) in JSON format, for which you can assign privileges and define rules and policies to access the published TDV resources. Multiple Claims can be given as comma-separated. The Claim values can also be in an array. You must define this field while creating the domain in order to automatically prepopulate the claim list when you use the “Edit External Claims” option in the Domain Management page to add the claims into your domain. URL: The Claim Info Endpoint is an OAuth2 protected resource that returns Claims about the authenticated End-User. Note: You can use tools like jsonlint.com to validate your JSON string.
Allow OAuth login via TDV	Check this option if you need to use an OAuth login to access the TDV tools (Currently supported in TDV Studio, Web Manager and Web UI).

Field	Description
tools	If you check this option, then you must provide the following properties:
Auth URL	This is the authorization URL. This is part of the endpoint metadata.
Token URL	This is the token URL. This is part of the endpoint metadata.
Client Secret	The client secret is provided by your Identity Provider. Contact your organization's administrator to find the Client Secret.
Client ID	The client id is provided by your Identity Provider. Contact your organization's administrator to find the Client Id.
Audience Url	The audience Url is provided by your Identity provider (Provided by the IdP Auth0).
Scope	Authorization scope, which is typically limited to the protected resources under the control of the client or as arranged with the authorization server (IdP). Limited scope is necessary for an authorization grant. Format: one or more strings, separated by spaces. In the IdP, scopes are defined when registering an application (TDV). This is usually done by your organization's administrator. TDV requires 'openid' scope at minimum and this is the default value when scope field is empty in TDV domain config. You may specify "email" as a scope when there is no human-readable username claim (for example, google as an IdP does not provide a username claim that can be used by TDV). When using "email" as a scope, remember to indicate "email" in the User ID Claim field.
UserInfo URL	This is the Userinfo URL. This is part of the endpoint metadata. Note: In case of IdP providing an Opaque access token (instead of a JWT token), it is important to give a valid UserInfo URL, especially if the ID token is not available.
Revocation URL	This is the Revocation URL. This is part of the endpoint metadata.

Field	Description
Annotation	This is an optional description for the domain.

7. Click OK.

Generating a Keystore Certificate

In order to login to the TDV applications (TDV Studio, Web UI and Web Manager) into an OAuth domain (using Auth0 IdP), you need to check if your truststore has the certificate. You may have to generate one if needed. If the certificate is missing, you may encounter an error such as:

```
Authentication Error: PKIX path building failed:  
sun.security.provider.certpath.SunCertPathBuilderException:  
unable to find valid certification path to requested target"
```

To generate a Keystore Certificate and resolve the above issue, follow these steps:

1. Generate the certificates based on the "domain name" in Auth0. (You can locate this in Auth0 application setting)
2. Run this command:

```
openssl s_client -showcerts -connect <domain name:443>
```

3. Copy the section of each cert and name them as A1.crt.....A4.crt (There should be 4 cert files).
4. Run the following command. You will need to run this 4 times since you have 4 files:

```
keytool -importcert -keystore cis_server_truststore.jks -file A1.crt -alias "A1A"
```

Note: You will be prompted for password. Use the option "changeit" and confirm the change.

5. Restart the server.
6. Login to your TDV application.

TDV to IdP Field Mapping

The table below is a reference for the fields on new Domain screen for an OAuth2 domain and its corresponding IdP fields. Note that these are only intended to help you find the information you need from your Identity Provider and should not be used as an exact map. The fields in your IdP are subject to change and you should rely on your IdP documentation for a more accurate and up-to-date information.

TDV Field	IdP Field	Sample Value
Auth0: Once you login to Auth0 and register the TDV application, open the application settings, where you can find the basic information such as the Client ID, Client Secret and your IdP domain. Using the IdP domain, you can access the metadata by adding <code>"/.well-known/openid-configuration"</code> to the IdP domain url. The metadata has the other endpoints that you will need, in order to create a the new OAuth2 domain in TDV.		
Issuer Value	issuer	<code>https://<sub-domain>.auth0.com/</code>
When Validation is "JWKS URI" or "Auto"	jwks_uri	<code>https://<sub-domain>.auth0.com/.well-known/jwks.json</code>
Auth URL	authorization_endpoint	<code>https://<sub-domain>.auth0.com/authorize</code>
Token URL	token_endpoint	<code>https://<sub-domain>.auth0.com/oauth/token</code>
Client ID	Client ID	Alphanumeric string from the application settings page
Client Secret	Client Secret	Alphanumeric string from the application settings page
UserInfo URL	userinfo_endpoint	<code>https://<sub-domain>.auth0.com/userinfo</code>
Revocation URL	revocation_endpoint	<code>https://<sub-domain>.auth0.com/oauth/revoke</code>

TDV Field	IdP Field	Sample Value
Okta - Once you login to Okta and register the TDV application, you can find the basic information such as Client ID and Client Secret from the My Applications page. Open the Default API Authorization Server from the Security settings of the application. You find the Url for the Metadata from which, you will be able to get the other details that are needed to create a new OAuth2 domain in TDV.		
Issuer Value	issuer	https://<sub-domain>.okta.com
When Validation is "JWKS URI" or "Auto"	jwtks_uri	https://<sub-domain>.okta.com/oauth2/v1/keys
Auth URL	authorization_endpoint	https://<sub-domain>.okta.com/oauth2/v1/authorize
Token URL	token_endpoint	https://<sub-domain>.okta.com/oauth2/v1/token
Client ID	Client ID	An alphanumeric string from the <i>My App</i> page.
Client Secret	Client Secret	An alphanumeric string from the <i>My App</i> page.
Userinfo URL	userinfo_endpoint	https://<sub-domain>.okta.com/oauth2/v1/userinfo
Revocation URL	revocation_endpoint	https://<sub-domain>.okta.com/oauth2/v1/revoke
Azure AD - Once you login to Microsoft Azure and register the TDV application, open the "App registrations" page, from where you can get most of the information you will need in order to create a new OAuth2 domain in TDV.		
	Directory (tenant) ID	An alphanumeric string from the App Registration page. You will need this to access the Metadata.

TDV Field	IdP Field	Sample Value
Issuer Value	issuer	https://login.microsoftonline.com/<tenant-id>/v2.0
When Validation is "JWKS URI" or "Auto"	jwtks_uri	https://login.microsoftonline.com/<tenant-id>/discovery/v2.0/keys
Auth URL	authorization_endpoint	https://login.microsoftonline.com/<tenant-id>/oauth2/v2.0/authorize
Token URL	token_endpoint	https://login.microsoftonline.com/<tenant>/oauth2/v2.0/token
Client ID	Application (client) ID	An alphanumeric string from the App Registration page.
Client Secret	Client credentials (secret)	An alphanumeric string from the App Registration page.
Scope	A new scope value that is added while registering the app.	You will need to add a new Scope for which an authorization grant will be provided.
UserInfo URL	userinfo_endpoint	https://login.microsoftonline.com/<tenant-id>/oauth2/v2.0/userinfo
Revocation URL	revocation_endpoint	https://login.microsoftonline.com/<tenant>/v2.0/revoke

Google: Once you login to Google and register the TDV application, open the application settings, where you can find the basic information such as the Client ID and Client Secret. You can access the metadata at: <https://accounts.google.com/.well-known/openid-configuration>

TDV Field	IdP Field	Sample Value
The metadata has the other endpoints that you will need, in order to create the new OAuth2 domain in TDV.		
Issuer Value	issuer	https://accounts.google.com
When Validation is "JWKS URI" or "Auto"	jwks_uri	https://www.googleapis.com/oauth2/v3/certs
Auth URL	authorization_endpoint	https://accounts.google.com/o/oauth2/v2/auth
Token URL	token_endpoint	https://oauth2.googleapis.com/token
Client ID	Application (client) ID	An alphanumeric string from the App Registration page.
Client Secret	Client credentials (secret)	An alphanumeric string from the App Registration page.
Scope	A new scope value that is added while registering the app.	You will need to add a new Scope for which an authorization grant will be provided. For example, if you choose the email scope, you need to mention "email" in the User Id Claim field, so that the email claim can be read for authentication.
UserInfo URL	userinfo_endpoint	https://openidconnect.googleapis.com/v1/userinfo
Revocation URL	revocation_endpoint	https://oauth2.googleapis.com/revoke

Additional IdP documentation references for registering TDV application:

Auth0 - <https://auth0.com/docs/get-started/applications/application-settings>

Okta - <https://developer.okta.com/docs/guides/customize-authz-server/main/>

Azure AD - <https://learn.microsoft.com/en-us/azure/app-service/configure-authentication-provider-aad>

Google - <https://developers.google.com/identity/openid-connect/openid-connect>

Working with Claims from an OAuth2 domain

In OAuth/Open ID Connect, a claim appears as a key(name)/value pair where the name is always a string and the value can be any JSON value. TDV uses the name (or key) and value to create a mapping to privileges configured in TDV's Claim Management page of Web Manager.

Adding a Claim to an OAuth2 Domain

Adding claims to an OAuth2 domain gives the TDV system a way to support differentiated access, and use of TDV-defined resources for selected claims based on a given token. Claims are basically principals recognized by the domain.

All the TDV functionalities supported by Domain Groups are supported by Claims (For example, Setting Resource Rights and Privileges, Row-Based and Column-Based Security Policy Assignments, Workload Management, etc.). Refer to the Administration Guide for the details about these functionalities.

Follow the steps below to add the Claims to the TDV OAuth2 domain.

1. Select the OAuth domain.
2. Click on the button Edit External Claims button to open the Edit External Groups/Claims window. When the Edit or Add External Claims window opens, the currently available OAuth2 Claims (If any defined) are displayed, and those Claims already selected for use within TDV are shown with a marked check box.

Note: You can see a preview of the list of claim combinations in the window

displayed, only if the "Claim Info JSON" field was defined while creating the domain.

3. Choose the claim(s) you want to register in TDV from the list displayed in the Edit External Groups/Claims window.

You can also manually add the claims (provided by your IdP) by choosing the "Add Claim" option from the "Claim Management" screen. To do this:

1. From the SECURITY tab, choose Claim Management.
2. Click Add Claim
3. Enter the Domain Name.
4. Enter the Claim Name. You can get the claims from your Identity Provider.
Note: Make sure to include the separator(:) between the claim name level.
5. Assign Rights and Privileges as needed.
6. In the "Claim Restrictions" section,
 - Select or deselect "Direct DS Data Access" restriction for the selected claim. This access control option is given to prevent direct data access to the data source asset. The restricted claim is subject to direct data source data access restriction and cannot access the data in the TDV tools.
 - Select or deselect "Direct DS Invocation" restriction for the selected claim. This access control option is given to prevent stored procedure calls to the data source asset. The restricted claim is subject to direct data source invocation restriction and cannot make stored procedure calls to the data source from any TDV tools.
7. Click OK.

After adding an OAuth2 Claim to TDV, rights can be assigned and data sources can define privileges for the Claim to use resource definitions and data. In other words, the Claim dictates the rights and privileges to use the TDV resources.

Example:

For example, the following JSON object contains three claims (sub, name, groups):

```
{
  "sub": "user@tdv.com",
  "name": "Mike John",
  "groups": ["admin", "leads"]
}
```

Each claim should be converted to an one dimensional string (for example, the above JSON object. can be represented as "groups:admin", "groups:leads"). This can be manually added in the Claim Management page, using "Add Claims" option. When you use the "Edit External Claims" button, a preview of the automatically converted claims are listed, if you provided the sample JSON string in "Claim Info" field of the domain configuration.

Rights and Privileges

When adding the claims to the OAuth2 domain, set appropriate rights and privileges for OAuth2 Claims in the same way that TDV groups and users get assigned rights and privileges. Privileges are assigned at the individual resource level to Claims in order to access data through JDBC, ODBC, or Web services clients.

Note: For a new OAuth domain, default privileges are defined in the domain's all group.

See [Managing Security for TDV Resources](#) for more information on assigning Rights and Privileges and defining Row Based Security rules. TDV Claims are used to create mappings between claims in OAuth token and TDV privileges, and not used to configure memberships as TDV Groups. To adjust claims in OAuth token, please check documentations of each IdP as they generate the token.

Note: TDV retrieves claims from Access Token(OAuth/OIDC), ID Token(OIDC) and UserInfo Endpoint(OIDC), in that order. If claims are found in the access token, TDV does not read the ID token or the User Info Endpoint.

Example:

For example, the following JSON object contains the claim "roles":

```
{
  "roles": ["admin", "reader", "contributor"]
}
```

You can either manually add these claims in the Claim Management Page or through the Edit External Claims option after configuring the Claim Info JSON field when setting up the OAuth domain.

The claim 'roles' has specific roles of the current user which might be just one or multiple values. In this case, 'roles' claim in TDV could have the parent level privileges that commonly apply to everyone, and 'roles:admin' should have some specific privileges that apply to admin role only.

Here are some examples of the token's claims and TDV claim management configuration.

Case 1: "roles": null

IF 'roles' claim exists without values, THEN 'roles' claim in the OAuth domain will be assigned privileges when configured.

Case 2: "roles": "admin"

IF 'roles' claim has a single string "admin", THEN 'roles' claim AND 'roles:admin' claim in the OAuth domain will be assigned privileges when configured.

Case 3: "roles": ["reader", "contributor"]

IF 'roles' claim has multiple strings in array, THEN 'roles' claim AND 'roles:reader' claim AND 'roles:contributor' claim in the OAuth domain will be assigned privileges when configured.

Note: Multiple claims' privileges are merged to the current user at login for above multiple combinations.

Removing a Claim from an OAuth2 Domain

Removing a claim from an OAuth2 domain deletes the OAuth2 claim and all implicit rights and privileges on the TDV Server.

To remove a claim from an OAuth2 domain

1. In Manager, choose SECURITY > Claim Management
2. Choose your OAuth2 Domain from the drop down.
3. Choose the Claim that you want to remove.
4. Click the Remove Claim button to remove the claim.

Viewing Claim Configuration

The TDV administrator with Read All Users right can review and monitor Claim Configuration from the Manager.

To View the Claim Configuration in an OAuth2 Domain

1. In Manager, choose SECURITY > Claim Management.
2. Click on the Claim that you want to view and click on Edit Claim to view/edit the Rights and Privileges assigned to a Claim.

Editing OAuth2 Domain Connection Parameters

You can edit an OAuth2 domain to change the connection parameters required to connect to the authentication server using the bearer token. Everything but the domain name display text can be modified.

To edit an OAuth2 domain

1. In Manager, choose SECURITY > Domain Management
2. Select the Domain Name link for the OAuth2 domain that you want to edit.
3. Make your changes and click OK.

Removing an OAuth2 Domain

When you remove an OAuth2 domain, all claims, rights, and privileges associated with that domain are deleted and removed from TDV.

To remove an OAuth2 domain

1. In Manager, choose SECURITY > Domain Management.
2. Select the row for the OAuth2 domain that is to be removed.
3. Click Remove Domain.
4. A verification prompt will ask whether you want to remove the selected domain.
5. Click OK and the domain and claims from that domain are no longer configured for use of TDV resources.

Defining Security Policies for Claims

In order to access claim-specific data from the TDV published resources, Row Based and Column-Based Security policies must be defined. Refer to the chapter [Managing Security for TDV Resources](#) and [Managing Column-Based Security](#) for instructions on how to define these policies.

Note: The name of the Rule, when setting up the policies, must match the Claim name.

The following built-in procedures can be used to read the token while defining the Security policies:

hasClaim()

getClaim()

For example, if a customer “Apex Systems” is passing a token with the Claim name “customer” (the claim value of which will be “Apex Systems”, then in the Predicate Rule definition of a Row-Based Security rule, specifying the following will return the value “Apex Systems”:

```
CustomerName=getClaim("customer")
```

Refer to the *Application Programming Interface Guide* chapter *Built-in Procedures* for a description about these procedures.

Note: Prior to defining these policies, you must ensure that appropriate access rights have been given for the TDV resources for the Claim. You can do that in TDV Studio. For instructions on how to provide access to TDV Resources, refer to [Rights and Privileges](#).

Hybrid OAuth2 Domain

TDV supports a Hybrid OAuth2 domain to enable the use of existing group privileges that are configured in other domains, for example, LDAP. To create a hybrid OAuth2 domain, follow these steps:

1. Launch Manager.
2. From the SECURITY tab, choose Domain Management.
3. Click Add Domain.

4. 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).

5. Specify the domain type as Hybrid OAuth2.

6. The table below describes the different fields in the New Domain window for the OAuth2 Domain. The values for some of the fields are available in the metadata that your IdP provides when you register TDV as an application. Refer to [TDV to IdP Field Mapping](#) for more information about the IdP specific fields.

Field	Description
Domain Name	The name of your domain.
Domain Type	The domain type. Choose Hybrid OAuth2 to create a hybrid OAuth2 domain.
Issuer Value	Issuer indicates who issues the OAuth tokens. It typically matches with the value that appears in the oauth discovery url (For example, https://xyz-374535.okta.com/well-known/openid-configuration) It is a unique value in a TDV domain and therefore you cannot create two domains with the same issuer value. This is a required field.
User ID Claim	This is an optional field. However, entering the keyword "preferred_username" in this field will enable the user to login to the TDV applications (TDB Web UI, Studio and Manager) using the user name tied to the OAuth account.
Issuer Claim	Indicates the domain binding claim name. If this is not set, "iss" is the default value. This field is used to receive tokens that do not carry an "iss" claim.
Group Claim	Indicates the key that is part of the token and holds the different groups defined in TDV. Refer to the section Claim Management for Hybrid OAuth2 Domains for information about how to add the TDV claims in your IdP.

Field	Description
Group Format	The domain and principal given in a specific format. if left blank, the default format is domain/principal. A delimiter is required.
Group Separator	Indicates the separator used in the list of TDV groups. The default is space.
Validation	The method used for validating the token. It can be Secret, Public Key or JWKS. JWKS - JSON Web Key Set endpoint containing information about public keys. The public keys are used to verify the JSON Web Token (bearer token) issued by the authorization server. Secret - The Secret is part of the signature in the bearer token. The signature is a hash generated by a cryptographic algorithm looking at the header and payload. The hash will be used to verify that the token created by the authorization server has not been tampered. Public Key - This is the authorization server's Public Key. Public keys are in JSON Web Key (JWK) format and is used to verify the bearer token issued by the authorization server.
Claim Info	JSON: The specific Claim as a name and value pair in JSON format, for which you can assign privileges and define rules and policies to access the published TDV resources. Multiple Claims can be given as comma-separated. The Claim values can also be in an array. URL: The Claim Info Endpoint is an OAuth2 protected resource that returns Claims about the authenticated End-User.
Auth URL	The authorization URL. This is part of the endpoint metadata.
Token URL	The token URL. This is part of the endpoint metadata.
Client Secret	The client secret is provided by your Identity Provider. Contact your organization's administrator to find the Client Secret.
Client ID	The client id is provided by your Identity Provider. Contact your organization's administrator to find the Client Id.

Field	Description
Scope	Authorization scope, which is typically limited to the protected resources under the control of the client or as arranged with the authorization server (IdP). Limited scope is necessary for an authorization grant. Format: one or more strings, separated by spaces. In the IdP, scopes are defined when registering an application (TDV). This is usually done by your organization's administrator. TDV requires 'openid' scope at minimum and this is the default value when scope field is empty in TDV domain config.
Annotation	This is an optional description for the domain.

7. Click OK.

Claim Management for Hybrid OAuth2 Domains

Hybrid OAuth2 domains enable the use of existing groups (claims) and the group privileges that are configured in other TDV domains. To do this,

1. Create a domain in TDV using the TDV Manager tool.
2. Create groups(claims) for the domain and assign privileges.
3. Open your IdP's Claim Management page for the registered TDV application.
4. Add the Claim(s) that you have created in TDV. Qualify the claim name with the domain name. If there are multiple claims from different domains, use a delimiter.

For more details about adding Claims in your IdPs, refer to your IdP documentation:

Okta - <https://developer.okta.com/docs/guides/customize-tokens-returned-from-okta/main/>

Auth0 - <https://auth0.com/docs/secure/tokens/json-web-tokens/create-custom-claims>

Azure AD - <https://learn.microsoft.com/en-us/azure/architecture/multitenant-identity/claims>

5. Create a Hybrid OAuth Domain using the instructions provided for the fields in the

above section [Hybrid OAuth2 Domain](#)

6. In the **Group Claim field**, enter the Claim name as defined in the IdP.
7. In the **Group Format field**, enter the format you have used to define your domain-claim name (For example, composite@admin, composite/admin)
8. In the **Group Separator field**, enter the delimiter you have used in the Group Claim field (For example, ":" or "," or something else.
9. Once other fields are entered Click on OK to save the Hybrid OAuth2 Domain.

TDV and SSL Authentication

TDV supports the Secure Socket Layer (SSL) protocol for authenticated data transfer among all components in the TDV environment, including:

- TDV Server (each instance) including inbound/outbound connections to data sources that support SSL
- Business Directory
- JDBC, web service, ODBC, ADO.NET and TDV Studio clients.

This section describes how to set up the working parts of SSL authentication. The following topics are covered:

- [Overview of TDV and SSL](#)
- [Keystore and Truststore Files for TDV](#)
- [Setting Which Protocols to Disable When Creating an SSL Connector](#)
- [Setting Up SSL](#)

Overview of TDV and SSL


SSL protocol lets you enforce a secure authentication regime to regulate access to data resources.

If you plan to build secure Java programs, it is recommended that you configure each TDV instance and each component with its own JKS (Java keystore) certificate prior to deployment. After using TDV to define user and group access profiles, you can begin to layer authentication protocols.

TDV includes a generic JKS file so that you can use it for development and testing of Web services and for JDBC, ODBC, or ADO.net clients secured over HTTPS ports.

You need Read All Resources and Modify All Resources rights to change the JKS file location, file type, or password.

The diagram shows the main components that communicate with each other and therefore require SSL authentication if the installation is to be secure.



Keystore and Truststore Files for TDV

Keystore and truststore files are the places where keys for secure communications are stored.

Each system component participating in SSL communication requires:

- A keystore file for its own key, which it furnishes to any other component that requests that it authenticate itself
- A truststore file for the keys of each other component that it trusts and needs to authenticate


Keystore and truststore files of the same type (for example, Java Keystore) have the same format.

- [Keys Passed between System Components](#)
- [Default Locations of Keystore and Truststore Files](#)
- [Keystore and Truststore Configuration Parameters](#)

Keys Passed between System Components

The figure below shows which keys are stored in each keystore and truststore file in a SSL-secured system, and which keys are sent to authenticate each component to another, and to itself (in the case of the TDV server).

In an environment with multiple instances of the TDV server, each server has its own key, and source and client truststores contain keys for each server instance with which they communicate.



Default Locations of Keystore and Truststore Files

The following table shows the default location and names of keystore and truststore files for components of a TDV installation.

Component	Default Location of Keystore and Truststore Files	Filenames
TDV and Business Directory	<TDV_install_dir>/conf/server/security	cis_server_keystore.jks
		cis_server_truststore.jks
Studio	<TDV_install_dir>/conf/studio/security	cis_studio_keystore.jks

Component	Default Location of Keystore and Truststore Files	Filenames
		cis_studio_ truststore.jks
JDBC	<TDV_install_dir>/apps/jdbc	cis_jdbc_ keystore.jks cis_jdbc_ truststore.jks

Keystore and Truststore Configuration Parameters

To access the keystore and truststore configuration parameters for TDV and its data sources, select Administration > Configuration from the main Studio menu, and in the Configuration window navigate to:

- Server > Communications

The following observations make it easier to understand the many keystore and truststore configuration parameters:

- The values of keystore and truststore parameters are all locally defined (that is, by TDV instance). They are not altered when restoring a backup and are not replicated in a cluster.
- Many of these parameters come in pairs ending with “Current” and “On Server Restart.” Changing the value of any “On Server Restart” parameter has no effect until the next server restart.
- Trusted certificate entries in truststore files can have any number of bits.
- The TDV Server configuration keystore key alias has a default value that names a sample keystore, so that TDV server can authenticate itself to sources and clients immediately upon installation.

The table below lists the keystore and truststore configuration parameters for the TDV server and its data sources.

Note: JDBC clients store SSL keys as values in JDBC parameters.

Parameter Name	Description
Keystore Key Alias (Server only; not Data Sources)	The alias name of the key entry used in SSL authentication to establish the identity of the server to external clients. For TDV server authentication to data sources, this value is optional. If a value is set, the key entry corresponding to the provided alias is used for client authentication, regardless of the contents of the foreign server's truststore or the results of any security callbacks.
Keystore File Location	The location of the keystore file used in SSL authentication to establish the identity of the server to external clients. The keystore file must contain exactly one key entry (a private key/certificate pair). It can also contain certificate entries from trusted certificate authorities that are used to validate the certificates that are presented by external clients.
	The password of the keystore file (and of the entries within it, which must be the same).
	The type of the keystore file. It must be a valid keystore type, such as JKS or PKCS12.
Truststore File Location	The location of the truststore file used in SSL authentication to decide what external clients the server should trust. The truststore file can contain certificates from trusted Certificate Authorities. These are used to validate the certificates that are presented by external clients. The TDV JDBC client driver uses the client system's truststore properties to validate the certificate: • javax.net.ssl.trustStore • javax.net.ssl.trustStorePassword • javax.net.ssl.trustStoreType The TDV Server certificate must be added to this client's truststore; otherwise, validation fails. The placeholder TDV certificate does not work after the client system truststore is enabled, unless it is added to the client truststore.

Parameter Name	Description
	The password of the truststore file (and of the entries within it, which must be the same).
	The type of the truststore file. Valid truststore types include JKS or PKCS12.

Setting Which Protocols to Disable When Creating an SSL Connector

To avoid known security flaws, it is best to disable SSLv3. By default, a Studio configuration parameter named Disabled Protocols for SSL Connectors does this for you. You can change the list of protocols to disable by modifying this parameter's comma-separated value string.

Removing the value string causes the default JRE settings to take effect. Under the default JRE settings, SSLv2, SSLv2Hello, and SSLv3 protocols are disabled for SSL sockets for incoming connections, and TLSv1 protocol is used for outgoing connections.

Note: Changing this value has no effect until the next server restart.

To change which protocols to disable when creating an SSL connector

1. Select Administration > Configuration from the main Studio menu.
2. Navigate to Configuration > Server > Communications > Disabled Protocols for SSL Connectors.
3. Change the comma-separated list of the protocols to disable.
4. Restart the server so that the changes take effect.

Setting Up SSL

The following topics describe the procedures used to set up TDV components for secure communications.

- [Using the Keytool Utility](#)

- [Installing a Truststore Certificate](#)
- [Setting Up Authentication between Studio and the TDV Server](#)
- [Setting Up Authentication between Client Applications and TDV Server](#)
- [Creating a JDBC Client Application with SSL Capability](#)
- [Setting Up Authentication between Client Applications and TDV Server over JDBC](#)
- [Setting Up Client Authentication for Web Data Sources](#)
- [Setting Up Client Authentication for Relational Database Sources](#)
- [Example - How to Obtain a third-party SSL Certificate and install into your Server and Studio Truststore?](#)
- The Business Directory Guide contains a section titled, “Keystore and Truststore Files for Business Directory.”

Using the Keytool Utility

Keytool is a publicly available command-line utility for managing public/private key pairs and associated Certificate Authorities (CAs). It is replicated in the following locations in TDV Server and Business Directory folders:

```
<TDV_install_dir>/jdk/bin
```

```
<BD_install_dir>/jdk/bin
```

Installing a Truststore Certificate

This topic describes how to check for and install a certificate in a truststore.

To check for and if necessary install the certificate in the truststore

1. In a browser, type the HTTPS URL of the TDV server.
2. Click the browser’s lock icon to view the certificate.

This icon is usually to the left of the URL field in the browser header.

3. Click the link to Certificate Information or View Certificate to see details.
4. Note the name of the party that the certificate is issued to.
5. Check the certification path to see how many certificates are in the chain.
6. In Studio, go to Administration > Configuration to open the configuration window.
7. Navigate to Server > Communications and find the following values:
 - Truststore File Location (Current)
 - Keystore Key Alias (Current)
8. Navigate to the location of the keytool utility:

```
cd <TDV_install_dir>\jdk\bin
```

9. Example of adding a certificate to the TDV Server truststore

Windows: open cmd.exe as Administrator privilege. <TDV_install_dir>\jdk\bin\keytool.exe -list -keystore <TDV_install_dir>\conf\server\security\<truststore_file_name> | findstr <certificate_alias>

UNIX: <TDV_install_dir>/jdk/bin/keytool -list -keystore <TDV_install_dir>/conf/server/security/<truststore_file_name> | grep <certificate_alias>

note: example provided is for installing a certificate to the TDV Server truststore. path for -keystore would need to change for doing this operation for TDV BD, Studio, JBDC, ODBC or ADO.NET clients.

10. Type the keystore password.

The result should be a line with the name of the certificate, the date it was installed, and “trustedCertEntry.”

The string “trustedCertEntry” confirms that the certificate is a trusted root in the truststore. If that string is not present, continue with the next steps to copy the certificate chain to the truststore.

11. Save the certificate chain (which you found in an earlier step) by copying it to a CAR file.
12. Use the browser’s utility (for example, its certificate export wizard) to save the file in a directory location where you can retrieve it later.
 - DER-encoded binary X.509 (.CER) is a recommended format.

13.Example of importing the certificate chain into the TDV Server truststore

Windows: open cmd.exe with Administrator privilege. <TDV_install_dir>\jdk\bin\keytool.exe -keystore <TDV_install_dir>\conf\server\security\<truststore_file_name> -import -alias <certificate_alias> -trustcacerts -file <CER_file>

UNIX: <TDV_install_dir>/jdk/bin/keytool -keystore <TDV_install_dir>/conf/server/security/<truststore_file_name> -import -alias <certificate_alias> -trustcacerts -file <CER_file>

Troubleshooting

You might encounter situations where you cannot make an SSL connection to the TDV server. This topic discusses a few of them.

- If you repeatedly receive an error like “PKIX path building failed” or “Unable to find valid certification path to requested target,” go back to [Installing a Truststore Certificate](#) and repeat the steps in which you use keytool to see whether the certificate is present in the truststore file.
- If the certificate entry in the truststore file is marked “trustedCertEntry” but you are still receiving certificate errors, probably your browser has not exported the complete certificate chain into C:\temp\mycertificate.cer.
- If the existing truststore contains too many certificate entries, you may want to remove it and create a new one. For the procedure, refer to [Creating a New Truststore File](#).

Creating a New Truststore File

Under certain circumstances you can remove the truststore and create a new one.

To create a new truststore file

14.If you want to remove an existing truststore file, back it up first and then remove it.

15.Use keytool to create the new truststore file:

```
<TDV_install_dir>\jdk\bin\keytool
```

```
-genkey
```

```
-alias <alias_for_your_truststore_file>
```

```
-keystore <TDV_install_dir>\conf\studio\security\<truststore_file_name>
```

note: example provided is for TDV Studio.

16. Check the contents of the new file:

```
<TDV_install_dir>\jdk\bin\keytool
```

```
-list
```

```
-keystore <TDV_install_dir>\conf\studio\security\<truststore_file_name>
```

The new file should contain one entry:

```
cis_studio, May 7, 2016, PrivateKeyEntry,
```

```
Certificate fingerprint (MD5): 01:12:23:34:45:56:67:78:89:9A:AB:BC:CD:DE:EF:FE
```

note: example provided is for TDV Studio.

Setting Up Authentication between Studio and the TDV Server

The encryption_util.bat can be used to update the authentication between Studio and the TDV Server. The utility will change and encrypt all the passwords for all the Studio installs in your environment. You must continue to use the keytool to update the passwords on the TDV Server.

If you decide not to use the encryption_util, you need to configure the JKS digital certificate that you intend to use for secured Web services and secured JDBC communications. The JKS digital certificate initiates and establishes SSL communication over HTTPS and LDAP ports.

If a truststore location is not specified, search for a keystore file in the following locations:

- <TDV_install_dir>/jdk/conf/security/jssecacerts
- <TDV_install_dir>/jdk/conf/security/cacerts

Assumptions

- TDV assumes that all passwords stored in all the keystore and truststore files are the same.

To use encryption_util.bat to validate Studio side authentication password is valid

1. On the Windows machine where Studio is installed, locate the encryption_util.bat script.
2. Use a command window to run the script using the following command:

```
encryption_util.bat -studioKeyStoreVerify
```

To use encryption_util.bat to update Studio side authentication

3. On the Windows machine where Studio is installed, locate the encryption_util.bat script.
4. Use a command window to run the script using the following command:

```
encryption_util.sh -toolsKeyStore -keyStorePassword somepassword -  
trustStorePassword somepassword -keyStoreChange
```

To configure SSL between Studio and the TDV Server

5. Obtain a JKS digital certificate from a Certificate Authority, or generate your own using keytool.
6. For Studio authentication to TDV, add the certificate to these files:

```
<TDV_install_dir>/conf/server/security/cis_server_truststore.jks
```

```
<TDV_install_dir>/conf/studio/security/cis_studio_truststore.jks
```


7. In Studio, and select Administration > Launch Manager (Web) to open the TDV Manager Web interface.
8. Log in to Manager.
9. In Manager, choose CONFIGURATION > SSL to display the SSL MANAGEMENT page.
10. In the New Value column next to Java Keystore File Location, enter the full path to the new JKS file on the server.
11. Click APPLY.

The REVERT button recovers the current value.

12. Change the Java Keystore File Type and the Java Keystore Password so that their values when the server restarts match the digital certificate being installed.
13. Change the passwords of the TDV Server and Studio truststores:

```
<TDV_install_dir>\jdk\bin\keytool.exe -storepasswd -new <your_password> -keystore
<TDV_install_dir>/conf/server/security/cis_server_truststore.jks
```

```
<TDV_install_dir>\jdk\bin\keytool.exe -storepasswd -new <your_password> -keystore
<TDV_install_dir>/conf/studio/security/cis_studio_truststore.jks
```

14. Restart the TDV Server to load the keystore and apply the changes.

Setting Up Authentication between Client Applications and TDV Server

Follow these instructions to set up authentication for non-JDBC connections.

To establish authentication for non-JDBC connections

1. Obtain a JKS digital certificate from a trusted Certificate Authority (CA), or generate your own using keytool.
2. For client authentication to TDV, add the certificate to the TDV Server and Studio truststore files:

```
<TDV_install_dir>/conf/server/security/cis_server_truststore.jks
```

```
<TDV_install_dir>/conf/studio/security/cis_studio_truststore.jks
```

3. Change the password of the TDV Server and Studio truststore files:

```
keytool -storepasswd -new <your_password> -keystore cis_server_truststore.jks
```

```
keytool -storepasswd -new <your_password> -keystore cis_studio_truststore.jks
```

4. Restart the TDV Server.

Creating a JDBC Client Application with SSL Capability

These are the general steps to enable a custom-developed client application to integrate with the SSL authentication capabilities of TDV.

1. Create your client application and declare a connection URL, using the following syntax:

```
jdbc:compositesw:dbapi@<fully_qualified_hostname>:<portnumber>
```

```
?domain=<cis_domainname>&dataSource=<data_source_name>&encrypt=true
```

For example, for Java you might add:

```
String url = "jdbc:compositesw:dbapi@localhost:9401?"
```

```
+"domain=composite&dataSource=cdspt&encrypt=true"
```

```
String user = "compUser";
```

```
String pass = "compPassword";
```

```
// Load driver
```

```
Class.forName("cs.jdbc.driver.CompositeDriver")&encrypt=true
```

```
// Create connection
```

```
conn = DriverManager.getConnection(url, user, pass);
```

For other URL properties, see JDBC Driver Connection URL Properties in the *TDV Client Interfaces Guide*.

2. Declare the username and password variables for use in the connection statement.
3. Optionally, find the JDBC driver name on the Data Source tab of the JDBC Data Source Administrator.
4. Optionally, write a sample program to test the connection URL.
5. Create or modify your client program so that it includes the connection syntax. For example, you must include a statement similar to the following to establish the connection:

```
conn = DriverManager.getConnection(URL, userName, password);
```

6. To set up authentication between JDBC client applications and the TDV server, you must declare a connection URL. This URL contains the following JDBC parameters where the keystore information can be specified.

JDBC Parameter	Description
validateRemoteCert	This property initiates a validation handshake of the TDV server-side certificate (on Windows, Linux is supported on TDV 8.3 or later). This validation includes checking that the cert is not expired, that the root certificate is installed on your machine, that the certificate is signed with the fully qualified dns name of the server, etc. In general, it applies all the validation rules of IETF RFC 8446 (Transport Layer Security 1.3). Note that this means you can not use self-signed certificates with validateRemoteCertificate=true (but it will work without this property).

JDBC Parameter	Description
validateRemoteHostname	False (default): No host name validation is performed. True: The csjdbc.jar compares the value of host in JDBC URL with the subject CN (common name) value in the certificate received from the targeted TDV Server. If host name validation fails, the connection is not established.

7. Restart the TDV Server.

Setting Up Authentication between Client Applications and TDV Server over JDBC

Client applications, including Studio, can connect to TDV Server over JDBC connections. For secure communications, you need to define secure authentication.

The steps are included in this section for convenience. For a full description of the URL properties, refer to these topics in the *TDV Client Interfaces Guide*:

- Defining a JDBC Client using a Connection URL
- JDBC Driver Connection URL Properties

Make sure your application has been designed to accommodate TDV SSL authentication for JDBC. See [Creating a JDBC Client Application with SSL Capability](#).

To define authentication between JDBC client applications and the TDV Server

The JDBC driver (.jar) file will be referenced by the client tool during configuration of the connection. For some client tools, the JDBC class name will be automatically discovered, for others the class name may need to be entered by the user as “cs.jdbc.driver.CompositeDriver”.

The fully populated connection strings are:

For a secure connection:

```
jdbc:compositesw:<userid>@<servername>:9401?domain=<domainname>&dataSource=<datasourcename>&encrypt=true&validateRemoteCert=true
```

For a non-secure connection:

```
jdbc:compositesw:<userid>@<servername>:9401?domain=<domainname>&dataSource=<datasourcename>
```

The 4 Java properties needs to be set for any process to configure access to the Trust Store that contains the public certificates from the Digicert certificate authority. The Trust Store provided should be placed in a secure location on disk and the two properties `javax.net.ssl.trustStore` and `javax.net.ssl.trustAnchors` should be adjusted to point to the file.

Java instance properties:

```
-Djavax.net.ssl.trustStore=<fully qualified trust store location on the client machine>
```

```
-Djavax.net.ssl.trustStorePassword=<password provided by administrative team>
```

```
-Djavax.net.ssl.trustStoreType=JKS
```

```
-Djavax.net.ssl.trustAnchors=<fully qualified trust store location on the client machine>
```

Instructions for an Encrypted Connection

The JDBC connection URL adds the following two properties which enables a security connection to TDV's published security configuration database using trusted certificates that result in all traffic including payload to be encrypted applying the SHA-2 256-bit encryption algorithm.

encrypt: The encrypt property instructs the JDBC driver to transparently switch to the secure SSL port (the configured data services port +2) for all communications.

validateRemoteCert: This property initiates a validation handshake of the TDV server-side certificate (on Windows, Linux is supported on TDV 8.3 or later). This validation includes checking that the cert is not expired, that the root certificate is installed on your machine, that the certificate is signed with the fully qualified dns name of the server, etc. In general, it applies all the validation rules of IETF RFC 8446 (Transport Layer Security 1.3). Note that this means you can not use self-signed certificates with `validateRemoteCertificate=true` (but it will work without this property).

Setting Up Client Authentication for Web Data Sources

When Web data sources require client authentication, a keystore must be specified to identify the TDV Server to the provider. The TDV Server configuration keystore key alias has a default value that names a sample keystore, so that you can use client authentication immediately upon installation.

If the TDV configuration settings for keystore alias (or for keystore alias) are set to null, the method described below to comply with client authentication requirements is used for Web data sources. The TDV configuration to use a specific keystore key alias overrides keystore specification defined on individual data sources.

To specify a keystore to comply with client authentication requirements

1. Open the Web data source in Studio.
2. Click the Advanced tab in the New Physical Data Source window.
3. Click Import Certificate Key Store from File to import the certificate.
Studio displays a dialog to specify the certificate.
You can choose a JKS or PKCS12 certificate keystore for authentication between TDV and any Web data source that requires a trusted certificate.
4. If you want to remove a keystore file, select it from the list and click Clear Certificate Key Store.
5. If you want to export the current certificate keystore to a JKS or PKCS12 file, click Export Certificate Key Store to File.
6. Optionally, set the Channel Pass-through field to a name or names that correspond to values passed in the HTTP request header for login authentication or for other purposes.

The Channel Pass-through is a comma-separated list of the names of HTTP request header properties that are to be passed through to the WSDL, XML, or HTTP data source.

If the data source expects a property with a name different from what was originally sent in the HTTP request header, you can change the property name. Put the name expected by the data source on the left side of an equal sign, and the original property name on the right.

7. Optionally, on the Advanced tab, add one or more environment variables the Environment Pass-through field to pass through to the WSDL, XML, or HTTP data source for login authentication or other purposes.

You can set environment variable names and values by calling the SetEnvironment procedure. See the Info tab for /lib/util/SetEnvironment in the Studio resource tree, or the *TDV Application Program Interface Guide*, for more information.

Property names in the Environment Pass-through field can be renamed before they are passed to the data source, just as they can with channel pass-through.

8. Optionally, specify the Execution Timeout (msec) period for REST data sources.

Execution Timeout is the number of milliseconds that an execution query on the data source is allowed to run before it is canceled. A value of zero (default) disables execution timeout, which you can use, for example, for resource-intensive cache updates set to run at non-peak processing hours.

Setting Up Client Authentication for Relational Database Sources

If all TDV configuration parameter for keystore alias are set to NULL but a registered relational data source requires client authentication, use the method described below to comply. The TDV configuration to use a specific keystore key alias overrides keystore specification defined on individual data sources.

You can put all of the data source keystore and truststore certificates in one file, even for multiple types of data sources.

To configure SSL between TDV Server and data sources registered in Studio

1. Obtain a JKS digital certificate from a Certificate Authority (CA), or generate your own using keytool.
2. For data source authentication to TDV, add the certificates to these files:

```
<TDV_install_dir>/conf/server/security/cis_server_truststore.jks
```

3. If necessary, change the password of the same two files:

```
keytool -storepasswd -new <your_password> -keystore cis_server_truststore.jks
```

4. Restart the TDV Server.

SSL Mutual Authentication

Any client that has SSL mutual authentication enabled with TDV server can access TDV Server resources directly without a password. User and Domain name is necessary for authorization.

Note: SSL Mutual Authentication is not supported for Composite Domain Users.

The server configuration setting “Enable SSL Mutual Authentication Login” can be tuned to enable or disable the SSL mutual authentication. By default, this setting is False (SSL mutual authentication is disabled). To enable this, click on Administration -> Configuration -> Server -> Configuration -> Security -> Enable SSL Mutual Authentication Login and set the value to True.

To enable SSL Mutual Authentication Login for Published Web services(Odata v4, SOAP, REST) and Data Services (OData v2 and OData v4):

1. Set Server configuration "Enable SSL mutual Authentication Login" to “true”.
2. Import Client Certificate to cis_server_truststore.jks.
3. Access published service with “https” and “username@domain” without a “password”

SSL Mutual Authentication Using Certificate

Certificate-based SSL Mutual Authentication provides an extra level of security by allowing only users with a certificate attribute defined in a mapping table. The mapping table can be any table or file datasource that your organization's administrator creates and maintains. It stores the TDV principals associated with the domain and the attribute of the certificate which will be referred by the SSL mutual authentication.

This certificate based SSL Mutual authentication is not supported for composite domain.

The server configuration setting “Enable SSL Mutual Authentication Login” can be tuned to enable or disable the SSL mutual authentication. By default, this setting is False (SSL mutual authentication is disabled). To enable this, click on Administration ->

Configuration -> Server -> Configuration -> Security -> Enable SSL Mutual Authentication Login and set the value to True.

Additionally, to globally configure the mapping table path and the mapping field name, tune the following configuration options found in Administration -> Configuration -> Server -> Configuration -> Security -> Authentication -> Certificate-based Authentication

- Certificate Field name - the field name in the truststore file that will be mapped with the user name.
- Certificate Mappings Table Path - the location of the mapping table.

Certificate Mapping

To provide an additional layer of security to the SSL Mutual Authentication, a mapping table can be maintained by your organization's administrator.

Certificates have multiple attributes in it (For example, the TDV bundled keystore has the attributes Common Name (CN), Organization Unit (OU), Organization Name (O), Locality Name (L), State Name (ST) and Country (C).

The table can be used to maintain a certificate-to-principal mapping. In addition to the users found in the certificate key chain, TDV provides an additional security to match the user with an attribute of the certificate that can be stored in the mapping table.

The mapping table can be a table or a file data source, and the path to the table or file will be provided in the LDAP domain definition. Following are the fields in the mapping table:

Field	Description
cert_field	Indicates the field name in the truststore file that will be mapped with the user name (For example, CN, OU, etc from the keystore file).
principal	Indicates the user corresponding to the certificate that will be used for authentication. Can be either the group or an individual user.
principal_type	Indicates whether the authentication is for all users in the group or only for the user.

Example - How to Obtain a third-party SSL Certificate and install into your Server and Studio Truststore?

For the purposes of illustration, we will assume the following:

- Your certificate provider sends you a certificate chain comprised of three certificates:
 1. clu_win64.com.cer
 2. sub1.clu_win64.com_clu_win64.com_.cer
 3. sub2.clu_win64.com_sub1.clu_win64.com_.cer
- You are using your a keystore file named 'root.jks' to store the Private Key.
- Your TDV server is running at : localhost:9400 and it is installed in the folder C:\apps\tdv.

Follow the steps below to obtain an SSL certificate and connect to Studio :

1. Create a Private key:

```
C:\apps\tdv\jdk\bin\keytool -genkey -alias AliasForMyCertificates -keyalg RSA -keystore
KeyStoreForMyCertificates.jks -keysize 2048
```

2. Use the Private key to create a CSR request (i.e. a file to Request a New Certificate

```
C:\apps\tdv\jdk\bin\keytool -certreq -alias AliasForMyCertificates -keystore
KeyStoreForMyCertificates.jks -file RequestTheCertificate.csr
```

3. To get an SSL certificate, submit the CSR file to a certificate provider (e.g Verisign or Thawte)*

The Certificate provider will respond by sending back a Private key file (e.g. *root.jks*). It is strongly recommended to have the provider send it as a JKS file as this is the format that TDV expects. You can verify that the file contains a Private key by using keytool to search for a "PrivateKeyEntry" as below:

```
"C:\apps\tdv\jdk\bin\keytool -list -v -keystore root.jks -storepass changeit
```

```
Alias name: clu_win64.com
```

```
Creation date: Jun 18, 2017
```

```
Entry type: PrivateKeyEntry
```

4. Copy root.jks to C:\apps\tdv\conf\server\security

5. Open Studio and set these 2 configuration settings:

Server >> Communications >> Keystore Key Alias (On Server Restart) = clu_win64.com

Server >> Communications >> Keystore File Location (On Server Restart) = C:/apps/tdv/conf/server/security/root.jks

6. Open a browser, connect to TDV and view the certificate details in the browser*. You should see the new certificate i.e. clu_win64.com)

7. Open Studio, click the "Encrypt" checkbox", and attempt to connect*. As the certificates are not yet in the Studio truststore, an error dialog will pop up stating that there is an RMI exception. This will be accompanied by a "PKIX path building" error in the cs_studio.log. This error is expected. It verifies that TDV is using the certificate to open an SSL connection with the client (in this case, Studio).

8. Import the certificates into the Studio truststore as below:

```
cd C:\apps\tdv\conf\studio\security
```

```
C:\apps\tdv\jdk\bin\keytool -import -alias firstalias -file
```

```
C:\apps\cert\clu_win64.com.cer -keystore cis_studio_truststore.jks -storepass changeit
```

```
C:\apps\tdv\jdk\bin\keytool -import -alias secondalias -file
```

```
C:\apps\cert\sub1.clu_win64.com_clu_win64.com_.cer -keystore cis_studio_truststore.jks -  
storepass changeit
```

```
C:\apps\tdv\jdk\bin\keytool -import -alias thirddalias -file
```

```
C:\apps\cert_test_for_CIS-66774\sub2.clu_win64.com_sub1.clu_win64.com_.cer -keystore  
cis_studio_truststore.jks -storepass changeit
```

9. Shut down and re-open Studio, click the "Encrypt" checkbox" and attempt to connect once more*. This time, Studio should connect.

Configuring Kerberos Single Sign-On

This topic introduces some of the configuration tasks you can perform to track information and control TDV behavior for Kerberos Single Sign-On (SSO). It describes how to integrate Kerberos authentication so that TDV can recognize Kerberos authenticated users and provide them with access to secured applications and resources.

- [About Kerberos Authentication and TDV](#)
- [Using Kerberos Authentication with TDV](#)
- [Using Kerberos Authentication with Published Resources](#)
- [Using Kerberos SSO Authentication with Data Sources](#)
- [Configuring Kerberos with Hive and Impala Data Sources](#)
- [Tip from an Expert on SSO Connection Issues](#)

About Kerberos Authentication and TDV

Enterprise users can leverage Kerberos infrastructure to authenticate just once to secure access to TDV-defined resources. The duration of an authenticated session is set by the Kerberos administrator. TDV supports pass-through of the Kerberos tokens from the authenticated client through TDV to the Kerberos server and to the data sources. The TIBCO® Data Virtualization Server, data sources, and clients of the TDV Server must be configured to support Kerberos token pass-through and SSO.

Kerberos must already be up and working in your environment prior to TDV Server and Studio installation. This includes the requirement that within the Kerberos installation every user, computer, and service has a Principal Name assigned to it. A Kerberos Principal represents a unique identity in a Kerberos system to which Kerberos can assign tickets to access Kerberos-aware services (in this case, a TDV Server). A Domain Administrator creates the Service Principal Name (SPN).

Supported Platforms and Requirements for Kerberos

TDV supports Kerberos SSO authentication so that Studio designers, developers, and end-users can log on to their operating system and use that authenticated login identity as established by the Kerberos authentication system. User identity and group affiliations that are established and maintained by an associated LDAP server can authorize use of applications, resources, and data provided that the authentication and authorization match.

Use of a Kerberos authentication system for Single Sign-On requires configuration of all components that makes use of the Kerberos authentication provider. If you use Active Cluster, each TDV node of the cluster must have an SPN (Service Principal Name) identifier, and a KeyTab file generated from that SPN.

Data Sources and Kerberos

In order to configure a data source for introspection with Kerberos, you need to first make sure that your TDV Server is configured for Kerberos authentication. Refer to [Configuring Kerberos for Use with TDV](#).

Login to Studio and open the New Datasource connection window for your datasource.

Note: If you do not have a valid ticket, you may need to get one. Refer to [Obtaining a Kerberos Ticket](#) for instructions on obtaining a Kerberos ticket.

The following table describes the required information in the Basic tab of the New datasource window.

For Kerberos specific fields refer [About Configuring Kerberos SSO for Data Sources](#) . Also refer to the Datasource-specific *Adapter Guide* for other connection properties.

Field	Description
Hostname	Name of the machine that is part of the domain which has been configured in the TDV Manager for use with Kerberos authentication. Refer to Configuring TDV for Use with Kerberos Authentication for setting up the domain and user groups in the TDV Manager.

Field	Description
Port	The Port in the domain that you need to connect to, in order to access the database.
Database Name	Name of the database
Login/Password	When using Kerberos authentication this is not required.

Using Kerberos Authentication with TDV

Using Kerberos authentication with TDV requires a number of configuration steps as described in these sections:

- [Configuring Kerberos for Use with TDV](#)
- [Configuring TDV for Use with Kerberos Authentication](#)
- [Setting Up SSPI Kerberos SSO](#)
- [Setting Up JGSS Kerberos SSO](#)
- [About Studio and SSO with Remote Desktop](#)

Configuring Kerberos for Use with TDV

The KDC Kerberos v5 Server must already be installed and running in your environment before you install TDV Server and Studio. You then configure the Kerberos system to use with TDV, establishing a security context in which Kerberos and the TDV identify each other.

Creating a Service Keytab

A Domain Administrator must follow these steps to create a KeyTab file:

1. Run the following ktpass command line utility:

```
ktpass -princ <servicename>/<hostname>.<domain>@<REALM> -mapuser <username>  
-pass <password> -crypto All -pType
```

```
[KRB5_NT_PRINCIPAL|KRB5_NT_SRV_INST|KRB5_NT_SRV_HST]
-out <name>.keytab
```

Note: The exact ktpass utility syntax depends on the environment you have set up. The following is a sample ktpass command line to create the keytab file for a QA environment:

```
ktpass -princ HTTP/krb5-win.sample.net@sample.NET -mapuser qa1 -pass tiger -crypto
All -pType KRB5_NT_PRINCIPAL -out krb5cis.keytab
```

A keytab file contains pairs of Kerberos principals and encrypted keys derived from the Kerberos password. The keytab file is used to identify TDV to Kerberos so that automated service processes can be run in this secure environment.

2. Copy the KeyTab file to a local directory accessible to Server.

In a later procedure ([Configuring TDV for Use with Kerberos Authentication](#)) you set a TDV configuration parameter value to the KeyTab file's directory.

3. Make sure that each Kerberos client has a Kerberos configuration file.

All clients (end-user computers, data sources, and the TIBCO® Data Virtualization Server) require a Kerberos configuration file to define the realm and the domain for authentication to the Kerberos Key Distribution Center (KDC).

Default locations for the Kerberos configuration file are shown in the table.

Operating System	Default Location and Filename
Windows	The location of the krb5.ini file varies with Windows version. For example: it can be C:\Winnt\krb5.ini, C:\Windows\krb5.ini, and so on.
Linux	/etc/krb5.conf
UNIX-based	/etc/krb5/krb5.conf

The Kerberos configuration file contains definitions like the following, where default_realm, kdc, default domain, and domain_realm have your implementation values:

```
[libdefaults]
```



```
default_realm = SUPPORT.NET
```

```
[realms]
```

```
SUPPORT.NET = {
```

```
    kdc = qaad.support.net
```

```
    default_domain = SUPPORT.NET
```

```
}
```

```
[domain_realm]
```

```
.support.net = SUPPORT.NET
```

A sample configuration file in Linux (krb5.conf) is given below:

```
[libdefaults]
```

```
dns_lookup_realm = false
```

```
ticket_lifetime = 24h
```

```
renew_lifetime = 7d
```

```
forwardable = true
```

```
rdns = false
```

```
pkinit_anchors = FILE:/etc/pki/tls/certs/ca-bundle.crt
```

```
default_ccache_name = FILE:/tmp/krb5cc_%{uid}
```

```
default_realm = SUPPORT.NET
```

4. (Optional) If you are using Active Directory, the server user account must enable the delegation property for the realm. The realm has a value like: SUPPORT.NET\qa1 if the Kerberos token from the client is used to access Kerberos-enabled data sources.

Creating a User KeyTab

Follow these steps to create a KeyTab file on the client. These steps are necessary for any data source requiring a user KeyTab.

5. Install TDV Server and Studio.
6. Change directory to JDK bin path
7. Type the following command:

```
ktab -a <username>
```

Note: The <username> must be the same user logged in to the machine and also a domain account.

8. Type the password for your user account.
9. Type the following command to get the Principal:

```
ktab -l -e -t
```

10. Type the following command to obtain and cache the Kerberos Ticket Granting Tickets (TGT):

```
kinit <username@REALM> and give your password.
```

11. To get a TGT ticket, execute the "klist" command.

Note: If a file cache is used as the default cache setting in the Kerberos configuration file, then the klist command returns a Ticket Cache that looks like:

```
FILE:/tmp/krb5cc_uid number
```

12. Restart the TDV Server.

13. Open the Studio and make a connection to a data source that supports Kerberos. Use the authentication you have used in the above steps.

Note: If you are unable to connect, run the “klist” command in your bin directory (TDV_InstallDir\bin\jdk) to check if you have a valid ticket cache. If it expired, run the “kinit” command in step 7 above to get a new ticket cache and try to connect again.

Configuring TDV for Use with Kerberos Authentication

The Studio Configuration window lets you map Windows domains to LDAP domains. The domain mappings link the authenticated users to the appropriate external group. Authentication is performed by the Kerberos system. Authorization to use TDV system, shared, and published resources depends on privileges assigned to users either directly or through their membership in LDAP groups. Kerberos-authenticated users with LDAP group affiliations are *implicitly* granted only those user rights and privileges that have been *explicitly* associated with the group.

By default all group and user rights and privileges are set to their most restrictive values. Rights and privileges must be set explicitly for Kerberos authenticated users to gain implicit rights and privileges by LDAP group membership. For further information, see the *TDV User Guide*.

Adding Domain and User Group

Follow these steps to add domain and user groups in the TDV Manager for use with Kerberos authentication:

1. Open Manager in a Web browser using a TDV administrative login that has Read and Modify All Users rights.
2. Choose SECURITY tab > Domain Management to access the DOMAIN MANAGEMENT page.
3. Add a domain and its LDAP-defined information.

TDV requires an administrative login to view externally available groups on the LDAP server.

4. Add external LDAP groups (using the Edit External Groups button) from the configured domain.
5. Once the group is created, choose the group, click on Edit Resource Privileges and grant Access Tools and Read All Resources rights.
6. Add a Windows Registry Key to enable Ticket-Granting-Ticket (TGT) Session Keys.

Change the `allowtgtsessionkey` registry REG_DWORD value to 1 to include a session key in the TGT. For Windows 10 or above, the registry location of `allowtgtsessionkey` is:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\Kerberos\Parameters
```

A value of 1 requires that a session key be returned with the TGT, and enables use of Kerberos TGT sessions.

Understanding Studio Kerberos Properties File

Each Studio client that is to be configured for use with Kerberos SSO must have a local copy of the `krb5.properties` file located in the `<TDV_install_dir>/conf/studio` directory. When Studio is starting up, the presence of this file triggers display of an SSO check box on the Studio login window.

Note: If Studio does not detect this file, or if the SPN value is set to a different TDV node, the Studio login uses Basic authentication, which requires the user to enter a valid user name, password, and domain for that server instance.

The Studio `krb5.properties` Service Principal Name (SPN) is derived from the TDV SPN. The TDV Server uses the Required Principal Name configuration parameter to authenticate the TDV service to Kerberos.

All Studio clients that connect to that TDV Server instance must use an SPN derived from the TDV instance's SPN. For example, if the Required Principal Name is `HTTP/krb5-win.support.net@SUPPORT.NET`, the derived SPN is `HTTP@krb5-win.support.net`. If a user of a Studio instance wants to use Kerberos SSO authentication to connect with a different TDV Server instance, the `krb5.properties` file SPN value must be changed to use that TDV instance's SPN name.

For more information about the `krb5.properties` file, see the `Krb5LoginModule` Java documentation.

For specific details on configuring this file for Security Support Provider Interface (SSPI), see [Preparing the Studio Kerberos Properties File for SSPI SSO](#).

For specific details on configuring this file for Java Generic Security Services (JGSS), see [Preparing the Studio Kerberos Properties File for JGSS SSO](#).

Setting Up SSPI Kerberos SSO

TDV and Studio can use Security Support Provider Interface (SSPI) on Windows for Kerberos Single Sign-On (SSO) for accounts that have a service principal name (SPN).

Kerberos SSO setup assumes the customer has used the Windows Services window to create a domain name service (DNS) account under Windows for the TIBCO® Data Virtualization Server. If you have no local DNS account for the TDV server, you need to set one up (for example, `SUPPORT <domain> + <user> + <password>`).

ODBC, JDBC, Studio, published Web services, and OData support SSPI-based Kerberos authentication.

Note: Data sources support only JGSS based Negotiate and Kerberos authentication. If you introspect a data source set up with SSPI Kerberos authentication, you get a 401 Authentication Error message.

To prepare TDV Server and Studio for SSPI Kerberos SSO, follow these procedures:

- [Setting up the TDV Service for SSPI Kerberos SSO](#)
- [Configuring TDV Server for SSPI Kerberos SSO](#)
- [Preparing the Studio Kerberos Properties File for SSPI SSO](#)

Setting up the TDV Service for SSPI Kerberos SSO

You need to configure services, import groups and assign privileges to set up the server side for SSPI Kerberos SSO.

Note: SSPI Kerberos Windows clients cannot authenticate the connection from TDV to the underlying data source. SSPI Kerberos Windows clients can, however, authenticate the connection to TDV.

Follow these steps to set up the TDV service for SSPI Kerberos SSO:

1. Open the Services window in your Windows environment.
For example in Windows 10, select Start > Settings. Type “services” in the Settings Search bar and click on “View Local Services” to open the Windows Services Manager.
2. Scroll to the TDV instance you are setting up for Kerberos SSO
3. Right-click the instance and select Properties from the context menu.
You need to configure SSPI Kerberos in this Properties window.
4. On the Log On tab, select the "This account" radio button and give the Service account credentials that was used to setup the instance. A Service account is used to access the network resources with service-specific rights and permissions while minimizing the permissions required for individual users using the application server. A Domain administrator will usually create this account.
5. In the Studio menu bar, select Administration > Launch Manager (Web).
6. Log into the Web Manager and go to the Domain Manager page to create an LDAP domain.
7. Click Domain > Domain Name.
Ordinarily you would select the Active Directory radio button for the server URL.
8. Enter an LDAP name and password.
9. Click the Add External Group button to import the group containing TDV.
Add External Group pulls in all groups, for which you then set appropriate privileges. Once the group is created, choose the group, click on Edit Resource Privileges and grant Access Tools and Read All Resources rights.

Configuring TDV Server for SSPI Kerberos SSO

The TDV Server supports Kerberos SSO authentication for the convenience of users who have already authenticated their identity to a Kerberos domain controller.

On the server side, you need to configure TDV for SSPI Kerberos SSO.

Follow these steps to configure the TDV parameters for Kerberos SSO authentication:

10. Log into Studio as the admin user.
11. Select Administration > Configuration from the Studio menu bar.

12. In the tree pane, navigate to the Server > Configuration > Security > Authentication folder.

13. Make the following change within that folder.

Parameter	Action and Description
Windows Domain Mapping	Enter a key-value pair. - The key is the reported Active Directory domain of an authenticated user. - The value is the name of the LDAP domain name you set up in the Web Manager.—the domain you set up in Setting up the TDV Service for SSPI Kerberos SSO. Often the Windows domain key and the LDAP name value are the same. Keys and values are case-sensitive. Provide all the case-sensitive combinations of the key-value pair to avoid any SSO authentication issues.

14. Navigate to the Kerberos subfolder.

15. Make the following changes within that folder.

Parameter	Action and Description
Allow Kerberos Authentication	Change this value to True. A warning helps you avoid inadvertently changing this without implementing Kerberos first.
Native	Make sure this is set to True for SSPI Kerberos.

16. Click OK.

17. Restart the Server.

Preparing the Studio Kerberos Properties File for SSPI SSO

On each Studio client that is to be configured for use with SSPI single sign-on, you need to set up the krb5.properties file.

Follow these steps to set up the krb5.properties file for SSPI single sign-on:

18. In <TDV_install_dir>\conf\studio, make a copy of krb5_sample.properties and rename it krb5.properties.

19. Open an editor such as Wordpad to edit krb5.properties.

20. Make sure Native is set to true for SSPI:

```
Native = true
```

21. Uncomment the lines that apply to SSPI, and fill in the values appropriate to the current TDV instance:

```
#####
```

```
#      SSPI      #
```

```
#####
```

```
##Service Principal Name or Service account
```

```
spn=HTTP/FullyQualified_HostName@Realm
```

```
spn=[domain name]\\[account name]
```

```
spn=[account name]@[domain name]
```

22. Restart Studio.

Setting Up JGSS Kerberos SSO

TDV and Studio can use Java Generic Security Services (JGSS) for Kerberos SSO.

Note: Kerberos SSO setup assumes the customer has created a domain name service (DNS) account under Windows for the TIBCO® Data Virtualization Server using the Windows Services window. If you have no local DNS account for the TDV server, you need to set one up (for example, SUPPORT <domain> + <user> + <password>).

To prepare TDV Server and Studio for JGSS Kerberos SSO, follow these procedures:

- [Setting Up the Windows Service for JGSS Kerberos SSO \(TDV Server on Windows only\)](#)
- [Configuring TDV Server for JGSS Kerberos SSO](#)
- [Preparing the Studio Kerberos Properties File for JGSS SSO](#)

Setting Up the Windows Service for JGSS Kerberos SSO (TDV Server on Windows only)

You need to configure services, import groups and assign privileges to set up the server side for JGSS Kerberos SSO.

Note: JGSS clients using Kerberos can authenticate both the connection to TDV and the connection to the underlying data source.

Follow these steps to set up the TDV service for JGSS Kerberos SSO:

1. Open the Services window in your Windows environment.
For example in Windows 10, select Start > Settings. Type “services” in the Settings Search bar and click on “View Local Services” to open the Windows Services Manager.
2. Scroll to the TDV instance you are setting up for Kerberos SSO.
3. Right-click the instance and select Properties from the context menu.
You need to configure JGSS Kerberos in this Properties window.
4. On the Log On tab, select the "This account" radio button and give the Service account credentials that was used to setup the instance. A Service account is used to access the network resources with service-specific rights and permissions while minimizing the permissions required for individual users using the application server. A Domain administrator will usually create this account.
5. In the Studio menu bar, select Administration > Launch Manager (Web).
6. Log into the Web Manager and go to the Domain Manager page to create an LDAP domain.
7. Click Domain > Domain Name; usually select the Active Directory radio button for the server URL.
8. Enter an LDAP name and password.
9. Click the Add External Group button to import the group containing TDV.

Add External Group pulls in all groups, for which you then set appropriate privileges. Once the group is created, choose the group, click on Edit Resource Privileges and grant Access Tools and Read All Resources rights.

Configuring TDV Server for JGSS Kerberos SSO

The TDV Server supports JGSS Kerberos SSO authentication for the convenience of users who have already authenticated their identity to a Kerberos domain controller.

Note: Data sources support JGSS based Negotiate and Kerberos authentication, but they do not support SSPI Kerberos authentication.

On the server side, you need to configure TDV for JGSS Kerberos SSO authentication.

Follow these steps to configure the TDV parameters for JGSS Kerberos SSO authentication:

10. Log into Studio as the admin user.
11. Select Administration > Configuration from the Studio menu bar.
12. In the tree pane of the Configuration window, navigate to the Server > Configuration > Security > Authentication folder.
13. Make the following change within that folder.

Parameter	Action and Description
Windows Domain Mapping	Enter a key-value pair. - The key is the reported Active Directory domain of an authenticated user. - The value is the name of the LDAP domain name you set up in the Web Manager.—the domain you set up in Setting Up the Windows Service for JGSS Kerberos SSO (TDV Server on Windows only). Often the Windows domain key and the LDAP name value are the same. Keys and values are case-sensitive.

14. Navigate to the Kerberos subfolder.
15. Make the following changes within that folder.

Parameter	Action and Description
Allow Kerberos Authentication	Change this value to True. A warning helps you avoid inadvertently changing this without implementing Kerberos first.
Debug Output Kerberos Authentication Enabled	Set this to True to have TDV write JDK's Kerberos implementation output messages to the cs_server.out in the logs directory.
KeyTab File	Enter the value point to the generated keytab file in the TDV Server. For example, when TDV is installed on a Linux server, the keytab file is in <TDV_install_dir>/kerb5cis.kt.
Kerberos Configuration File	The Kerberos configuration file contains the locations of Key Distribution Centers (KDCs) and admin servers for the Kerberos realms of interest, defaults for the current realm and for Kerberos applications, and mappings of host names onto Kerberos realms. This file is usually: • c:\WINDOWS\krb.ini (Windows) • /etc/krb.conf (UNIX) Changes to this value do not take effect until server restart.
Native	Make sure this is set to False for JGSS Kerberos.
Required Principal Name	Enter the SPN value established by invocation of the Kerberos setspn utility. TDV must know the SPN to address the Kerberos domain server.

16.Click OK.

17.Restart the Server.

Preparing the Studio Kerberos Properties File for JGSS SSO

On each Studio client that is to be configured for use with JGSS single sign-on, you need to set up the krb5.properties file.

Follow these steps to set up the krb5.properties file for JGSS single sign-on

18. In <TDV_install_dir>\confstudio, make a copy of krb5_sample.properties and rename it krb5.properties.
19. Open an editor such as Wordpad to edit krb5.properties.
20. Make sure Native is set to false for JGSS.
21. Copy the Specific User -- JGSS section for each user who intends to log in using SSPI Kerberos for single sign-on.

```
Native = false
```

22. Uncomment the lines that apply to JGSS, and fill in the values appropriate to the current TDV instance:

```
#####
```

```
# Default User -- JGSS #
```

```
#####
```

```
##Service Principal Name
```

```
spn=HTTP@dev-krb5-win.support.net
```

```
native=false
```

```
doNotPrompt=true
```

```
useKeyTab=false
```

```
debug=true
```

```
useTicketCache=true
```

```
renewTGT=true
```

```
krb5.conf=c:/krb5.conf
```

```
#####
```

```
# Specific User -- JGSS      #
```

```
#####
```

```
##Service Principal Name
```

```
#spn=HTTP@dev-krb5-win.support.net
```

```
#native=false
```

```
#principal=principalName
```

```
#doNotPrompt=true
```

```
#storeKey=true
```

```
#debug=true
```

```
#useKeyTab=true
```

```
#keyTab=keytab file
```

Studio is now ready.

23.Restart Studio.

About Studio and SSO with Remote Desktop

Studio SSO access can fail when used with certain local configurations of the Microsoft Remote Desktop. When users access Studio through Microsoft Remote Desktop, an SSO authentication failure can occur when the user elects to use Remote Desktop's "remember credentials" feature. When the "remember credentials" feature prompts the user for credentials before accessing the remote machine, the local Windows operating system presents a Kerberos ticket to the remote machine that is unable to be forwarded for use in SSO authentication for Studio.

To use SSO with Studio on a remote machine with Microsoft Remote Desktop, the sign-on must occur on the remote machine itself using that machine's login dialog prompt. This creates a user ticket that Studio can use for SSO access to a remote TIBCO® Data Virtualization Server.

Refer to [Obtaining a Kerberos Ticket](#) for information on how to obtain a Kerberos ticket.

Using Kerberos Authentication with Published Resources

Kerberos authentication is tightly integrated with the TDV authorization schemes that secure both the data and published resources. Data and derived resources are made available only to the authenticated and authorized users as identified by the Kerberos system and an associated LDAP server.

The following topics are covered:

- [Configuring New Web Services for Kerberos Authentication](#)
- [Verifying Kerberos for an OData Data Service](#)

Configuring New Web Services for Kerberos Authentication

If you are creating a new REST, SOAP, WSDL, or XML/HTTP data source that needs to use Kerberos authentication, follow the steps in this section.

Follow these steps to implement Kerberos authentication where TDV is the client:

1. Configure Kerberos as described in [Using Kerberos Authentication with TDV](#).
2. Create a new Web service for the REST, SOAP, WSDL, or XML/HTTP data source and publish a resource to the new Web service.
For information about publishing Web resources, see “Publishing Resources to a Web Service” in the *TDV User Guide*.
3. For a REST Web service, follow these steps:
Open the REST Web service that you want to configure for Kerberos authentication.
Select the REST tab.
Set the following Service properties to configure the Web service for Kerberos:
 - Enabled: true
 - Enable HTTP Negotiate: true
4. For a SOAP or WSDL Web service, follow these steps:
Open the SOAP or WSDL Web service that you want to configure for Kerberos authentication.
Select the SOAP tab.
Set the following Service properties to configure the Web service for Kerberos:
 - Enabled: true
 - Security Policy: /policy/security/system/Http-Negotiate-Authentication.xml
5. In Studio, create a new REST, SOAP, WSDL, or XML/HTTP data source, specifying the following parameters on the Basic tab.
REST connection parameters are shown in the following table:

Connection Type	Parameters to Specify
REST	Base URL: URL to access this REST data source using the syntax: Login: <LDAP login for this domain> Password: <LDAP password for this domain> Pass-through Login: Disabled Authentication: NEGOTIATE

Connection Type	Parameters to Specify
	Domain: not available Service Principal Name: HTTP@<machine>.<domain> Method: For the XML/HTTP protocol, under Operations, the specification for HTTP Verb must be POST or GET.

SOAP connection parameters are shown in the following table:

Connection Type	Parameters to Specify
SOAP	URL: <URL to access this SOAP data source> Login: <LDAP login for this domain> Password: <LDAP password for this domain> Pass-through Login: Disabled Authentication: NEGOTIATE Domain: <LDAP domain name> Service Principal Name: HTTP@<machine>.<domain>

WSDL connection parameters are shown in the following table:

Connection Type	Parameters to Specify
WSDL Connection Information	URL: <URL to access this WSDL> Login: <LDAP login for this domain> Password: <LDAP password for this domain> Pass-through Login: Disabled Authentication: NEGOTIATE Domain: not available Service Principal Name: HTTP@<machine>.<domain>

XML/HTTP connection parameters are shown in the following table:

Connection Type	Parameters to Specify
XML/HTTP Connection Information	URL: <URL to access this WSDL> Login: <LDAP login for this domain> Password: <LDAP password for this domain> Pass-through Login: Disabled Authentication: NEGOTIATE Domain: not available Service Principal Name: HTTP@<machine>.<domain> Method: For the XML/HTTP protocol, under Operations, the specification for HTTP Verb must be POST or GET.

6. Verify that the connection works:
Introspect the REST, SOAP, or WSDL data source.
Open the Web service operation and run it.

Verifying Kerberos for an OData Data Service


Follow these steps to verify Kerberos for an OData data service

1. Configure Kerberos as described in [Configuring TDV for Use with Kerberos Authentication](#).
2. Publish a table with primary key to a database in Data Services/Databases.
For example, publish the /shared/examples/ds_inventory/products table to a database such as Data Services/Databases/examples.
3. Open the database that contains the resource you published.
4. On the OData tab, check the Negotiate check box.
5. Verify the OData service for Kerberos authentication:
Create an XML/HTTP data source.

Introspect the XML/HTTP data source. TDV creates a new Web Service Operation for this data source in the resource tree.

Execute the Web Service Operation. TDV should display the results of the operation in the Results tab.

Click Details to see the XML.



Using Kerberos SSO Authentication with Data Sources

TDV supports Kerberos pass-through so that JDBC, ODBC, and ADO.NET clients and their users can directly authenticate themselves to the data sources to gain authorization to use secured data resources.

The following topics are covered in this section:

- [About Configuring Kerberos SSO for Data Sources](#)
- [About JDBC Clients and Kerberos SSO](#)
- [Setting the DSN for ODBC Clients and Kerberos SSO](#)
- [About ODBC Linux Clients and Kerberos SSO](#)
- [About ADO.NET Clients and Kerberos SSO](#)

About Configuring Kerberos SSO for Data Sources

When you add a new data source, you can specify Kerberos authentication for the data sources that support it. For further information, see *Working with Data Sources* in the *TDV User Guide*.

See the following table to understand pass-through authentication for each authentication protocol.

Authentication	Pass-through login: Enabled	Pass-through login: Disabled
BASIC	Basic login information such as the user name and password are passed through from the client to the data source to create a connection.	This setting is not recommended for Kerberos SSO as the client credentials are not passed through to the data source for negotiation of a connection. The data source adapter configuration settings are used to negotiate shared connections and used again for all users.
KERBEROS NEGOTIATE	If Kerberos tokens are present because they were generated by Kerberos SSO, then they are used to connect to the data source directly. Alternatively pass-through login information can be used to connect to the data source based on Kerberos authentication.	The login and password of the data source adapter configuration is used to login to the Kerberos KDC and then those credentials are used to connect to the data source.

For data sources that support Kerberos authentication, the following configuration parameters are important for use of a Kerberos authentication system, when adding a new data source using the New Physical Data Source dialog. Only the parameters that are appropriate for the specific data source need to be specified. For example, for Oracle only two parameters are required.

Parameter	Description
Pass-through Login	Must be Enabled for identification and use of the Kerberos authentication credentials of a client. With pass-through enabled, the

Parameter	Description
	client's Kerberos token is used to negotiate a connection with the data source. If pass-through login is not enabled, data source connection are negotiated with the Studio login and password (if saved) or with the TDV Server authentication status. When data is requested from a data source for the first time, pass-through login connection negotiation is used. Subsequent requests or executions sent to the same data source by the same user use the existing connection on an exclusive and restricted basis. Connections are not reused if they have been established with a data source configured to use pass-through login with a client-specific username and password. Only the user who created a connection can reuse that connection.
Authentication	Choose the KERBEROS option to use Kerberos authentication credentials with pass-through login to negotiate client connections to data sources. Client submission of the Kerberos credential through JDBC requires the code implementation of two properties from the krb5 login module. Refer to About JDBC Clients and Kerberos SSO . Choose NEGOTIATE to gain access to WSDL and the XML over HTTP data sources using Kerberos SSO authentication.
Service Principal Name	If you select KERBEROS or NEGOTIATE authentication, you need to provide a Service Principal Name (except in the case of an Oracle data source, which has separate configuration settings that point to the Kerberos Service Principal). The Service Principal Name (SPN) is the Principal name associated with the service account. It is a unique identifier that authenticates a service to Kerberos. The SPN for each data source is unique to that service. To get this value, run the following command from the command prompt: 1. Change you location to TDV_Install\jdk\bin 2. Type ktab and press ENTER.

About JDBC Clients and Kerberos SSO

The JDBC client must be written to call the `cs.jdbc.driver`. The driver class to connect with the TDV Server is named in the JDBC URL.

You can specify the location of the kerberos configuration file (`krb5.conf`) in the JDBC URL as a property in the form:

```
kerberos.krb5.conf={<path_to_krb5.conf>}
```

Sample JDBC Client Code

The following code example for a JDBC client can be adapted for your clients.

```
jdbc:compositesw:dbapi@Host_  
Name:9401?domain=MyDomain&dataSource=MyDataSource&authenticationMethod=  
kerberos&kerberos.  
spn=HTTP@FullyQualified_Host_Name
```

```
import java.security.PrivilegedExceptionAction;
```

```
import java.sql.Connection;
```

```
import java.sql.DriverManager;
```

```
import java.sql.ResultSet;
```

```
import java.sql.Statement;
```

```
import java.util.HashMap;
```

```
import java.util.Properties;
```

```
import javax.security.auth.Subject;
```

```
import javax.security.auth.spi.LoginModule;
```

```
public class TestCompositeKerberos {  
  
    static String loginModule = "com.sun.security.auth.module.Krb5LoginModule";  
  
    public static void main(String[] args) throws Exception{  
  
        // System.setProperty("java.security.krb5.conf","C:\\WINDOWS\\krb5.ini");  
  
        connectWithDefaultUser();  
  
        connectWithSpecificUser();  
  
    }  
  
  
  
    public static void connectWithDefaultUser(){  
  
        Connection con = null;  
  
        Statement stat = null;  
  
        try {  
  
            Class.forName("cs.jdbc.driver.CompositeDriver");  
  
            String url = "jdbc:compositesw:dbapi@Host_Name:9401?  
domain=MyDomain&dataSource=MyDataSource";  
  
            Properties props = new Properties();
```

```
props.put("authenticationMethod", "kerberos");
```

```
props.put("kerberos.spn", "HTTP@FullyQualified_Host_Name");
```

```
con = DriverManager.getConnection(url, props);
```

```
stat = con.createStatement();
```

```
ResultSet rs = stat.executeQuery  
("SELECT * FROM test.test.C_CUSTOMER");
```

```
rs.next();
```

```
System.err.println(rs.getString(2));
```

```
}catch (Exception except) {
```

```
    except.printStackTrace();
```

```
}finally{
```

```
    try{
```

```
        if(stat != null){
```

```
            stat.close();
```

```
        }
```

```
    }if(con != null){
```

```
con.close();
```

```
}
```

```
}catch(Exception e){}
```

```
}
```

```
}
```

```
public static void connectWithSpecificUser() throws Exception{
```

```
    Subject subject = getSubject(username, password);
```

```
    Subject.doAs(subject, new PrivilegedExceptionAction(){
```

```
        public Object run(){
```

```
            Connection con = null;
```

```
            Statement stat = null;
```

```
            try {
```

```
                Class.forName("cs.jdbc.driver.CompositeDriver");
```

```
                String url = "jdbc:compositesw:dbapi@Host_Name:9401?  
domain=MyDomain&dataSource=MyDataSource";
```



```
Properties props = new Properties();
```

```
props.put("authenticationMethod", "kerberos");
```

```
props.put("kerberos.spn", "HTTP@FullyQualified_Host_Name");
```

```
con = DriverManager.getConnection(url, props);
```

```
stat = con.createStatement();
```

```
ResultSet rs = stat.executeQuery  
("SELECT * FROM test.test.C_CUSTOMER");
```

```
rs.next();
```

```
System.err.println(rs.getString(1));
```

```
return null;
```

```
}catch (Exception except) {
```

```
except.printStackTrace();
```

```
return null;
```

```
}finally{
```

```
try{
```

```
if(stat != null){
```

```
stat.close();
```

```
}
```

```
if(con != null){
```

```
con.close();
```

```
}
```

```
}catch(Exception e){}
```

```
}
```

```
}
```

```
});
```

```
}
```

```
private static Subject getSubject(String principle,  
String password)throws Exception{
```

```
    LoginModule krb5Module = (LoginModule)Class.forName  
(loginModule).newInstance();
```

```
    Subject subject = new Subject();
```

```
    HashMap sharedState = new HashMap();
```

```
sharedState.put("javax.security.auth.login.password", password.toCharArray());
```

```
sharedState.put("javax.security.auth.login.name", principle);
```

```
HashMap options = new HashMap();
```

```
options.put("principal", principle);
```

```
options.put("debug", "true");
```

```
options.put("storeKey", "true");
```

```
options.put("useFirstPass", "true");
```

```
krb5Module.initialize(subject, null, sharedState, options);
```

```
try{
```

```
    krb5Module.login();
```

```
    krb5Module.commit();
```

```
}catch(Exception e){
```

```
    e.printStackTrace();
```

```
    krb5Module.abort();
```

```
    return null;
```

```
}
```

```
return subject;
```

```
}
```

```
}
```

Setting the DSN for ODBC Clients and Kerberos SSO

Using Kerberos with an ODBC client requires driver configuration. ODBC Client applications can connect with TDV using a 32-bit or a 64-bit TDV driver. The computer on which the ODBC client application resides must have one of the TDV drivers installed and configured.

Follow these steps to add the Kerberos system DSN:

1. Log in as an administrator.

NOTE: The ODBC manager may truncate the password at 14 characters.
2. Run one of the following to install the ODBC driver on a Windows client:
 - CsOdbcInstall<version>.exe installs cis<version>.dll
 - CsOdbcInstall<version>_x64.exe installs cis<version>_x64.dll
3. Open the ODBC Data Source Administrator, accessible through the Windows control panel named Data Sources (ODBC) under Administrative Tools.
4. Select the System DSN tab, click the Add button, select the system data source you just installed, and click Finish.

The ODBC Driver Configuration window appears.

5. In the ODBC Driver Configuration window, set the Integrated Authentication field to Kerberos.

This System DSN enables TDV to recognize the specified data source and catalog.

6. Define the Kerberos SPN using the Microsoft format:

```
HTTP/FullyQualified_HostName@Realm
```

7. Click Test.
8. Once the Test Connection is successful, click OK.

About ODBC Linux Clients and Kerberos SSO

The TDV Linux ODBC driver can support Kerberos as long as the following conditions are met:

- Header files of a GSSAPI library that implements Kerberos (MIT, Heimdal, Centrif, or another) are available for dynamic loading at run time.
- The system has LIBDL and a GCC compiler.
- You run kinit <user>.

If you do not run kinit, you get a GSS library error message that reads, 'Failed to connect to DSN 'ktest': Exception: Problem initializing context.'

About ADO.NET Clients and Kerberos SSO

ADO.NET clients can be configured to use Kerberos authentication to connect with and use published TDV data sources. See the *Client Interfaces Guide* for instructions on how to install, configure, and use the ADO.Net driver.

When adding a new data connection to the TDV Server:

- Use the Advanced properties to enable Kerberos in the Integrated Authentication field.
- Set the Kerberos SPN field to the Microsoft format of the TDV SPN:

```
HTTP/FullyQualified_HostName@Realm
```

Configuring Kerberos with Hive and Impala Data Sources

This section contains instructions for how to configure your Hive data connection, including Impala, for use with Kerberos:

Follow these steps to configure Hive/Impala datasources with Kerberos:

1. If you are using Hive 0.13, make sure that you have the Apache HIVE-6486 patch installed for the JDBC driver.
2. Open Studio.
3. Open or add a new Hive or Impala data source.
4. Select the Basic tab.
5. Select Kerberos for the Authentication field.
6. Select the Advanced tab.
7. In the Connection URL Pattern field add a semicolon to the end of the URL.
8. In the Connection URL Pattern field add the following elements to the end of the connection URL.

```
jdbc:hive2://<HOST>:<PORT>/<DATABASE_
NAME>;<Principal>;<auth>;<kerberosAuthType>;
[hive.server2.proxy.user=<DELEGATED_USER>]:
```

Property	Description of Necessary Value
<Principal>	The Kerberos SPN for the Hive instance. For example, principal=hive/DBName-016.kt.support.net@KT.SUPPORT.NET.
<auth>	Specifies that the authentication method is Kerberos. For example, auth=kerberos.
<kerberosAuthType>	Specifies to use the private credentials inserted into the Subject by Kbr5. For example, kerberosAuthType=fromSubject

Property	Description of Necessary Value
[hive.server2.proxy.user=<DELEGATED_USER>]	An optional token that if used will be replaced by the cis user at run time. The proxy.user is used to access the Hive Server2 proxy functionality.

For example, if the value in the Connection URL Pattern field is:

```
jdbc:hive2://<HOST>:<PORT>/<DATABASE_NAME>
```

Modify it to become:

```
jdbc:hive2://<HOST>:<PORT>/<DATABASE_NAME>;<Principal>;<auth>;<kerberosAuthType>
```

For example:

```
"jdbc:hive2://HiveHost:10000/default;principal=hive/localhost.localdomain@EXAMPLE.COM;auth=kerberos;kerberosAuthType=fromSubject"
```

9. Save the data source.

Tip from an Expert on SSO Connection Issues

This section explains some troubleshooting tips while configuring Kerberos authentication for TDV.

Kerberos Test Utility

TDV has a Kerberos Test utility that will help you troubleshoot any errors you encounter while configuring Kerberos authentication. To run this utility:

1. Open the command line.
2. Change directory to TDV_Install\bin
3. Run test_kerberos.bat. In unix, you can run the utility from <TDV_Install>/bin/test_kerberos.sh

4. Fill in the appropriate fields and click on “Start”. The Debug pane below displays the log which will help you to debug any issues encountered.

The following table describes the different fields in the Test tool:

Field	Description
Server/Client	Choose Server or Client mode depending on whichever Kerberos authentication you are trying to debug.
Keytab File Name	The name and location of the Server Key tab file. Refer to Configuring Kerberos for Use with TDV for instructions on how to create this file. The Domain Administrators are generally responsible for creating this file. While testing a Client Kerberos authentication, this field is not required.
SPN	This is the Service Principal Name. You can get this from the Properties file. For Server authentication test, it is given in the form: <code>HTTP/FullyQualified_HostName@Realm</code>. For example, <code>HTTP/krb5-win.sample.net@sample.NET</code> For Client authentication test, it is of the form: <code>HTTP@FullyQualified_HostName</code>. For example, <code>HTTP@krb5-win.sample.net</code>
Port number	In a Server mode, this will be the port that the Server will listen. In a Client mode, this is the Server port.
Server Host	The host name/IP where TDV Server is running. If the server is running in your localhost, check the “Use localhost” option. This field is disabled when testing a Server.
Kerberos Conf File	The Kerberos configuration file with the full path. For example, “c:\windows\system32\krb5.ini”.
Username/Password	This is not required when testing the Server authentication. In a client mode, this is the Client authentication credentials (not

Field	Description
	the windows or the ticket-cache authentication).
Use Cached Credentials	Check this option if you want to use the cached user credentials.

Troubleshooting Tips

Buffer Size

You may encounter a buffer size related error when connecting to TDV with SSO, check the Exception stack trace. If the source of the exception is “java.io.IOException: FULL head”, you need to configure a larger header buffer size. The default size is 4096.

Follow these steps to modify the setting:

1. Log into Studio as the admin user.
2. Select Administration > Configuration from the Studio menu bar.
3. Locate the two Head Buffer Size configuration parameters:
 - HTTP > Header Buffer Size
 - HTTPS > Header Buffer Size
4. Increment the value by 4096 until you no longer get the error.
5. Restart the Server.

Authentication Errors

Obtaining a Kerberos Ticket

If the Kerberos ticket is invalid or expired you may encounter an error “Unable to obtain Principal Name for authentication”.

Follow these steps to obtain and cache the Kerberos Ticket Granting Tickets (TGT):

6. Open the Command-line tool as an administrator.
7. Change directory to JDK bin path.

8. Type the following command:

```
kinit <username@Domain> and give your password.
```

9. Execute the "klist" command to get the ticket.

Once you have a valid Kerberos ticket, you can login to TDV Studio using the SSO option.

Using File Cache in Linux

In the Linux environment if the default cache setting is KEYRING, you may encounter SSO Authentication errors. Changing it to a File Cache in the configuration file (krb5.conf) will solve the issue. For example,

```
default_ccache_name = FILE:/tmp/krb5cc_%{uid}
```

Data Source Authentication Error

Data sources support only JGSS based Negotiate and Kerberos authentication. If you introspect a data source set up with SSPI Kerberos authentication, you get a 401 Authentication Error message.

SSO Authentication Failure

There can be many number of reasons when you get an SSO Authentication error. Running the Kerberos Test utility will help you identify the exact problem. Refer [Kerberos Test Utility](#) for instructions on how to run the utility.

Managing Security for TDV Resources

This topic documents several TDV security features which help you ensure that information is available only to authenticated, authorized individuals who have appropriate rights and privileges.

- [Overview of TDV Security Features](#)
- [Summary of Password Encryption and Security in TDV](#)
- [Summary of Internet Security Options](#)
- [Rights and Privileges](#)
- [Configuring Account Security for TDV](#)
- [Row-Based Security](#)
- [Encryption Settings for TDV Server](#)
- [Configuring Samba and Winbind for NTLM \(Tips from an Expert\)](#)

Overview of TDV Security Features

TDV provides many layers and types of security: password security, Internet security, domain/group/user security, row-based security, authentication protocols like Kerberos and NTLM. These security features are described in this topic and also elsewhere in the TDV documentation set as follows:

- [Summary of Password Encryption and Security in TDV](#)
- [Summary of Internet Security Options](#)
- Right and privileges to control levels of access to TDV resources are described in [Rights and Privileges](#).
- Row-based security lets you control access to specific rows of data and is described in [Assigning Users to TDV Groups or Identities](#).
- For information about setting up users and groups and their passwords, see [Composite Domain Administration](#)
- For information about authentication protocols, see:

- [Configuring Kerberos Single Sign-On](#)
- [Configuring NTLM Authentication](#)
- For descriptions of the security features available within and between TDV components, and between TDV components and their clients and data sources, see the *TDV Security Features Guide*.

Summary of Password Encryption and Security in TDV

This section summarizes TDV password security features for TDV components and data sources. For details, see the *TDV Security Features Guide*.

TDV encrypts the passwords used to access TDV components and external data sources. Specifically:

- Passwords sent by JDBC and ODBC to TDV are encrypted.
- Passwords passed between TDV components are encrypted.
- Passwords in metadata are encrypted.
- Passwords for LDAP and dynamic domain users are encrypted or not stored. Case-sensitive user sign-in for external LDAP is supported.
- The DBA password for the repository is not stored.
- Repository password and the repository connection with TDV are encrypted.
- Passwords are not shown in the log files.
- Passwords in HTTP SOAP headers for admin functions are encrypted.

Options are available to include or exclude encrypted user, repository, LDAP, and data source passwords in export files.

See [Changing Passwords for Other Composite Domain Users](#) for how to change a TDV password.

Summary of Internet Security Options

All communication between TDV and other TDV components through the Internet can be encrypted using SSL or HTTPS.

- Web Services security (WSS) Web service client security is supported.
- Pluggable Authentication Modules (PAM) can be deployed to mediate user/client sign-in. Authentication modules can be created, deployed, and enabled to secure access based on tokens, connections, physical assets, location, and biometrics. See [About Pluggable Authentication Modules](#).
- TDV to data source SSL is supported, with or without WS client authentication.

For information on SSL configuration with TDV, see [TDV and SSL Authentication](#).

Rights and Privileges

Two managed security layers—rights and privileges—ensure that only those with an appropriate security profile can access and manipulate TDV tools and native resources.

Default rights and privileges are as follows:

- The administrator has all rights and privileges.
- Everyone else has no rights or privileges. Rights and privileges must be assigned explicitly to groups or individual users.

These following sections describe setting up TDV resource rights and privileges for users and groups:

- Rights—Let you define access to tools and the ability to read or modify system-level characteristics. See [Resource Rights](#) for more information.
- Privileges—Let you control access to and manipulation of specific resources. See [Resource Privileges](#) for more information.

Resource Rights

Rights are security features that give groups and users the ability to perform TDV actions by letting them use associated tools and options. By default, no rights are given to any user except the administrator, who has rights to view and change everything in the TDV system.

This section covers the following topics:

- [Overview of Rights-Based Security](#)
- [Group and User Rights](#)
- [Installed Users and Groups and Their Rights](#)

Overview of Rights-Based Security

Rights-based security architecture creates a division of labor and TDV access management by functional group responsibilities, as described in [Group and User Rights](#). Users by default have no rights, because they access TDV through client connection rather than connecting directly to the server. For a description of these rights and the default groups to which they are assigned, see [Summary of TDV Rights and Restrictions](#).

The rights available on the TDV system are:

- ACCESS_TOOLS
- MODIFY_ALL_CONFIG
- MODIFY_ALL_RESOURCES
- MODIFY_ALL_STATUS
- MODIFY_ALL_USERS
- READ_ALL_CONFIG
- READ_ALL_RESOURCES
- READ_ALL_STATUS
- READ_ALL_USERS
- UNLOCK_RESOURCE

Group and User Rights

In the TDV system, rights determine which parts of TDV each user can access and use.

Rights are best specified at the group level. Because users automatically inherit all rights assigned to the groups to which they belong, we recommend that you manage enterprise rights at the group level. Role-based management can be more efficient than assigning rights individually.

Assign the Access Tools right to those user groups who should have access to Studio or other TDV components.

TDV does not introspect new LDAP domains to obtain lists of potential users. When you set up LDAP for TDV, use LDAP tools to choose the specific users and groups who are to have access to TDV.

Installed Users and Groups and Their Rights

The following default users and groups are created in TDV during installation. These users and groups cannot be removed from TDV.

- The "composite/admin" group is precreated with all rights. The "composite/admin" user is pre-created as a member of this group and cannot be removed from this group.
- The "composite/nobody" and "composite/system" users are pre-created with no rights and cannot be given rights or placed into groups.
- The "composite/all" and "dynamic/all" groups and the "composite/anonymous" user are pre-created with no rights. They can be granted rights, but we strongly recommend against doing so.

User and Group	Description
nobody user	Nobody is a special user who cannot be assigned rights or made a member of groups.
system user	System is a special user who cannot be assigned rights or made a member of groups.
composite/ anonymous user	The anonymous user is not a member of the all group, and does not inherit rights or privileges from that group. However, you can add rights and privileges for the anonymous user explicitly. The default TDV configuration setting does not allow anonymous users to sign in (TDV Server > Configuration > Security > Enable Anonymous sign-in: false).
composite/all group	The composite/all and dynamic/all groups are created during TDV

User and Group	Description
dynamic/all group	installation. They have no rights, and it is strongly recommended that no rights be given them, because this would give rights to all users without appropriate differentiation. All users that authenticate using a composite or LDAP domain and log into Studio are automatically members of the composite/all group.

Resource Privileges

Privileges determine which groups and users are able to view or act upon data from defined resources using the TDV suite of products. Privilege specification provides a comprehensive security layer to safeguard access to resources defined within TDV.

No default privileges are granted for newly defined resources, except to administrators, the resource owner, and users with the Modify All Resources right so that object ownership rights to grant privileges can be controlled solely by selected users.

Privileges can be assigned to an entire domain, selected groups, or individual users. Privileges can be set for any object exposed through TDV: containers (folders or parent objects) and individual resources, down to individual table columns. Privileges can be propagated to subordinate objects (child objects or dependent objects).

If you restrict access to a view in the published layer, the shared area, and in the introspected data source, the column has the same restrictions.

This section contains:

- [Initial Default Resource Privileges in Studio](#)
- [Resource Ownership and the Grant Privilege](#)
- [Assignment of Privileges](#)
- [Container and Resource Privileges](#)
- [Column-Level Restrictions on Privileges](#)
- [About Managing Dependency Privileges](#)
- [Setting and Viewing Privileges](#)
- [Propagation of Privileges](#)

- [Privileges for Non-Studio TDV Users](#)
- [Copying Privileges](#)
- [Finding and Editing Resource Privilege Dependencies](#)

Initial Default Resource Privileges in Studio

By default, every object resource defined in Studio is initially created with full privileges for the object creator. Except for administrative users and users with Modify All Resources rights, no other users are granted privileges on new resources unless those privileges are added later—added either to users or to the groups to which they belong.

Default Read privileges are given to members of the all group in the composite domain for all sample resources, system resources, and parent containers of those resources.

Any user who is part of an LDAP domain is automatically given access to the objects that belong to the all group in the composite domain.

By default, anonymous users and users in the dynamic domain are disabled in the TDV installation. They must be explicitly enabled.

The following is a summary of default privileges assigned to system resources.

Studio Tree Category	Default Resource Privileges		
	Group: All	Anonymous	Dynamic
localhost	Read		
services	Read		
databases	Read, Write	Read	Read
databases > system	Read		
databases > system > <table>	Select		
webservices	Read, Write	Read	Read
webservices > system	Read		

Studio Tree Category	Default Resource Privileges		
	Group: All	Anonymous	Dynamic
lib	Read		
lib > debug	Read		
lib > resources	Read		
lib > services	Read		
lib > services > <specific_service>	Read, Execute		
lib > sources	Read		
lib > users	Read		
lib > util	Read		
shared	Read, Write	Read	Read
shared > examples	Read, Execute, Select		
users	Read	None	None
users > composite	Read	None	None

Resource Ownership and the Grant Privilege

Each resource in TDV has an owner. The user who creates a resource is initially the default owner. An owner of a resource automatically has all privileges on that resource.

The resource owners can define privileges for groups and users who need to view, access, and use a resource. Privileges can be defined for a parent object in the TDV directory and they can be applied to child resources and subfolders recursively. Child object resources are not available or even visible to users who do not have Read privileges on all of the resource's parent containers. The owner or administrator has the option to grant or revoke privileges on the owned resource at any time. The owner can revoke his own privileges (for example, to prevent accidental deletion of data), and later

re-grant those privileges. A user who is given ownership of a resource can share all the privileges of ownership by giving the Grant privilege to other users or groups. Users who are not administrators or owners of the resource cannot change those privileges.

Administrators, users or groups with the Modify All Resources right can:

- See all resource definitions and associated privileges.
- Assign or remove privileges of groups and users.
- Change the owner of a resource from Studio's Administration menu.
- Change privileges on all resources that they have access to view, but they might not have access to read all resources.

Note: For resources that are likely to be called and invoked by other resources, you can give the Grant privilege to distribute access to other developers.

Assignment of Privileges

We recommend that you assign privileges by groups rather than by individual users. This style of access control lets future developers manage large numbers of users by adding them to or removing them from groups that combine easily understood sets of role-based privileges.

For LDAP domains, set the privileges for all members in the group at the same time, because individual members of the domain do not appear until they have logged into the system for the first time.

For the composite or dynamic domain, you can increase control by creating additional groups to manage subsets of rights and privileges from the Manager. For the LDAP domain, you must use the LDAP tools to modify groups and their rights.

The following shows a typical privileges dialog for a resource with both implicit privileges (those assigned by group definitions) and explicit privileges (those assigned directly to a user or group). Privileges assigned explicitly are shown by a green check mark. An amber check mark shows privileges assigned implicitly (acquired by either group membership or possession of a right like Modify All Resources).

Container and Resource Privileges

Privileges fall into two groups, design-time and run-time. Read privilege belongs to both groups, and must be granted to users for all enclosing containers (parent or folder objects) of the resource the user wants to access or manipulate.

Read, Write, and Grant are design-time privileges that determine access to objects for users of TDV and other TDV components.

Read, Execute, Select, Insert, Update, and Delete are run-time, data manipulation privileges that determine resource security for client-interface access to data through TDV.

The following tables describe what can and cannot be done with different combinations of privileges. In the table, an X means that the privilege is granted on the resource, and N/A means that the privilege is not applicable to the resource.

Privileges and Resources								What User Can Do with Resource or Container
Read	Write	Execute	Select	Update	Insert	Delete	Grant	
All resources								Not view, modify, or add resources to it.
X								View but not modify it.
X	X							View, modify, delete, move, reconfigure, rename, or create resources inside of it.
Folders, data sources, published databases, catalogs, schemas, TDV Web services (and								

Privileges and Resources								What User Can Do with Resource or Container
Read	Write	Execute	Select	Update	Insert	Delete	Grant	
their services, operations, and ports)								
X	X	X	X	X	X	X	X	View, modify, query, delete, move, reconfigure, rename, and execute.
Tables, source privileges (second tab for published database resources), views								
X	X	N/A	X	X	X	X	X	View, modify, delete, move, rename, reconfigure, or query it.
Columns								
X	X	N/A	X	X	N/A	N/A	X	View, modify, run select, run update.
SQL Script procedures, Java procedures, packaged queries, XSLT and XQuery procedures, parameterized queries, transformations								

Privileges and Resources								What User Can Do with Resource or Container
Read	Write	Execute	Select	Update	Insert	Delete	Grant	
X	X	X	N/A	N/A	N/A	N/A	X	View, modify, delete, move, reconfigure, rename, and execute.
Published database resources (which have a second tab for source privileges), Web service definitions, triggers, definition sets, models								
X	X	N/A	N/A	N/A	N/A	N/A	X	

Column-Level Restrictions on Privileges

Column-level restrictions set in TDV (Table or View) causes the restricted column and metadata to not appear through the client interface, and an access error is generated.

About Managing Dependency Privileges

Managing privilege settings for resources can become complex when many data resources contribute to the output of a resource. When resources use another resource, every prospective user and group must have adequate privileges to access and invoke those dependencies.

Even if a Studio developer has the privilege of using and viewing a dependency resource, that person might not be able to assign privileges on resource dependencies to other users and groups.

At run time, when using resources contained within other resources, the Read privilege must be present in all parent containers.

Privileges Required for Resource Dependencies

When a resource has a dependency, the user who requests the view or invokes the procedure must have privileges to view the dependency, and to Select or Execute to retrieve data from the dependent resource.

Note: You can use the lineage feature in Studio to learn about the resource relationships, dependencies, and references. See “Exploring Data Lineage” in the *TDV User Guide*.

For example, a user might want to see the data of View_A, where View_A executes Procedure_B (which draws on data from physical source Table_D) and also selects from physical source Table_C. To perform a SELECT on View_A, the user needs these privileges.

On the Resource	Privilege Required
All parent containers of View_A, Procedure_B, Table_C, Table_D	Read
View_A, Table_C, Table_D	Select
Procedure_B	Execute

Setting and Viewing Privileges

Privileges can be assigned to selected groups or individual users. Privileges can be set for an entire folder, including all child objects. Privileges let you restrict access and changes to data down to individual table columns.

When a privilege is granted to a user, the corresponding privilege check box is selected in the panel. See [Container and Resource Privileges](#).

- Select is offered, but not Execute, if the resource is a view. Execute is offered, but not Select, if the resource is a procedure.
- Studio presents all privileges when you edit the privileges on a container, even though only Read and Write are relevant to a container.
- The Properties window for a published resource has a tab named Privileges (for the published object) and another tab named Source Privileges (for the corresponding unpublished object). Only the unpublished object can be given the Execute privilege.

Individual LDAP users can be directly assigned rights and privileges for any resource, but only after they appear in the TDV system.

When designing client interface applications, you might experience difficulties with resource availability and access privileges if you do not have the necessary Read privileges. You must have Read privileges for all levels of the resource tree to enable clients to view and use a resource. Developers, resource owners, and administrators should work together to make sure that appropriate groups and users can both view and access contained resources.

To set privileges on a resource

1. Make sure you have one of the following privilege conditions:

- You own the resource.
- You have the Grant privilege on the resource.
- You have the Modify All Resources right.

See [Resource Ownership and the Grant Privilege](#) for more information.

2. Right-click the resource name and select Privileges, or select the resource and choose Resource > Privileges. For example, when selecting privileges from a data source resource you would see the following screen.

The Resource and Your Privileges portions of the panel are for information purposes only.

3. Select the Show All radio button, to display or to set privileges explicitly for users who are not shown.

For boxes with an amber check mark, you might not be allowed to change the privilege setting. Privileges explicitly assigned are shown by a green check mark. An amber check shows privileges obtained implicitly (by group membership or through possession of a right like Modify All Resources).

4. Select the Read and Write boxes to set Studio design-time privileges. The Read privilege on a resource lets the user check whether the resource exists. The Write privilege lets the user modify the TDV resource definition, which determines what native resource can be used, and how it can be used.

Make sure that groups and users:

- Who are given new privileges on a resource also have the Read privilege on all parent containers in the path of that resource.

- Who have the Read privilege can view the object design, projections, schema, SQL, and annotations.
 - Who have the Read and Select privileges can view dependencies, view the execution plan, and run the SQL contained in the view.
 - Who have the Read, Write, and Select privileges can save, cache, and publish the view.
 - Who need to modify an existing resource definition have the Write privilege on that resource.
 - Have the Read privilege on views or views used to build another view.
5. Select the Execute, Select, Update, Insert, or Delete boxes, to set run-time privileges.
 - The Select privilege lets the user submit SQL SELECT statements to retrieve data.
 - The Execute privilege lets the user execute a procedure.
 - The Insert, Update, and Delete privileges let the user change table data.

At run time, Read privileges are used for folders and their contents, but not for tables or procedures. However, a user with Select privileges on a table but no Read privileges can still select from that table.
 6. Check the Grant box for resources where you want to share ownership privileges. The Grant privilege gives other Studio users the same privileges as the original resource owner.
 7. Use the following fields and buttons to filter the users and groups that appear on the Privileges panel.

Field or Button	Description
Hide users without explicit privileges	Default. Does not show users who have not been granted privileges on the resource from the user interface window.
Hide users and groups without explicit privileges	Does not show users or groups who have not been granted privileges on the resource from the user interface window.

Field or Button	Description
Show All	Show all users and groups, whether or not they have explicit privileges.
Filter	Lets you type a filter string to apply to user and group names. A wildcard is added before and after your string; so for example “ea” finds users Jean and Bea, and group Minneapolis.

8. Select the Apply recursively to dependencies check box to apply the setting selections to all resources that this resource *depends on*.
Resources that this resource depends on can reside anywhere.
9. Select the Apply recursively to dependents check box to apply the setting selections to all resources that *depend on* this resource.
Dependent resources can reside anywhere.
10. Select the Apply recursively to child resources and folders check box to apply the setting selections to all the child resources that are *owned by* this parent resource.
Child resources reside within the parent container.
Note: Many resources, such as procedures, queries and transformations, do not make Apply recursively to child resources and folders available for selection.
11. Select one of the radio buttons either to apply changes to privilege settings, or apply changes to privilege settings *and* clone user and group privileges to all child and/or dependent objects of the selected resource.

Radio Button	Description
Only apply modification	Saves the specific changes made to the privilege settings of the selected resource, but preserves all other privilege settings of the child and/or dependent resources.
Make child resources look like this resource	Saves changes made to the privilege settings of the selected resource, and propagates all of the object’s privilege settings to child and/or dependent resources. Changes applied using this radio button can both <i>add</i> and <i>remove</i> child and/or dependent resource privileges.

12. Click OK.

For information on how to grant privileges to a source on which a resource depends, see [Privileges Required for Resource Dependencies](#).

Propagation of Privileges

The privileges Insert, Update, Delete only need to be set for the top-level view—the view that is published to a virtual database. These privileges do not need to be set separately for other views in the dependency lineage.

Users who have Select, Insert, Update, and Delete privileges on a view, but do *not* have Read access on that view, can see the view when signed in to Studio, but if they try to open the view, an error message is displayed. After clicking OK in the error message box, they can open and execute the view in Studio, but they cannot save any changes to the view.

Implicit Container Privilege Assignment

READ privileges are automatically assigned on parent and ancestor Containers when SELECT, INSERT, UPDATE, DELETE or EXECUTE privileges are detected as assigned on a given resource.

The same implicit assignments are applied when privilege assignment is propagated on dependencies and/or dependents.

This behavior can be tuned by editing the configuration setting Administration - > Configuration -> Server -> Configuration -> Security -> Authentication -> Implicit Container Privilege Assignment. The setting controls whether missing READ privileges will be automatically assigned on parent and ancestor containers of resources whose privileges are being set to include any of SELECT, INSERT, UPDATE, DELETE or EXECUTE access. By default this is set to TRUE.

Privileges for Non-Studio TDV Users

This topic describes privileges for users who do not use Studio to connect to TDV. One category of such users is those who connect to TDV through JDBC, ODBC, or ADO.NET. Insert, Update and Delete privileges are of course needed for such users to take those actions, but other conditions for these actions may not be as obvious:

- A published resource and its lineage of views and tables do not require the Read privilege.
- The published resource must have the Select privilege.

- *Intermediate* views and tables do not require Insert, Update and Delete privileges for the user to take such actions.

Copying Privileges

Users or administrators who have the Grant privilege on resources can copy privilege settings from a selected resource to one or more other resources.

To copy privileges

1. Right-click the resource whose privileges you want to replicate.
2. Select Copy Privilege.
3. Select one or more target objects to which to copy the privileges.
4. Select target resources carefully. All privilege settings of the initial resource overwrite all of the privilege setting of the selected target resources.
5. Select the Copy Privileges into Target Descendants, if you want to copy privileges to the descendants of the target objects.
6. Click OK.

Finding and Editing Resource Privilege Dependencies

The dependency privilege analysis checks users and groups against all resources defined by the TDV Server. Specifically, it checks whether those with Select or Execute privileges also have the appropriate privileges to access and use dependency resources when they are present. The Dependency column displays an aggregate status icon that tells whether any dependency privilege settings need review and correction.

Administrative users with the Modify All Resources right can click in any box and thereby assign a privilege to all selected users and groups. Clicking on any box again sets the privilege back to its former state. You can click the ADD PERMISSIONS button to assign an explicit privilege to correct all resource dependency privilege deficiencies.

If the administrator currently using Manager does not have the Modify All Resources right, privilege changes are possible only on those resources and containers for which the administrator is the owner or has the Grant privilege.

The Manager Resources pages provide tools that automatically analyze privilege sets on resources and their dependencies for groups and users. TDV administrators can add

edit dependency privileges, and view, edit, or remove privileges by resource, user, or group.

The Manager Resources pages let you do the following:

- Find missing privileges by resource for any group or user.
- Check for privilege inconsistencies on dependent resources.
- Add or remove privileges on a resource or dependency.
- Find and assign privileges for new dependencies on a revised resource.
- Perform a resource-based privilege security audit.

The Manager Resources pages are best used with the following administrative rights:

- Access Tools—To launch and use the Manager.
- Read All Config—To see the Manager resource page.
- Read All Users—To see privileges beyond those for one's own sign-in ID and groups.
- Modify All Resources—To add, remove, or automatically correct dependency privilege settings (except for resources for which a user is the owner or has the Grant privilege).

Note: Manager displayed in IE8 compatibility mode is slow. Also, IE8 warns that the JavaScript can cause the computer to become unresponsive. It is safe to ignore this warning. When using IE8, change the Refresh Rate.

To find resource dependencies

1. Launch Manager using:
 - Launch Manager option in the Studio Administration menu
 - <http://<hostname>:9400/manager/login>
2. Select Resource Management from the SECURITY tab menu, to open the RESOURCES page.
 - The RESOURCES page displays published, user-accessible folders and resources.
3. Select a radio button and click Analyze Dependency Privileges (radio button in the

first column shows which is selected).

4. The Dependency column displays:

Icon	Indicates
question mark	Dependency analysis has not been performed.
clock	Analysis is in progress.
check mark	All users and groups with invocation privileges for the resource also have invocation privileges for all of the resource dependencies (where present).
red circle x	One or more users or groups with invocation privileges for this resource are lacking adequate invocation privileges on one or more dependencies. Any resource with an inconsistent privilege assignment has a user or group with a Select or Execute privilege allowing invocation of the parent resource, but without the required privileges to access/use dependencies required for execution.
yellow triangle	One or more users or groups with partial privileges for this resource lack adequate privileges to invoke it or dependent resources. Where privilege settings do not match normal usage patterns, the dependency analysis marks them with a yellow warning triangle—for example, a user with the Read or Write privilege but not a privilege to Select or Execute on that resource or any of its dependencies.

Analyze Dependency Privileges analyzes all resources on the page for privilege settings.

5. Click View Privileges to display a detailed resource privilege report that you can use to check resource privilege settings.

The View Privileges button displays a “RESOURCE PRIVILEGES for <resource>” page containing up to 100 users and groups with their privileges, and a dependency analysis status indicator for each user and group with a privilege on the selected resource. This display is suitable for a resource-based security audit that shows who does and who does not have access to a given resource.

6. Optionally, modify the Refresh Rate near the upper right of the Manager home page, to control how often information on the page is refreshed.
7. Select the check box for a row to edit, remove or modify the privileges. The following actions can be performed.

Action	Description
Edit Privileges	Directly modify privileges on the selected resource for one or more users and groups. When you select more than one user or group, privilege settings show an aggregated value of the privilege settings.
Edit Dependency Privileges	Analyze dependency resources (and their parent containers) and required privileges for access and invocation. Also opens the Edit Resource Dependencies window, which displays more detail about aggregated privilege settings. Select those users and groups that show dependency privilege errors or inconsistencies.
Remove Privileges	Remove explicitly-assigned privileges from selected users or groups. This function cannot remove users from groups that give the user privileges implicitly, or remove administrative rights to remove privileges. Implicit privileges can be removed only by removing the user from the group that gives those privileges, by changing the group assigned privileges, or by removing user rights.
Reset the Resources pages	Click the SECURITY tab > Resource Management selection.

To correct a privilege setting

1. Select SECURITY > Resource Management.
2. Click Analyze Dependency Privileges.
3. Click View Privileges.
4. Click Edit Dependency Privileges.
5. Click any of the icons to change privileges for all the users and groups selected.

6. Click Add Permissions.
7. Add Permissions identifies all privilege deficiencies, and explicitly grants all missing privileges needed to access and invoke a resource and its dependencies.
8. Click OK or Cancel when done.

Configuring Account Security for TDV

Account security is an important enterprise tool to provide secure access to applications and data across your enterprise. This section includes the following topics:

- [Configuring Account Lockout for TDV](#)
- [Locking and Unlocking TDV User Accounts](#)
- [Setting IP Restrictions](#)

Configuring Account Lockout for TDV

TDV allows you to control the number of password failures that a user has before their account is denied access to TDV. You can set the number of invalid log on attempts that are allowed before an account is locked out. Choose a value that prevents valid users from routinely getting locked out of the system and deters invalid users from accessing the system.

If you use TDV in a cluster environment, TDV records the number of failed attempts across all server nodes in the cluster.

Lower threshold numbers might require more frequent intervention to unlock valid accounts. Higher threshold number make systems easier for invalid users to access.

To configure account lookout thresholds

1. Sign-in to Studio as a user with the Modify All Config privilege. For example, the Admin user typically has the modify all config privilege.
2. Select Administration > Configuration.
3. Navigate to the Implicit Lock Threshold parameter.
4. Type a value that prevents valid users from routinely getting locked out of the system and deters invalid users from accessing the system. For example, 5.

5. Click OK.

After an account lockout threshold has been set, user accounts are locked immediately after executing the specified number of invalid sign-in attempts.

Locking and Unlocking TDV User Accounts

With the ability to automatically lock user accounts comes the need to unlock accounts. TDV provides the ability to lock and unlock user accounts through the Web Manager interface.

To immediately lock or unlock a TDV user account

1. Sign-in to Studio as a user with the Modify All Users and Access Tools privileges. For example, the Admin user typically has the modify all users privilege.
2. Select Administration > Launch Manager.
3. Type the password, if you are prompted.
4. Select the SECURITY tab and User Management.
5. In the Locked column, click the padlock icon to immediately lock or unlock a user account.

Changes take effect immediately.

Setting IP Restrictions

You can configure TDV to grant or deny specific computers, groups of computers, or domains access to TDV. IP restrictions can be important when you allow employees to remotely sign-in to the TDV system. Remote access can be exploited by invalid system users. Preventing invalid users from accessing your TDV system can help keep you and your customers secure from hostile use of your system data.

If a system attempting access is on the list of forbidden IP addresses, users would receive an authentication failure message, but would not be notified that they are on a list of forbidden IP addresses.

To configure IP restrictions for TDV

1. Sign-in to Studio as a user with the Modify All Config privilege. For example, the Admin user typically has this privilege.
2. Select Administration > Configuration.
3. Navigate to the Forbidden IP Addresses or Required IP Addresses parameter in the hierarchy.
4. To add values to the lists, use the green plus-icon to add a new editable text field, and type a value for the IP address that you want to allow or deny.
 - Asterisks (*) can be used as wildcard characters.
 - *.*.*.* is a valid value string, but TDV ignores it.
5. Click OK.

Row-Based Security

This section contains:

- [About Row-Based Security](#)
- [Enabling Row-Based Security on TDV Resources](#)
- [Creating or Editing Row Filter Policies](#)
- [Creating or Editing a Row Filter Policy Group](#)
- [Assigning Row Filter Policies to TDV Resources](#)
- [Testing the Security Filter Policies Within Studio](#)

About Row-Based Security

Control of access to sensitive data is a fundamental part of TDV. Row-Based security is a form of access control that can be applied to specific rows. Row-Based security prohibits access to restricted rows. Typically, row-based security is controlled by defining different data viewing permissions based on user and group.

TDV has a guided method to help you create a row filter policy that is based on user or group. TDV also provides access to the SQL script that drives the row filter policy. You can create and modify the associated SQL script files to define a custom row filter policy.

For example, the salaries for employees or the locations of specific assets, might be data that you want to keep restricted from certain groups or users. Or, you could restrict access based on time of day or region.

If you are using row-based security policies in combination with column-based security policies, TDV applies the row-based policy before applying the column-based policy.

Row Filter Policies

A row filter policy is a SQL script that is used to define how you want the data access controlled. The two major row filter policy categories in TDV are:

Categories	Description
Tabular	TDV user interface guided method for creation of policies for row-based security. By using the default, Tabular option, TDV creates all the SQL scripting necessary to define the policy that is based on user or group and containing rules.
Free-form	Customized SQL script written by you.
Group	This category allows you to assign multiple policies to any number of resources.

Regardless of your row filter policy definition, the table or view that it is assigned to will never provide more data than it would without the row filter policy applied to it.

Rule

If you are defining tabular row filter policies, you can use TDV to define multiple rules. Each tabular row filter policy can contain one or more rules. Each rule is associated with a specific user or group that is defined within TDV. When defining rules, you can choose:

- All rows
- No rows
- Selected rows as defined by a predicate
- Selected rows as defined by another filter procedure

Assignments

Assignments are used to associate particular row filter policies to a table or view resource. A single row filter policy can be applied to one or more tables or views. After row filter policies are defined and applied to resources, data viewing can be controlled based on the criteria of the row filter policy.


Behavior of the WHERE Clause

The row filter policy that you define does not alter the SQL code of your view or table. The row filter policy only temporarily adds to the WHERE clause of SQL statements that consume the table or view.

Row filter policies are added at run time using the AND operator. Rules defined are added together at runtime using an OR operator.

Row Filter Policies and Caching Restrictions

You can define row filter policies and apply them to views that are part of your TDV data caching environment. However, if a view definition uses the results of a table or view that has a row filter policy assigned to it, the cache cannot be created.



Within Studio, if you have views with a row filter policy and caching enabled and you open the cache associated with that view, you can see all the data in the view as if the row filter policy had not been applied. For example, if you have a row filter policy and cache defined on ORDERS, when you view data for ORDERS you will see the data as constrained by your row filter policy. However, if you view the data for ORDERS_CACHE, you will see all the data.


Behavior of Row Filter Policies with Group Row Filter Policies

Because it is possible to define individual row filter policies and group row filter policies and assign them to the same resources, you could run into some behavior that might not seem predictable.

In the case of 2 policies assigned to specific resources that are also used in a group and assigned to the same resources, TDV would apply the policies using an algorithm like, (p1 and p2) or (p1 and p2).

To add a further example, given the V_SALES view with region, product and customer data and polices set up for REGION, PRODUCT, and CUSTOMER with filters in each to segment each data set based on ID. For V_SALES, if a group row filter policy is created that combines each policy into one and then assign the group to it.

If the ALL selection is chosen, the three policies are combined in the SQL with AND operators.



If the ANY selection is chosen, the three policies are combined in the SQL with OR operators.

For more information on how to define groups, see [Creating or Editing a Row Filter Policy Group](#).

Enabling Row-Based Security on TDV Resources

The row-based security for TDV requires that the TDV system-wide option be enabled through Manager.

Note: Manager displayed in IE8 compatibility mode is slow. Also, IE8 warns that the JavaScript can cause the computer to become unresponsive. It is safe to ignore this warning. When using IE8, change the Refresh Rate.

To enable row-based security

1. Make sure that you are signed in as an Administrative user with the Modify All Resources right.
2. Launch Manager using:
 - Ctrl+M from Studio
 - Launch Manager option in the Studio Administration menu
3. Sign-in as a user an Administrative user with the Modify All Resources right.
4. Select Row-Based Security from the SECURITY tab menu, to open the ROW-BASED SECURITY page.
5. Click Change Enabling to toggle between Enable and Disable for the Row-Based Security field.

On this page you can view all the row filter policies that have been created. You can add new row filter policy, delete existing row filter policy, or edit an existing row filter policy this page. For each row filter policy listed, there are Assignments.

6. To define policies, see [Creating or Editing Row Filter Policies](#) .
7. To create policy groups, see [Creating or Editing a Row Filter Policy Group](#).
8. To assign policies to resources, see [Assigning Row Filter Policies to TDV Resources](#).

Creating or Editing Row Filter Policies

You can use the TDV user interface to lead you through the creation of policies for row-based security. By selecting the default, Tabular option, TDV takes care of creating all the SQL scripting necessary to define the policy.

Column names referenced in the filter expression are not validated. You must test each row filter policy that you define to make sure that it works as expected. For example, if a filter is on REGION_NAME in the LOCATION table and the column is changed to REGION_CODE, the filter must be updated to reflect that change.

The SQL script for free-form row filter policies can be edited using Studio. All other definitions, editing, or deleting of row-based security objects must be done using Manager.

Note: There is no ranking or priority of the rules within a row filter policy. If a user or group is assigned to more than one rule within a row filter policy, the two predicates are combined using OR statements.

To add or edit a row-based security policy

1. Follow the instructions in [Row-Based Security](#).
2. To add a new row filter policy, click Add Policy. Or, to edit an existing row filter policy, select it and click Edit Policy.
3. Select or specify values for the following fields:

Field	Specify
Name	Specifies the name you want to give to the row filter policy. Spaces are not allowed.
Folder	Specify the location within /shared for the SQL Script procedure in the Folder field when adding a new row filter policy. If you are editing an existing row filter policy, the Folder field is read only. The default value for this field is /shared.
Enabling	• Enable—Allow the use of the specific row filter policy. • Disable—Disallow the use of the specific row filter policy. Regardless of this setting, you can still add and edit the row filter policy. If the policy is disabled, it is not used against the data, even if the policy is part of a group. For example, if an RBS_group contains three policies (pig, hive, and gnu) and the gnu policy is disabled, when the RBS_group is applied only the pig and hive policies are used.
Specification	• Tabular—Select to have TDV create the SQL script for you. • Free-Form—Select to create the SQL script on your own. • Group—Select to create a container that can include several row filter policies.

Field	Specify
Description	(Optional) Type an explanation of the row filter policy.

If you selected Free-Form, skip to Step 7. If you selected tabular, the screen displays additional fields and buttons.

- If you selected Tabular, you can add or edit a rule for a particular user or group using the Manager screens. Click Add or Edit, and then edit or specify values for the fields listed in the table below.

Field	Specify
User/Group	Specifies whether this identity is for a Group or a User.
Domain	Specifies the Domain to which you want the row filter policy to apply.
Name	Specifies the user or group name to which you want the row filter policy to apply.
Rule	Specify All, None, Predicate, or Procedure.
Data/Path	If you selected Predicate or Procedure, then specify the string used to filter the data. For example, Region='EUR'.

- Click OK.

- If you selected Tabular, you can specify values for the following fields:

Field	Specify
Default Rule	Specify rules to be carried out if an identity is matched, or a default rule to be used when no identity is matched. All Rows—Allow all rows to be returned. No Rows—Do not allow any rows to be returned. Predicate—Return only rows satisfying an explicit SQL predicate. A predicate is a part of a WHERE clause.

Field	Specify
	Procedure—Return only rows satisfying a predicate returned by a different SQL script procedure.
Data/Path	If you selected Predicate or Procedure, then specify the string used to filter the data, or a SQL script procedure that returns a valid filter expression (typically this is a WHERE-clause fragment). The SQL Script procedure must be a valid row filter procedure. For example, Region='EUR'. Or you can input the path to a procedure.

7. Click OK.

8. To edit the SQL script associated with the custom (free-form) row filter policy:
Open Studio.

Navigate to the /shared folder and then to the SQL script resource object that has the name of the custom row filter policy you created.

Edit the SQL script procedure. For row filter policies defined using the tabular method, you can view the SQL script, but you cannot edit it. For information on how to edit the script, see “Creating a SQL Script” in the *TDV User Guide*.

The following is a sample SQL script that defines a row filter policy:

```
PROCEDURE "policy1" (IN alias VARCHAR, OUT result VARCHAR)
```

```
BEGIN
```

```
DECLARE temp VARCHAR;
```

```
DECLARE test BOOLEAN;
```

```
SET result = ";
```

```
CALL /lib/users/TestUserIdentity('USER','admin','composite',test);
```

```
IF(test) THEN
```

```
    SET result = '(' || alias || '.OrderID = 20)';
```

```
END IF;
```

```
    CALL /lib/users/TestUserIdentity('USER','test','composite',test);
```

```
IF(test) THEN
```

```
    SET temp = '(' || alias || '.OrderID < 20)';
```

```
    IF (result = "") THEN
```

```
        SET result = temp;
```

```
    ELSE
```

```
        SET result = result || ' OR ' || temp;
```

```
    END IF;
```

```
END IF;
```

```
    IF ( result = "") THEN
```

```
        SET result = 'FALSE';
```

```
    END IF;
```

END

Creating or Editing a Row Filter Policy Group

Row filter policy groups allow multiple policies to be assigned to one or more resources. It is an organizational tool that you can use to save time when assigning multiple row filter policies to one or more resources.

The definition of the row filter policy groups provides few constraints so that you can design the policies as you need to. Because of that flexibility, you must perform testing of your policies. TDV will attempt to detect obvious errors in policy behavior, but, TDV cannot detect conflicts between multiple policy definitions or logic errors of policies that are defined by dynamic variables.

To add or edit a row-based security policy group

1. Follow the instructions in [Row-Based Security](#).
2. To add a new row filter policy group, click Add Policy. Or, to edit an existing row filter policy, select the policy group and click Edit Policy.
3. Select or specify values for the following fields.

Field	Specify
Name	Specifies the name you want to give to the row filter policy group. Spaces are not allowed.
Folder	Specify the location within /shared for the policy when adding a new row filter policy group. If you are editing an existing row filter policy group, the Folder field is read only.
Enabling	• Enable—Allow the use of the specific row filter policy group. • Disable—Disallow the use of the specific row filter policy group. Regardless of this setting, you can still add and edit the row filter policy group.

Field	Specify
	If the policy is disabled, it is not used against the data, even if the policy is part of a group. For example, if an RBS_group contains three policies (pig, hive, and gnu) and the gnu policy is disabled, when the RBS_group is applied only the pig and hive policies are used.
Specification	Group—Select this to add a row filter policy group.
Description	(Optional) Type an explanation of the row filter policy group.

The screen displays additional fields and buttons.

4. Click Add Policy.
5. Select or type the full path and name of the policy you want to add to the group.
6. Click OK.
7. Select or type the full path and name of the policy you want to add to the group.
8. Click OK.
9. Continue adding policies to the group until it is defined as you want.
10. Specify a value for the Group Requirement field.

Value	Description
Row data must satisfy ALL policies listed above	Combines all of the policies in the group with AND operators. Use this setting to display a record if both the all conditions are met.
Row data might satisfy ANY policy listed above	Combines all of the policies in the group with OR operators. Use this setting to display a record if any of the conditions are true.

11. Optionally, type a description of the policy group.
12. Click OK.

Assigning Row Filter Policies to TDV Resources

For each resource within TDV, you can assign a specific row filter policy. This procedure assumes that row filter policies have been defined that you can assign to the resources. If you need to define row filter policies, see [Creating or Editing Row Filter Policies](#).

Note: Long lists in Manager and online help might not display as expected in Chrome. You can switch to another browser to resolve the issue.

To assign row filter policies to resources

1. Follow the instructions in [Row-Based Security](#).
2. Click Assignments on the ROW-BASED SECURITY page.
3. Click Add or Edit.
4. Type or select the full-Studio navigation tree path and name of the resource to which you want to assign the row filter policy. You can use the navigation on the window to browse and locate the object to which you want to assign the row filter policy.
5. Click Delete to remove the row filter policy association from the resource.
6. Click OK to save your changes.

Testing the Security Filter Policies Within Studio

TDV provides the ability to test the row filter policy within Studio, but recommends that you perform thorough testing of your row filter policies by setting up a working test environment that includes all the data sources and client applications that would be interacting with the data for which you have defined row-based security.

Within Studio, the Test Identity tab lets you execute a view and simulate the results according to row filter policies that you have defined. You can test the row filter policy and the resulting data set by executing the view as different users and groups. After using Studio to test your row filter policies, you are also advised to run tests of your TDV systems before going to production with row-based security.

The test on this tab changes how row filter policies, based on identity, alter SQL statements. For security reasons, the Test Identity tab cannot change data source connections, so the candidate rows that might be filtered are the same candidate rows

that the administrator would see. Use of this tab does not alter any pass-through credentials.

To test your row filter policy within Studio

1. Make sure that you are signed in as an Administrative user with the Modify All Resources right.
2. Identify any existing table or view resources for which you want to test row-based security.
3. Create a view of those table or view resources so that you can apply filters to them.

For example, to filter the data in the EMP view, create an EMP_VIEW_FILTER view with the definition of `SELECT * FROM EMP`. Define row-based security on the EMP_VIEW_FILTER view.
4. Select and open the view of the resource for which you want to test row filter policies.
5. Select the Test Identity tab.
6. Select one of the following.

Option	Description
Current Identity	Run the view and retrieve a result set filtered according to the row filter policy, as the user currently signed in to Studio.
Simulated Identity	Run the view and retrieve a result set filtered according to the row filter policy, as the specified user or group. Type the User@Domain or Group@Domain values. For example, admin@composite or QA@finance.

7. Click Execute.
8. Review the data returned in the Result portion of the screen and determine if the row filter policy is displaying the rows as you expected for your row filter policy definition.

Encryption Settings for TDV Server

TDV uses a symmetric key to encrypt credentials and other sensitive data that is stored in the server databases. During the installation of TDV, the following property files are created in the specified location, with the encryption settings.

1. **boot.properties** - <InstallDir>/conf/
2. **encryption.properties** - <InstallDir>/conf/server/

It is important to note that, since the property files contain the encryption key information, it should not be exposed to end users. During TDV installation process, it sets a restricted permission to the files for this reason.

The files will be in the format specified below:

boot.properties

```
server.encryption.userKey=
```

encryption.properties

```
server.encryption.algorithm=AES
```

```
server.encryption.uuid=
```

```
server.encryption.keySize=128
```

```
server.encryption.userKey=
```

- The encryption algorithm used is AES.
- The default keySize is set to 128 bits.
- The uuid and the userKey values are generated for every instance of the TDV Server.

After the installation of TDV, manually backup these files to avoid loss of data during unforeseen system issues. Note that TDV support will need this backup (along with the password you used while creating the backup) in order to help with the issue.

Defining or Editing Encryption to Protect TDV Server Data

You can use the TDV Manager web interface to lead you through the creation of encryption for your TDV Server. TDV uses a symmetric key to encrypt credentials and other sensitive data that is stored in the server databases.

The Encryption page of Manager can be used to:

- Set encryption of your TDV Server (this topic)
- Manage encryption settings by exporting or importing an encryption settings file. (see [Exporting or Importing an Encryption Settings File From Manager](#))

To add or edit encryption

1. From Studio, select Administration > Launch Manager.
2. In Manager, select SECURITY > Encryption.
3. Select or specify values for the following fields:

Control	
Encryption Algorithm	Value: AES AES (Advances Encryption Standard) is the industry standard encryption algorithm and replaces the TEAV algorithm entirely.
Encryption Password	Type a password with a minimum of 6 characters. Spaces are not allowed. Or, click Generate to have TDV create a unique string.
Unique Identifier	Type a unique identifier for the server export with a minimum of 6 characters. Spaces are not allowed. Or, click Generate to have TDV create a unique string.
Encryption Key Size	Larger sizes imply a longer encryption key and therefore stronger encryption. • 128 bits

Control

- 192 bits
- 256 bits

4. Click Save Configuration.

This saves and uses the password and unique ID to encrypt the current running TDV Server and all the data stored there.

Backup of Encryption Settings

In some cases, system related issues such as insufficient disk space, intermittent disk reading failure, etc. may cause TDV server to shut down unexpectedly. During such instances, check the details of the error message and rule out any hardware, software or environment related issues. If the problem persists, contact TDV support.

It is recommended that users backup encryption key periodically. This will help restore data during unforeseen circumstances. There are two ways to backup the encryption settings:

1. **Backup through web manager** - Refer to [Exporting or Importing an Encryption Settings File From Manager](#) for instructions on how to backup the encryption settings using Web manager.
2. **Manual backup** - After the installation of TDV, manually backup the files boot.properties and encryption.properties to avoid loss of data during unforeseen system issues.

Note that TDV support will need this backup (along with the password you used while creating the backup) in order to help with the issue.

Exporting or Importing an Encryption Settings File From Manager

You can manage encryption settings by exporting or importing an encryption settings file. The Encryption page of Manager can be used to:

- Set encryption of your TDV Server
- Manage encryption settings by exporting or importing an encryption settings file.

To export or import a password protected CAR file

1. From Studio, select Administration > Launch Manager.
2. In Manager, select SECURITY > Encryption.
3. Select or specify values for the following fields:

-
E x p o r t	4. Click the Export button to save your encryption settings in the form of an encrypted text file. 5. Type a new password with 6 or more characters, and type the password again to confirm it. An encrypted server back up file named backup_encryption_settings.txt is created.

Contents of the file are scrambled, for example:

```
46a4f7c7-470e-4415-a5e8-6762d2d2eef1$$$ENC
(yByoTeUNmlJum6jHgBdYZA== $KC0/yKi46q4C8LhT5YsNWDdKODJB/LzfVD/GoG
0awEFE8jRmPi1JzOWAcBEAeNjsyo6TgZ56AOjBtWBmgMQBFtoA+0CnaIX1WqX6
whILxj6STd2QnIYbR2IkugMo0zonUi5+UubBC5Mj5bLR6c+2/cTWI3d5/iohPDflqmMk
InvMkV/XaquoQDzasmVFDNd5d7oQ9UpK3rzSKlqceOjFjgCQYnP54cPq8RduwH0
W/TFblgjEY2EnEYWwfGiEuSOgZaegkmueO8QiswYw9Ou0nNmGbowu3pJ5v/X2ED
YeJiCZqBtRxCOYMobgD7z9w+5c9F1qlhqxhuBaxhnLYMdg==)
```

-
I m p o r t	6. Click the Import button to import your encryption settings file. 7. Browse to the encrypted settings text file. Typically named backup_encryption_settings.txt. 8. Type the password that was used to export the text file. 9. Click Upload.
-

Configuring Pass-Through Security for HiveServer2

With Hive 1.1.0 (through HiveServer2) username and password authentication is supported. Pass-through security with TDV is supported for dynamic domains.

To set up pass-through security for HiveServer2

1. Follow the instructions for HiveServer2 Security Configuration that are available at:

http://www.cloudera.com/content/cloudera-content/cloudera-docs/CDH4/4.2.1/CDH4-Security-Guide/cdh4sg_topic_9_1.html

2. Enable pass-through security configuration of TDV as described in [Dynamic Domain Administration](#).

Configuring Samba and Winbind for NTLM (Tips from an Expert)

If your NTLM security environment includes Samba and Winbind, see [Implementing NTLM Authentication for UNIX](#) Samba and Winbind are not TDV products, so for details on their use refer to their product documentation.

Using Version Control and TDV

Typically, teams of developers work with TDV to implement corporate solutions. When several people all work on the same project, having a way to manage changes to the resources within TDV becomes important. Version control is also valuable for reverting changes or deploying a group of resources altogether at once.

- [About Version Control \(VCS\) for TDV](#)
- [Configuring Version Control for TDV Resources](#)
- [Creating a VCS connection with SSL Authentication](#)
- [Attach a TDV Folder to a VCS Instance](#)
- [Committing TDV Resources to the VCS](#)
- [View the History of a Resource](#)
- [Compare a Resource with Local](#)
- [Revert Changes to a Resource](#)
- [Checking In a Resource to the VCS](#)
- [Checking In Multiple Resources to the VCS](#)
- [Detach a VCS Folder from Your TDV Instance](#)
- [TDV Data Directory Management](#)
- [Manage Connections](#)

About Version Control (VCS) for TDV

You can use TDV with popular version control systems. Version control helps track changes and gives control over changes to TDV resources. For centralized version control, check-outs and check-ins are done to a repository. For distributed version control, changes can be merged into any branch.

Each version control system interacts with TDV a little differently. We recommend setting up some test projects so that you can get familiar with how your resources will

behave. For example, tasks like check in, attach, specifying connection information can vary depending on the version control tool you implement.

Supported Default Version Control Systems

By default, the supported version control systems are:

- SVN client version 1.9.5 and higher
- SVN server version 1.7 and higher
SVN, HTTP, HTTPS protocols supported. SVN+SSH not supported.
- GIT client version 2.11 and higher
- GIT server version 1.7 and higher
GIT, HTTP, HTTPS protocols supported (support one way SSL authentication with certificate). GIT+SSH not supported
- GitHub is used for OAuth authentication. Contact your organization's administrator for the GitHub enterprise account credentials.
- Visual Studio Team Foundation Server 2015 (TFS) with a Git Front-end is supported for HTTP protocol with NLTM authentication only.

Requirements

- TDV is packaged with Microsoft Visual VC++ version 2013. This is required for svn connections to be successful. TDV installer will install this version for you. When using a TDV patch upgrade, this version of VC++ needs to be manually installed. Execute the file vcredist_x64.exe which can be found in the directory <TDV_install_dir>\bin.
- Version control for TDV does not support clusters.
- The TDV VCS user must have access tools rights to access the VCS.
- Configure the GIT VCS temp directory using the Temp Directory configuration parameter (Studio - Configuration - Server - Configuration - Files - Temp Directory (On Server Restart)). For example, set it as D:/temp/vcs.
- If encrypting resources when they are checked into VCS, you must communicate and share the encryption key that you used with others that are authorized to access those resources. There is currently no automatic way to securely share encryption keys between different TDV users.

Limitations

- When reverting changes, uncheck the option “Include TDV Privileges configuration data when fetching from remote VCS”. If this option is selected, then the TDV privileges for the user will be considered and you will see an exception that restricts you from reverting the changes if you don’t have the required privileges.
- When archiving (backup or export) TDV resources if you are a non-admin user, security information is lost. When archiving (backup or export) TDV resources if you are an admin user, security information is retained.

Configuring Version Control for TDV Resources

Through the Manage connections window you can add, delete, and copy connections to your version control systems.

To configure version control

1. Open Studio.
2. Select VCS.
3. Select Manage Connections.
4. Click Add.
5. Type or select values for the following:

Field	GIT Example	SVN Example
Name	vcs-git	vcs-svn
Description	files from git	files from svn
Type	GIT	Subversion
URL	dv-vcs.beesknees.com/git/vcs.git	svn://192.25.5.76:3690
User	gituser2	svnuser1

Field	GIT Example	SVN Example
Password	foa23f9u	foa23f9u

6. Refer [Configuring VCS for GitHub OAuth Authentication](#) for instructions on making a GitHub OAuth VCS connection.
7. Select Verify.
This action sets the credential that TDV uses for VCS during the Studio session that you currently have open. If you close and reopen Studio, you might be asked to log in to the VCS repository again.
8. Select a branch (applicable for GIT connection).
9. Select Create.
10. Select Close when you are done.

Configuring VCS for GitHub OAuth Authentication

In order to configure a VCS connection using the GitHub OAuth authentication, you will need a valid GitHub account. Follow these steps to complete the configuration in TDV:

1. Choose Manage Connections from the VCS menu in TDV Studio. The “Manage Version Control Connections” window is displayed.
2. Enter a Name and Description for the connection.
3. Choose Git (Open Source) in the Type drop down.
4. Provide the Url for the GitHub repository. You will need to create a repository using your GitHub profile. For information on creating a repository, refer: <https://docs.github.com/en/get-started/quickstart/create-a-repo>
5. Before using GitHub OAuth access token in TDV VCS operation, you need to login to GitHub web site to build OAuth app on behalf of TDV GitHub users and get the 'Client ID' and 'Client Secret'. You will need this Client Id and Client Secret of your repository, if you are using OAuth authentication. For more information on creating the Id and Secret, refer: <https://docs.github.com/en/rest/guides/basics-of-authentication>
6. Create and Register the TDV OAuth app. Refer to the following link for information on how to do this - <https://docs.github.com/en/developers/apps/creating-an-oauth->

[app](#)

7. While creating and registering the TDV app as outlined in the above link, please use the sample URL `http://xxx.xxx.xxx.xxx:9400/oauth2callback`. in "Authorization callback URL" column. Authorization callback URL is used by GitHub to send authorization code

For example, you may use this as your callback url while registering TDV -
`http://localhost:9400/oauth2callback`

8. The GitHub user will then have to authorize the generated TDV OAuth app; please refer [build and authorizing OAuth application](#).
9. To complete the connection, provide your GitHub user id and Personal Access Token that you can retrieve from your GitHub profile. Enter the Personal Access Token in the Password field. For instructions on how to create Personal Access Token, refer to the instructions given in the section [Creating a Personal Access Token for GitHub Authentication](#).

Note: You can use the OAuth authentication without giving the personal access token. To do this, click on the OAuth checkbox. You will be transferred to the browser to make a secure connection to your GitHub repository.

10. Click on Verify to verify your connection.

Note: If you get connection issues, verify your credentials and ensure that the secret or the token has not expired. If you continue to have issues, clear your cookie and site data in the browser settings and re-verify your connection.

11. Select a branch of your repository.
12. Set the Encryption password.
13. Click Create to create the VCS connection.

Once a VCS connection is created, all operations (such as Attach, Detach, Commit, Check-in, etc.) can be performed on the VCS instance.

Creating a Personal Access Token for GitHub Authentication

To create a Personal Access token and use it to Manage your TDV VCS connection, follow these steps:

1. Create a Personal Access Token using the instructions outlined [here](#).
2. You will need to authorize the personal access token for use with SAML single

sign on (SSO). To do this, follow these [instructions](#).

3. To manage teams and people with access to your repository, follow these [instructions](#).

Creating a VCS connection with SSL Authentication

Enabling SSL certificate verification for Git can be tuned in TDV through the configuration parameter SSL verify (Server > Configuration > VCS > Git). By default this parameter is set to False. Set it to True to enable certificate verification.

Then follow these steps to create a VCS connection with SSL Authentication:

1. Select Manage Connections from the VCS menu.
2. Click Add.
3. Provide the optional Description field.
4. Choose Git from the VCS type drop-down.
5. Provide the GIT SSH url. This will enable the Certificate tab.
6. Click on the Certificate tab.
7. Click on the “Import Certificate KeyStore from File” button.
8. Choose the KeyStore file in the “Import a KeyStore Certificate” window.
9. Click Ok.

Note: Adding an invalid certificate or not providing a certificate to a Git SSH connection will throw an exception while performing any VCS operation on a resource attached to that instance.

10. Choose the Configuration tab again and provide the user and password to verify the connection.
11. Click Verify. In case of any error in the connection url or credentials, an exception will be thrown.

12. Once the connection is verified, choose a Git branch from the drop-down.
13. Click on Set Encryption button to provide an Encryption password.
14. Click Create.

Once a VCS connection is created, all operations (such as Attach, Detach, Commit, Check-in, etc.) can be performed on the VCS instance.

Change User Credentials for TDV Version Control

To change your credentials for version control

1. Select VCS.
2. Select Set Credentials.
3. Select your version control system.
4. Select Next.
5. Type the new username and password that you want to use.

Note: If you have a VCS connection that uses GitHub account and OAuth authentication, click on the OAuth checkbox in set credentials dialog and you will be transferred to the browser to make a secure connection to your GitHub repository.

Return to the Studio UI and continue with the set credentials process.

Attach a TDV Folder to a VCS Instance

Attaching a folder pulls data from a VCS repository or pushes data to a VCS repository and gives you access to the local files. You can add additional TDV folders to an existing connection at any time. **Te set up a workspace**

1. From the Studio resource tree, select a folder.
2. Select VCS > Attach “<resource>” to Version Control.
3. Select the VCS connection instance for which you want to attach.
4. Select Finish.

5. Select one of the following when asked how to set up your workspace:

Option	Description
Pull-down everything from Version Control and overwrite items in the TDV local workspace.	If your VCS uses a pull-based option, TDV mirrors the folder structure from the data in remote repository. Note: If you choose the Ignore Encryption option, then the password for the resource is left blank.
Use TDV local workspace as a new baseline revision for this connection's Version Control repository.	If your VCS uses a push-based option, TDV overwrites any data in the remote repository with the data from TDV.
Cancel. "<resource>" will no longer be managed by this Version Control connection.	

6. Select Next.
7. Select Download for pull based option or Create for push based option.

Committing TDV Resources to the VCS

These instructions assume that you have already established an attachment. For more information, see [Attach a TDV Folder to a VCS Instance](#).

To commit resources

1. Create new folders and resources within Studio.
2. From the Studio resource tree, select the resource or folder of resources for which you want to perform a commit.
3. Right click and select Version Control.
4. Select check-in <resource>.
5. Enter check-in comments.

6. Click commit changes button.
7. You will be prompted to sign on to Version Control (if you had not set the credentials before).
8. Click Sign in if you had to perform Step 7.

View the History of a Resource

Allows you to view a log of the changes you have made for the resource that you are interested in.

To view the VCS history of a resource

1. From the Studio resource tree, select the resource for which you want to see the change history.
2. Right-click and select Version Control > View History of "<resource>"... .
3. Review the information displayed.
 - Copy the changelist ID copies the Revision.
 - Fetch operation can be done to retrieve a snapshot of resources (current or historical). It can be done at a resource, folder, root or at a connection level.

Note: If you choose the Ignore Encryption option, then the password for the resource is left blank.

View the Full History of a Resource

Allows you to view a log of all the changes across a VCS connection that you are interested in.

To view the full VCS history of a resource

1. From the Studio menu bar, select VCS > Full History.
2. Review the information displayed.
 - Copy the changelist ID copies the Revision.
 - Rollback to a particular version of the resource.

Note: If you choose the Ignore Encryption option, then the password for the resource is left blank.

Compare a Resource with Local

Allows you to view a differences in the XML for the full resource or the SQL of the source only for the resource that you are interested in. You can also select two entries from the history and compare the revisions.

To compare a resource with local

1. From the Studio resource tree, select the resource.
2. Right click a resource and select Version Control > Compare resource with Latest Repository Version. Or:
 - Right-click and select Version Control > View History of “<resource>”... .
 - Select one of the revisions for the <resource>.
 - Select the Compare with local icon.
3. Review the changed, inserted, and deleted changes.
4. Use the buttons at the top to refresh and navigate quickly to the areas of difference.

Revert Changes to a Resource

If reverting a folder, all the changes will be reverted to the last commit.

To revert changes to a resource

1. From the Studio resource tree, select the resource.
2. Right click a resource and select Version Control > Revert Changes for “<resource>”... .

Note: If you choose the Ignore Encryption option, then the password for the resource is left blank.

3. Select OK.

Checking In a Resource to the VCS

Check in adds the selected resources into version control and creates a new version of each resource in the VCS.

To add a resource to VCS

1. From the Studio resource tree, select a folder that is part of your VCS.
2. Right-click and select > New > <resource_type>.
3. Define or edit the resource.
4. Right click the resource and select Version Control > Check-in "<resource>".

Note: If you used the Ignore Encryption option when fetching the data source, then the datasource password will be blank and therefore you will see a warning message while checking in the same resource, indicating that the password for your resource is NULL.

5. Add comments for the check-in.
6. Select Commit Changes.

Checking In Multiple Resources to the VCS

Check in adds the selected resources into version control and creates a new version of each resource in the VCS.

To add multiple resources to VCS

1. From the Studio resource tree, select a folder that is part of your VCS.
2. Define or edit the resources.
3. From Studio Menu, select VCS > Local Changes.
4. In the Local Changes tab that opens, select the check-in button.

Note: If you used the Ignore Encryption option when fetching the data source, then the datasource password will be blank and therefore you will see a warning message while checking in the same resource, indicating that the password for your resource is NULL.

5. Add comments for the check-in.
6. Select Commit Changes.

Detach a VCS Folder from Your TDV Instance

To decouple your work in TDV from your version control systems you can use the detach feature.

To manage your connections to the version control systems

1. From Studio, select the TDV folder that you want to decouple from your VCS.
2. Right-click and select Version Control > Detach from Version Control.
3. Select OK.

TDV Data Directory Management

By default, the TDV installation stores VCS and other TDV related data under <TDV_INSTALL_DIR>/data/vcs. If you do not have sufficient free storage in the location, then the “Data Directory” configuration option allows you to point to another location. To set this option in Studio, go to Administration -> Configuration -> Server -> Configuration -> Files -> Data Directory (On Server Restart) and set the value to a new location.

When this is done and TDV server restarted, the existing data directory will be copied to the new data directory. If there is not enough space for the new data directory, then an exception will occur.

Manage Connections

You can also copy and delete connections from the same area within TDV.

To manage your connections to the version control systems

1. From Studio, select VCS.
2. Select Manage Connections.
3. To delete a version control connection, select the red circle with the minus sign.
4. To edit the connection:
5. Select the connection for which you want to edit information.

6. Select the field that you want to edit and modify it.
7. Select Close to save your changes or select Revert to undo your changes.

Managing Column-Based Security

This topic documents several TDV security features which help you ensure that information is available only to authenticated, authorized individuals who have appropriate rights and privileges.

- [About Column-Based Data Obfuscation](#)
- [Column-Based Restrictions and Privileges](#)
- [Enabling Column-Based Security on TDV Resources](#)
- [Creating or Editing Column Filter Policies](#)
- [Mapping Column Filter Policies to TDV Resources](#)
- [Testing the Column Filter Policies Within Studio](#)

About Column-Based Data Obfuscation

Data obfuscation is a form of data masking where data is scrambled to prevent unauthorized access to data. This form of encryption results in unintelligible or confusing data. Data obfuscation techniques are used to prevent the intrusion of private and sensitive data.

TDV provides for masking of data at the column-based. This means that for selected Columns of data in a relational data organization, data can be obscured. For example, a table that contains a column for social security numbers can have the column data replaced with 10 asterisks (**-**-**** or *****).

Column filter policy server events are captured in the TDV log files.

If you are using row-based security policies in combination with column-based security policies, TDV applies the row-base policy before applying the column-based policy.

Column-Based Restrictions and Privileges

- The following privileges are required: access tools, read all resources, read all users, read all config, modify all resources, modify all config rights.

- Within Studio, privileges are read only for the CBS and Policies folders.
- CBS policies can be assigned to a resource that is impacted, but an error message will display.

Note: Manager displayed in IE8 compatibility mode is slow. Also, IE8 warns that the JavaScript can cause the computer to become unresponsive. It is safe to ignore this warning. When using IE8, you might need to change the Refresh Rate. For IE 11, to launch the Manager pages you might need to turn on the edge mode.

Enabling Column-Based Security on TDV Resources

The Column-Based security for TDV requires that the TDV system-wide option be enabled through Manager.

To enable Column-Based security

1. Launch Web Manager.
2. Sign in as a user an Administrative user with the Modify All Resources right.
3. Select Column-Based Security from the SECURITY tab menu, to open the COLUMN-BASED SECURITY page.
4. Click Change Enabling to toggle between Enable and Disable for the Column-Based Security field.

On this page you can view all the Column Filter policies that have been created. You can add new Column Filter policy, delete existing Column Filter policy, or edit an existing Column Filter policy this page. For each Column Filter policy listed, there are Assignments.

5. To define policies, see [Creating or Editing Column Filter Policies](#) .
6. To assign policies to resources, see [Mapping Column Filter Policies to TDV Resources](#).

Creating or Editing Column Filter Policies

You can use the TDV user interface to lead you through the creation of policies for Column-Based security.

Column names referenced in the filter expression are not validated. You must test each Column Filter policy that you define to make sure that it works as expected. For example, if a filter is on `REGION_NAME` in the `LOCATION` table and the column is changed to `REGION_CODE`, the filter must be updated to reflect that change.

Note: There is no ranking or priority of the rules within a Column Filter policy. If a user or group is assigned to more than one rule within a Column Filter policy, the two predicates are combined using OR statements.

To add or edit a Column-Based security policy

1. Follow the instructions in [Enabling Column-Based Security on TDV Resources](#).
2. To add a new Column Filter policy, click Add Policy. Or, to edit an existing Column Filter policy, select it and click Edit Policy.
3. Select or specify values for the following fields:

Field	Specify
Policy Name	Specifies the name you want to give to the Column Filter policy. Spaces and <code>'</code> are not allowed.
Data Type	• String—Valid assignments are: CHAR, VARCHAR, LONGVARCHAR, BOOLEAN. • Integer—Valid assignments are: TINYINT, SMALLINT, INTEGER, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, DOUBLE, CHAR, VARCHAR, BIT. • Decimal—Valid assignments are: DECIMAL, NUMERIC, FLOAT, REAL, DOUBLE, CHAR, VARCHAR. • Date—Valid assignments are: DATE, CHAR, VARCHAR. • Datetime—Valid assignments are: TIMESTAMP, CHAR, VARCHAR. • Unspecified —Valid assignments are any TDV data type.

Field	Specify
	- Decfloat - Valid assignments are: DECFLOAT, DECIMAL, DOUBLE, FLOAT, REAL, CHAR, VARCHAR. - Double - Valid assignments are: DOUBLE, DECIMAL, DECFLOAT, FLOAT, REAL, CHAR, VARCHAR. - Time - Valid assignments are: TIME, CHAR, VARCHAR.
Enabling	- Enable—Allow the use of the specific Column Filter policy. - Disable—Disallow the use of the specific Column Filter policy. Regardless of this setting, you can still add and edit the Column Filter policy. If the policy is disabled, it is not used against the data.
Annotation	(Optional) Type an explanation of the Column Filter policy.

4. Select the row in the table on the page.

5. Select the pencil to edit a rule or select Add Rule to add a new rule.

Field	Specify
Apply To	Specifies whether this identity is for a Group or a User. Not available for the default policy, because that policy governs all users and groups.
Domain	Specifies the Domain to which you want the Column Filter policy to apply. Not available for the default policy, because that policy governs all users and groups.
User/Group	Specifies the user or group name to which you want the Column Filter policy to apply. Not available for the default policy, because that policy governs all users and groups.

Field	Specify
Rule Type	Specify: – Original Value – Null – Static Value – Partial String Mask – Custom Function – Expression Partial String Mask is available only if you selected String as the Data Type. A default rule is required and is added by default for column-based security. The order of the default rules cannot be changed. You can, however, change the rule type for the default rule.

6. Depending on the Rule Type that you select, you can specify values for the following:

Rule Type	Specify
Original Value	No further fields to edit.
Null	No further fields to edit. All values for the column will display as Null values.
Static Value	Type the value to display for the column data. For example, always the same.
Partial String Mask	Type values for: • Prefix—Number of characters at the beginning of the string to leave alone. • Padding—Type any valid string value. • Suffix—Number of characters at the end of the string to leave alone.

Rule Type	Specify
Custom Function	Select a custom function from the list. This must be a custom function that you have defined in Studio and it must have at least one input and one output.
Expression	Type and expression in the text field. It can be any valid expression syntax.

7. Click Apply.
8. Click Save.

Example of Expression

Any expression can replace the column in select statement when it apply to column.

Example 1:

```
Select "columnname" from "tablename"
```

If you define expression to "columnname", the select statement is rewritten:

```
select "expression" as "columnname" from "tablename"
```

Example 2:

If you want to use the same expression for more than one column, you can use a placeholder instead of the column name.

```
CAST(LPAD('X', LENGTH($PARAM$)-4, 'X') || SUBSTRING($PARAM$, LENGTH($PARAM$)-3) AS VARCHAR(30))
```

At the time of policy assignment, you will be prompted to assign a value for the parameter \$PARAM\$. Use this parameter to assign the expression to any column you need.

Mapping Column Filter Policies to TDV Resources

For each column within a resource within TDV, you can assign a specific Column Filter policy. This procedure assumes that Column Filter policies have been defined that you can assign to the resources. If you need to define Column Filter policies, see [Creating or Editing Column Filter Policies](#) .

Parameters can be defined for assigned columns.

To assign Column Filter policies to resources

1. Follow the instructions in [Enabling Column-Based Security on TDV Resources](#).
2. Click the Edit Assignments tab on the COLUMN-BASED SECURITY page.
3. Expand and locate the resource you are interested in the Resource/Column field.
4. Select the policy you want to apply from the Assigned Policy field. A list of values is provided. The policies that you see in the list will be the policies that are valid to apply for that data type.

Testing the Column Filter Policies Within Studio

TDV provides the ability to test the Column Filter policy within Studio, but recommends that you perform thorough testing of your Column Filter policies by setting up a working test environment that includes all the data sources and client applications that would be interacting with the data for which you have defined Column-Based security.

To test your Column Filter policy within Studio

1. Log in as a user with row-based security (RBS) or column-based security (CBS) access rights or as the Administrative user with the Modify All Resources right.
2. Identify any existing table or view resources for which you want to test Column-Based security.
3. Assign a policy for a non-admin user to at least one column.
4. Use the test identity tab functionality to simulate identity and input non-admin@Domain to verify results.

OR

5. Restart or refresh Studio and log in as the non-admin user.
6. Open and execute the view for which the column-based security policy was applied.
7. Review the data.

Importing and Exporting Column Filter Policies

Column filter policies can be exported from Studio and imported into Studio when you perform a server backup and import.

Requirements

Export and import of column filter policies requires the following rights:

- Access Tools
- Read All Resources
- Read All Users
- Read All Config
- Modify All Resources
- Modify All Config Rights

To use package import and export of column filter policies requires the following rights:

- access tools
- read all resources
- read all users
- read all config
- modify all resources
- modify all config rights

System Event and Log Monitoring

This topic describes TDV's system event and log monitoring capabilities in the Manager Web browser interface as well as in Studio.

The following topics are covered:

- [Configuring Events](#)
- [About Events](#)
- [I/O Log](#)
- [Memory Log](#)
- [Storage Log](#)

Configuring Events

To configure event settings

1. Open Studio.
2. Choose Configuration from the Administration menu.
3. Navigate to: *Events and Logging*
4. Review and modify the event configuration parameters as necessary.

About Events

In Manager, you access information about the server events that have been logged by choosing Event Log from the LOGGING menu. The EVENT LOG page is displayed, providing information about all of the events generated by processes running on the TDV Server.

- [EVENT LOG Summary Information](#)
- [Work with the EVENT LOG Page](#)

- [The EVENT LOG Table](#)

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

Some of the information displayed on the EVENT LOG page is controlled by the configuration settings in Studio. Most of the event data displayed in Manager is also displayed in Studio Manager. For additional information about monitoring events in Studio Manager, see [Events Panel](#).

EVENT LOG Summary Information

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 Studio with the Administration menu Configuration option by editing 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 Studio with the Administration menu Configuration option by editing Maximum Log Entries.

Work 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 displays these columns for each event:

- **ID**- Unique event identifier.

- **Severity** - The severity level of the event, represented by a colored circle, can be one of DISABLED/OFF (gray circle), INFO (green circle), WARNING (yellow circle), or ERROR (red circle).
- **Category** - The type of event, such as REQUEST, SESSION, or TRANSACTION.
- **Type** - The type of event that occurred which can be anything in the event lifecycle including:
START, STOP, RESTART, CREATE, DELETE, ADD, REMOVE, ON, OFF, END, FAIL, CANCEL, COMMIT, ROLLBACK, COMPENSATE, DESTROY, REFRESH, REQUEST, RESPONSE, MISS, SUCCESS, INCREASE, DECREASE, CHECK_OUT, CHECK_IN, INVALID, TERMINATE, MODIFY, IMPACT, OVER, UNDER, PASS, RESET, ROLL, WRITE, WAIT, RUN, EXHAUST, UP, DOWN
- **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 Details in Manager

Every event has additional detailed information available. To view the read-only details, click the Show Row Details button for the row.



In addition to the information presented in the EVENT LOG table described in [The EVENT LOG Table](#), these details are provided:

- **Owner Domain** - The domain to which the Owner belongs.

- Parent Event ID - The ID of the parent event.
- Detail - All logged details about this particular event.

I/O Log

You can view the I/O log in Studio Manager by choosing I/O from the Manager pane. See [I/O Panel](#) for more information.

Memory Log

You can view the Memory log in Studio Manager by choosing Memory from the Manager pane. See [Memory Panel](#) for more information.

Storage Log

You can view the Storage log in Studio Manager by choosing Storage from the Manager pane. See [Storage Panel](#) for more information.

TDV Command-Line Utilities

This topic describes the TDV command-line utilities. Unless otherwise noted, you need to have Access Tools, Read All Resources, and Modify All Resources rights to run these utilities.

Note: Meta Integration Model Bridge (MIMB) for the import of external models is no longer supported.

- [The TDV Export and Import Utilities](#)
- [The TDV Package Import Utility](#)
- [The TDV Package Export Utility](#)
- [The TDV Server Utility Program](#)
- [Using the TDV Server Heap Dump Utility Program](#)

The TDV Export and Import Utilities

TDV provides command-line utilities to export and import CAR files, and back up and restore the system.

These utilities provide more granular control than File > Export and File > Import available from the Studio menu bar. If you need even more control over what is exported and imported, you can use the pkg_export and pkg_import utilities. For more information, see [The TDV Package Import Utility](#) and [The TDV Package Export Utility](#).

- Different versions of these utilities are available for use in different computing environments.

Utility	Description
backup_export	Saves TDV information to a CAR file.
backup_import	Imports CAR file information back into TDV.

- If encrypting resources when performing an export or import, you must communicate and share the encryption key that you used with others that are authorized to access those resources. There is currently no automatic way to securely share encryption keys between different TDV users.

For more information about using Studio to export and import the full TDV Server configuration, see “Using Studio for a Full Server Backup” in the *TDV User Guide*.

The export and import utilities and related topics are discussed in the following sections:

- [About the Backup Export Utility](#)
- [Rights Required for the Backup Export Utility](#)
- [Using the Backup Export Utility](#)
- [Using the Keystore File from an Exported CAR File](#)
- [Rules for the Backup Import Utility](#)
- [Rights Required for the Backup Import Utility](#)
- [Using the Backup Import Utility](#)
- [Modifying the PostgreSQL Repository Maximum Allowed Packets](#)

About the Backup Export Utility

The `backup_export` utility saves a single compressed file that contains all data-source metadata, user-defined resources (published, shared, and other), and server configuration settings. A full server backup file can be used later to restore an entire configuration, with the exception of local machine-based server configuration settings. You have the option to include or exclude custom Java and data source statistics regarding cardinality and table boundaries.

By default, the following are exported: domains, users, groups, all resources, security settings (ownership of resources and privileges on resources), keystore files, cache configurations, scheduling, driver configurations, and server-level configuration settings.

Settings and configurations for *local* computing environments are excluded from a full server backup export. These configuration settings are identified as locally-defined values in the Description area of the parameter pane in the Administration > Configuration window. Locally-defined TDV values include: repository configurations, local port settings, memory settings, log files, event settings and triggers, security settings for anonymous and dynamic users, LDAP properties files, passwords for users, capabilities files, and customized scripts.

Note: To use the keystore file from an exported CAR file, additional steps are necessary. For more information, see [Using the Keystore File from an Exported CAR File](#).

Rights Required for the Backup Export Utility

The backup_export utility requires the following rights:

- Access Tools
- Read All Resources
- Read All Users
- Read All Config

Using the Backup Export Utility

The backup export utility has the following restrictions:

- Export does not include run-time history such as 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.
- The System user of the Operating System where the TDV Server is running on, should have write access to TDV_INSTALL/conf/server directory

To use the backup export utility

1. Open a command prompt window.
2. Navigate to <TDV_install_dir>/bin.
3. Run backup_export:

```
./backup_export
```

```
-server <host_name> [-port <port_number>] [-encrypt]
```

```
-pkgfile <file_name>
```

```
-user <user_name> -password <password> [-domain <domain>]
```

```
-encryptionPassword <encryptionPassword>
```

```
[ -sso ] [ -sspi ] [ -spn <spn> ] [ -krb5Conf <krb5Conf> ]
```

```
[ -pkgname <name> ] [ -description <text> ]
```

```
[ -optfile <file_name> ] [ -optfilePassword <optpassword> ]
```

```
[ -newOptfilePassword <new optpassword> ]
```

```
[ -excludeJars ]
```

```
[ -includeStatistics ]
```

```
[ -genopt <filename> ] [ -optfilePassword <optpassword> ]
```

```
[ -verbose ]
```

The backup_export parameters are described in the following table.

Backup Export Parameters	Optional/ Required	Comments
-server <host_name>	Required	Target TDV server to which the utility is to connect.
-port <port_number>	Optional	Specifies the Web Services base port (HTTP) used to communicate with the TDV Server. The default value is 9400.
-encrypt	Optional	Encrypts communication between the command line and TDV using SSL sent over the dedicated HTTPS port. The HTTPS port is 9402
-pkgfile <file_name>	Required	Specifies the path and file name of the backup archive

Backup Export Parameters	Optional/ Required	Comments
		file (CAR). The file path must be accessible to the command-line client. TDV passes data to the command-line client, and then the client writes that CAR file to the specified location.
-encryptionPassword	Required	The password used to encrypt /decrypt sensitive data in your car file, for example, data source password, etc.
-user <user_name>	Required	User name of the TDV system administrator.
-password <password>	Required	Password of the administrative user who is performing the export. Note: When using an optfile, if the password contains a double quote character, it is necessary to escape the character with a double-quote.
-domain <domain>	Optional	User domain. The default value is composite.
-sso	Optional	Enables SSO authentication. Must be used with the -spn option. There is no need to input user_name or password.
-sspi	Optional	For Windows environments configured with sspi, you can use this parameter after -sso and -spn.
-spn <spn>	Optional	Use this parameter after specifying -sso, to indicate what the service principal name is. Use one of the following formats depending on the protocol you are using: • sspi is <ServiceName/Full_ComputerName@Realm> • JGSS is <ServiceName@Full_ComputerName>
-krb5Conf <krb5Conf>	Optional	For environments configured with multiple Kerberos authentication files, you can use this parameter after -

Backup Export Parameters	Optional/ Required	Comments
		sso and -spn to specify which authentication file to use. <krb5Conf> must specify the full path and filename to use. If this parameter is not specified, the default Kerberos system file is used.
-pkgname <name>	Optional	Names an attribute in the contents.xml within the exported backup file.
-description <text>	Optional	Description of the exported archive file set as an attribute of the contents.xml file within the exported CAR. This description is displayed prior to a Studio-based import.
-optfile <file_name>	Optional	Specifies a file to pass options without using the command line. The options file is useful for hiding password information. For example: <pre>backup_export -server localhost -user test -password <password> -pkgfile sample.car</pre> is the same as: <pre>backup_export -optfile sample.opt -optfilepassword password</pre> where sample.opt contains: <pre>-server localhost -user test -password <password> -pkgfile sample.car</pre> Notes: When the command-line utility is executed, the opt file will be modified and the password strings in the file will be encrypted.

Backup Export Parameters	Optional/ Required	Comments
		• An optfilePassword is required to access the optfile. • TDV versions prior to 8.4 do not have this feature.
-optfilePassword	Required if optfile is used.	Password required to access the option file.
-newOptfilePassword	Optional	Use this option to change the optfile password and re-encrypt the password strings in the option file.
-excludeJars	Optional	Suppresses export of custom Java procedure data source JAR files.
-includeStatistics	Optional	Includes any known cardinality statistics about data source table boundaries, column boundaries, and configurations when statistics gathering is enabled.
-genopt	Optional	By giving this option and an opt file name, the parameters will be stored in the opt file with the password encrypted. While doing an export or import operation the password saved will not be shown as a clear text. The file created can later be used in the -optfile option of this backup_export command. Note: An optfilePassword is required for the generated optfile. TDV versions prior to 8.4 do not have this optfile password feature.
-verbose	Optional	Generates output describing the export in the command-line window.

Using the Keystore File from an Exported CAR File

Keystore files can be exported and imported for reuse, provided you complete some additional configuration steps.

To use the keystore file from an exported CAR file

1. Import the CAR file into Studio. For information on how to export and import, see [Using the Backup Export Utility](#) and [Using the Backup Import Utility](#).
2. Locate the keystore file within <TDV_install_dir>.
3. Rename the file to make sure that it has a unique name.
4. Make sure that the keystore configuration parameters point to the correct keystore file. For more information on how to set the keystore configuration parameters, see [Configuring the Java Keystore File](#).

Rules for the Backup Import Utility

Importing follows these rules to resolve conflicts during import:

- If an imported resource does not exist, it is created. The person performing the import is given creator privileges (such as READ|WRITE|EXECUTE for a procedure or READ|WRITE for a folder). If the user is in the admin group and has specified the -includeaccess option, the resource owner and privileges are restored as well.
- If a resource is imported to a nonexistent folder, the folder and any parent folders that do not yet exist are created with the importing user being granted READ|WRITE privileges and ownership of the folders.

Note: Auto-creation of missing folders is not supported in the Data Services area.

- If an imported resource already exists, the old version is overwritten (assuming you have the WRITE privilege), except that:
 - 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, Abe's privileges are updated but Bob's are left intact.
 - If the resource is a folder or data source, its child resources are not removed.

- Physical data sources cannot be partially exported or imported.
- The Administration > Configuration settings are not carried over when you export or import a resource using Studio menu options.
- You must have WRITE privileges on a folder to create a resource in a folder.
- You must have WRITE privileges on a resource to overwrite that resource.
- You cannot import anything that was exported from the Data Services area to a location outside of that area.
- You cannot import anything that was exported from a location outside of the Data Services area into that area.

Rights Required for the Backup Import Utility

The backup_import utility requires the following rights:

- Access Tools
- Read and Modify All Resources
- Read and Modify All Users
- Read and Modify All Config

Using the Backup Import Utility

The backup import utility imports the entire contents of an archive file created with the backup_export utility or with Studio, subject to the [Rules for the Backup Import Utility](#). Execution of the backup import command requires administrative privileges and a full server export CAR file.

The backup import utility has the following restrictions:

- The -pkgfile target file must be a CAR file exported using either the backup_export utility or the Studio facility for full server backup.
- The -verbose option causes information messages to be displayed after importing. Without this option, only error messages are displayed.
- The -set option lets you specify or change data source connection information from the original.

```
[-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, and password.

- The System user of the Operating System where the TDV Server is running on, should have write access to TDV_INSTALL/conf/server directory

To use the backup import utility

1. Open a command prompt window.
2. Navigate to <TDV_install_dir>/bin.
3. Run the backup_import command:

```
.backup_import -server <hostname> [ -port <port> ] [ -encrypt ]
```

```
-pkgfile <filename> -user <username> -password<password>
```

```
[-domain <domain> ]
```

```
[-sso ] [ -sspi ] [ -spn <spn> ] [ -krb5Conf <krb5Conf> ]
```

```
[-relocate <oldPath> <newPath> ] ...
```

```
[-optfile <filename>] [-optfilePassword <optpassword>]
```

```
[-newOptfilePassword <new optpassword>]
```

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

```
[-encryptionPassword <encryptionPassword>]
```

```
[-reintrospect OR -reintrospectNone ]
```

```
[-createCacheTables ]
```

```
[-updateCacheTables ]
```

```
[-printinfo ]
```

```
[-overwrite ]
```

```
[-genopt <filename>] [-optfilePassword <optpassword>]
```

```
[-verbose ] [-ignoreEncryption ]
```

The backup_import parameters are described in the following table.

Backup Import Parameters	Optional/ Required	Comments
-server <host_name>	Required	Target TDV that the utility connects to for import of the CAR file.
-port <port_number>	Optional	Optionally specifies the Web Services base port (HTTP) used to communicate with the TDV Server. The default value is 9400.
-encrypt	Optional	Encrypts communication between the command line and TDV using SSL sent over the dedicated HTTPS port.
-pkgfile <file_name>	Required	Specifies the location and file name of the CAR file. The -pkgfile target file must be a CAR file exported with backup_export or a full server backup from Studio.
-user <user_name>	Required	User name of TDV system administrator.

Backup Import Parameters	Optional/ Required	Comments
-password <password>	Required	Password of the administrative user performing the import. Note: When using an optfile, if the password contains a double quote character, it is necessary to escape the character with a double-quote.
-domain <domain>	Optional	User domain. The default value is composite.
-sso	Optional	Enables SSO authentication. Must be used with the -spn option. No username or password is required.
-sspi	Optional	For Windows environments configured with SSPI, you can use this parameter after -sso and -spn.
-spn <spn>	Optional	Use this parameter after specifying -sso, to indicate what the service principal name is. Use one of the following formats depending on the protocol you are using: - SSPI format is <ServiceName/Full_ComputerName@Realm> - JGSS format is <ServiceName@Full_ComputerName>
-krb5Conf <krb5Conf>	Optional	For environments configured with multiple Kerberos authentication files, you can use this parameter after -sso and -spn to specify which authentication file to use. <krb5Conf> must specify the full path and file name to use. If this parameter is not specified, the default Kerberos system file is used.
-relocate <old_path>	Optional	Change the location of a resource from having

Backup Import Parameters	Optional/ Required	Comments
<new_path>		been in the old path to being in the new path. This option works on all kinds of resources at all levels and can be used to relocate individual resources or entire containers (folders). The only limitation is that the target location must be valid for the kind of resource being relocated. You can use the -relocate option to exclude specified resources from import by setting <new_path> to NOIMPORT. If you do this, the resources designated by <old_path> are not imported.
-optfile <file_name>	Optional	Specifies a file to pass options without using the command line. The options file is useful for hiding password information. For example: <pre>backup_import -server localhost -user test -password <password> -pkgfile sample.car</pre> is the same as: <pre>backup_import -optfile sample.opt -optfilepassword password</pre> where sample.opt contains:

Backup Import Parameters	Optional/ Required	Comments
		-server localhost -user test -password <password> -pkgfile sample.car Notes: - When the command-line utility is executed, the opt file will be modified and the password strings in the file will be encrypted. - An optfilePassword is required to access the optfile. - TDV versions prior to 8.4 do not have this feature.
-optfilePassword	Required if optfile is used.	Password required to access the option file.
-newOptfilePassword	Optional	Use this option to change the optfile password and re-encrypt the password strings in the option file.
-set <path> <type> <attribute> <value>	Optional	Enables you to change resource attributes during import. You can repeat this option to set different attributes or multiple class paths. - The <path> is the TDV resource name. - The <type> is DATA_SOURCE when the

Backup Import Parameters	Optional/ Required	Comments
		<attribute> is classpath, host, port, database, user, or password. - The <attribute> can be (depending on source type): - user <login> or <user_name> or error - password <password> or error - user2 <app_user_name> or error if not Oracle EBS - password2 <app_password> or error if not Oracle EBS - host <url_IP> or <dsn> or <server> or <appServer> or <url> or <root> or error - port <url_port> or <port> or error - database <url_database_name> or <enterprise> or <app_server> or error - path <root> or <url> or error - annotation - Set <value> to a valid entry for the selected attribute. String values with spaces can be enclosed in double-quotes. For Windows systems, use semicolon as delimiter: <pre>C:\DevZone\ATeam\Jars\my.jar; D:\Current\Ref\classes</pre> For UNIX systems, use colons as delimiter: <pre>/lib/ext/classes:/lib/src/jars</pre>
-encryptionPassword	Optional	The password used to encrypt /decrypt sensitive

Backup Import Parameters	Optional/ Required	Comments
		data in your car file, for example, data source password, etc. You do not need to specify this parameter if you are importing a car file from a TDV version prior to 8.0.
-reIntrospect	Optional	This option is used to control if it does re-introspections for all data sources at end of the import or not
-createcachetables		If used, the cache status, tracking and target tables required by cached resources are created, if not already present in the database for your data source as it is described in the metadata of the CAR file you are importing. The import process does not create the TDV resources for cache_status, cache_tracking, and target tables if they are not described in the CAR file metadata.
-updateCacheTables	Optional	Yes or no. Use this optional value to indicate whether you want to have the cache tables dropped and recreated with any potential changes that are included in the CAR file.
-printinfo	Optional	Disables the actual import process, and instead prints the archive file in the command window for you to view.
-overwrite	Optional	Clears all resources prior to importing the backup CAR. Ensures that the TDV resource tree matches the contents of the imported CAR file.
-genopt	Optional	By giving this option and an opt file name, the parameters will be stored in the opt file with the

Backup Import Parameters	Optional/ Required	Comments
		password encrypted. While doing an export or import operation the password saved will not be shown as a clear text. Note: An optfilePassword is required for the generated optfile. TDV versions prior to 8.4 do not have this optfile password feature.
-verbose	Optional	Generates output in the command-line window describing the export process. By default, only error messages are displayed or logged.
-ignoreEncryption	Optional	If this option is used, then all backup data will be imported regardless of whether a valid encryption key was provided. This means that the import will not fail. This option can be used to allow partially importing any backed up data. However, the import process will only import data that is not encrypted or can be decrypted using the provided encryption key. All encrypted portions of the backup data that cannot be decrypted will be imported as empty values and the import will otherwise succeed. This affects all encrypted values in the backup data, which includes, but is not limited to data source and LDAP domain connection passwords.

Here is a sample command line to change the password property of a data source:

```
backup_import
-user admin -password admin 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.

Modifying the PostgreSQL Repository Maximum Allowed Packets

If you use backup_import for a large CAR file, and if the number of packets being saved is too great for the current packet setting of the PostgreSQL repository, the following TDV error message appears:

```
com.compositesw.cdms.Webapi.WebapiException: One or more resources could not be
saved. Packet for query is too large ... You can change this value on the server by setting
the max_allowed_packet' variable.
```

If you encounter this error message, you can change a TDV PostgreSQL repository setting to allow a larger number of packets to be saved to the repository.

To change the PostgreSQL repository maximum allowed packets

1. Using a PostgreSQL admin tool, log in to the PostgreSQL client as the root user.
2. Change maximum allowed packets:

```
PostgreSQL> set @@max_allowed_packet = 25165824;
```

3. Verify that the setting was accepted:

```
PostgreSQL> show variables like '%max_allowed%';
```

```
+-----+-----+
```

```
| Variable_name          |-----| Value ---|
```

```
+-----+-----+
```

```
| max_allowed_packet | 25165824 |
```

```
+-----+-----+
```

The TDV Package Import Utility

The `pkg_import` command-line utility lets you import specified TDV Server resources.

- [Rules for the Package Import Utility](#)
- [Restrictions for the Package Import Utility](#)
- [Using the Package Import Utility](#)

The `pkg_import` utility is available in `<TDV_install_dir>/bin`. Different versions are provided for use in different computing environments.

Platform	Utility
Windows	<code>pkg_import.bat</code>
UNIX	<code>pkg_import.sh</code>

For details about using Studio to import selected resources, see the *TDV User Guide*.

Rules for the Package Import Utility

Importing follows these rules to resolve conflicts during import:

- If an imported resource does not exist prior to import, it is created. The user performing the import is given all privileges of 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, 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.

- If a resource is imported to a nonexistent folder, the folder and any parent folders that do not yet exist are created with the importing user being granted READ|WRITE privileges and ownership of the folders.

Note: Auto-creation of folders is not supported in the Data Services folder in the resource tree.

- If an imported resource already exists, the old version is overwritten (assuming you have the WRITE privilege), except that:
 - 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, Abe's privileges are updated but Bob's are left intact.
 - If the resource is a folder or data source, its child resources are not removed.

Restrictions for the Package Import Utility

The package import utility has the following restrictions:

- Server configuration parameter settings are not included when a resource is exported or imported using the pkg_export and pkg_import utilities, or when specifying a single resource for export using Studio.
- Importing into a folder requires the WRITE privilege on the destination folder and READ privileges on the parent directory path.
- Overwriting or deleting a TDV resource requires WRITE privilege on that resource.
- The System user of the Operating System where the TDV Server is running on, should have write access to TDV_INSTALL/conf/server directory
- Physical data sources cannot be partially exported or imported.
- You cannot import anything that was exported from the Data Services area to a location outside of that area.
- You cannot import anything that was exported from a location outside of the Data Services area into that area.

Using the Package Import Utility

The package import (pkg_import) utility imports directories and resources from a zipped CAR file. The pkg_import utility is available for execution from the bin directory of the TDV installation.

Use of pkg_import requires the following rights to restore TDV-defined resources:

- Access Tools
- Read and Modify All Config
- Read and Modify All Resources
- Read and Modify All Users

Import generally requires multiple administrative rights because it overwrites many types of resource objects, directories, user privileges, and other resource definitions.

Note: When you import using the `pkg_import` utility, a data source will be re-introspected if the "set" option is used to change its connection properties.

The package being imported, and the options specified, dictate what rights are actually required to perform the import. Most import procedures require Modify All Resources and Modify All Users.

The `pkg_import` utility imports the resources in the archive file into the server following the import rules ([Rules for the Package Import Utility](#)). If you use `pkg_import -pkgfile` on an archive created using `backup_export`, only the resource information is used.

Different combinations of `-overwrite`, `-includeusers`, and `-mergeusers` options have different outcomes for the user definitions, resource ownership, and usage privileges defined on the TDV target after import is complete.

This table shows the eight combinations possible with these three import options. In each row, check marks on the left show the import options used, and check marks on the right show what users are defined on the TDV target after the import with the given options is performed.

Package Import Options			Users Imported into Target			
-overwrite	-includeusers	-mergeusers	From CAR File		TDV Target	
			User A	User B (CAR)	User B (TDV)	User C
X					X	X
X	X		X	X		
X	X	X	X	X		X

Package Import Options			Users Imported into Target			
-overwrite	-includeusers	-mergeusers	From CAR File		TDV Target	
			User A	User B (CAR)	User B (TDV)	User C
X		X	X	X		X
	X		X	X		X
	X	X	X		X	X
		X	X		X	X

Note: The combination of -overwrite and -includeusers (without -mergeusers) removes all existing user definitions and replaces them with any users present in the CAR file being imported. The user performing the import must have Read and Modify All Users rights to be able to use this combination.

To use the package import utility

1. Open a command prompt window.
2. Navigate to <TDV_install_dir>/bin.
3. Run the pkg_import command:

```
./pkg_import -pkgfile <D:/directory/Path/and/File_Name.car>...
```

```
-server <host_name> [-port <port_number>] [-encrypt]
```

```
-user <user_name> -password <password> [-domain <domain>]
```

```
[-sso] [-sspi] [-spn <spn>] [-krb5Conf <krb5Conf>]
```

```
[-optfile <path_and_filename>] ...
```

```
[-optfilePassword <optpassword>]
```

```
[-newOptfilePassword <new optpassword>]
```

```
[-relocate <old_path> <new_path>] ...
```

```
[-rebind <old_path> <new_path>] ...
```

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

```
[-encryptionPassword <encryptionPassword> ] ...
```

```
[-printinfo] [-printroots] [-printusers][ -includeusers ] [ -mergeusers ]
```

```
[-printcontents] [-printreferences]
```

```
[-includeaccess] [-nocaching] [-nopolicy] [-createcachetables]
```

```
[-updateCacheTables][-excludejars] [-nosourceinfo]
```

```
[-overwrite] [-overrideLocks] [-messagesonly]
```

```
[-genopt <filename> ] [-optfilePassword <optpassword>]
```

```
[-verbose] [-quiet]
```

```
[-ignoreEncryption]
```

The table below describes the pkg_import parameters.

Package Import Parameters	Optional/ Required	Comments
-pkgfile <file_name>	Required	Specifies one or more import CAR files. The file names should be specified as absolute paths from mapped directories.
-server <host_name>	Required	TDV Server host to which the utility is to connect.
-port <port_number>	Optional	Specifies port for the target TDV instance. Default is 9400.
-encrypt	Optional	Encrypts communication between the command line and TDV using SSL sent over the dedicated HTTPS port.
-user <user_name>	Required	User name of profile used to import. User rights specified by the target TDV instance grant permission to import, and can restrict Write privileges to designated directories.
-password <password>	Required	Password for user profile used to export package. Note: When using an optfile, if the password contains a double quote character, it is necessary to escape the character with a double-quote.
-domain <domain>	Optional	Domain of the user performing the import. If it is omitted, the assumed value is composite.
-sso	Optional	Enables SSO authentication. Must be used with the -spn option. There is no need to input username or password.
-sspi	Optional	For Windows environments configured with SSPI, you can use this parameter after -sso and

Package Import Parameters	Optional/ Required	Comments
		-spn.
-spn <spn>	Optional	Use this parameter after specifying -sso, to indicate what the service principal name is. Use one of the following formats depending on the protocol you are using: - sspi is <ServiceName/Full_ComputerName@Realm> - JGSS is <ServiceName@Full_ComputerName>
-krb5Conf <krb5Conf>	Optional	For environments configured with multiple Kerberos authentication files, you can use this parameter after -sso and -spn to specify which authentication file to use. <krb5Conf> must specify the full path and filename to use. If this parameter is not specified, the default Kerberos system file is used.
-optfile <file_name>	Optional	Specifies a file to pass options without using the command line. The options file is useful for hiding password information. For example: <pre>pkg_import -server localhost -user test -password password</pre>

Package Import Parameters	Optional/ Required	Comments
		-pkgfile sample.car is the same as: pkg_import -optfile sample.opt -optfilepassword password where sample.opt contains: -server localhost -user test -password password -pkgfile sample.car Notes: • When the command-line utility is executed, the opt file will be modified and the password strings in the file will be encrypted. • An optfilePassword is required to access the optfile. • TDV versions prior to 8.4 do not have this feature.
-optfilePassword	Required if optfile is used.	Password required to access the option file.
-newOptfilePassword	Optional	Use this option to change the optfile password

Package Import Parameters	Optional/ Required	Comments
		and re-encrypt the password strings in the option file.
-relocate <old_path> <new_path>	Optional	Change the location of a resource from having been in the old path to being in the new path. This option works on all kinds of resources at all levels and can be used to relocate individual resources or entire containers (folders). The only limitation is that the target location must be valid for the kind of resource being relocated. Relocating a resource modifies (rebinds) references made by other resources being imported, but does not modify references to resources that are not part of the import.
-rebind <old_path> <new_path>	Optional	Sets a new resource path for a dependency resource. The option can be repeated. Resources can be rebound as a group during the import process. This option can be used when migrating from a development or test environment to a production server deployment. All imported resources are rebound. If rebinds caused by the relocate and rebind flags conflict, the rebind ones are performed first.
-set <path> <type> <attribute> <value>	Optional	Enables you to change resource attributes during import. You can repeat this option to set different attributes or multiple class paths. String values can be enclosed in double-quotes to allow for spaces. • The <path> is the TDV resource name. • The <type> is DATA_SOURCE when the <attribute> is classpath, host, port,

Package Import Parameters	Optional/ Required	Comments
		database, user, or password. The <attribute> can be one or more of the attributes of a data source. All the attributes of a datasource can be changed using the -set option. The attributes vary with each datasource. The system table SYS_DATASOURCE_ATTRIBUTE_DEFS lists all the attributes for each datasource. For Windows systems, use semicolon as delimiter: C:\DevZone\ATeam\Jars\my.jar;D:\Current\Ref\classes For UNIX systems, use colons as delimiter: /lib/ext/classes:/lib/src/jars
-encryptionPassword	Optional	The password used to encrypt /decrypt sensitive data in your car file, for example, data source password, etc. You do not need to specify this parameter if you are importing a car file from a TDV version prior to 8.0.
-printinfo	Optional	Causes the archive file to be examined and information about it to be displayed. The archive file is <i>not</i> imported when this option is specified.
-printroots	Optional	Prints the new paths to the imported resources. The archive file is <i>not</i> imported when this option is specified.
-printusers	Optional	Prints the user names of the owners of the imported resources and their associated user groups. The archive file is <i>not</i> imported when this option is specified.

Package Import Parameters	Optional/ Required	Comments
-includeusers	Optional	Imports all users present in the exported package CAR file unless mergeusers is specified. By default domain, groups, and user information are not included in export or import packages.
-mergeusers	Optional	Imports all users present in the CAR file who are not already present in the server target. This option takes precedence over the includeusers option, and can have different behavior when combined with the overwrite option.
-printcontents	Optional	Disables actual import, and prints properties of the CAR file to the command window.
-printreferences	Optional	Prints a list of resources referred to by the imported resources. The archive file is <i>not</i> imported when this option is specified.
-includeaccess	Optional	This 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.
-nocaching	Optional	Caching configuration is imported by default. This option must be used if you want to ignore cache configurations.
-nopolicy	Optional	CBS, RBS and Security policies are imported by default. Set this option if you want to ignore the configuration.
-createcachetables	Optional	If used, the cache status, tracking and target tables required by cached resources are created, if not already present in the database for your data source as it is described in the metadata of

Package Import Parameters	Optional/ Required	Comments
		the CAR file you are importing. The import process does not create the TDV resources for <code>cache_status</code>, <code>cache_tracking</code>, and target tables if they are not described in the CAR file metadata.
<code>-updateCacheTables</code>	Optional	Yes or no. Use this optional value to indicate whether you want to have the cache tables dropped and recreated with any potential changes that are included in the CAR file.
<code>-excludejars</code>	Optional	Does not import custom Java procedure data source's JAR files within the CAR file.
<code>-nosourceinfo</code>	Optional	(An overwrite safeguard.) Suppresses import of the following pre-existing connection attributes when an otherwise identical resource is already present in 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.
<code>-overwrite</code>	Optional	Ensures that TDV exactly matches the directories present in the CAR file. Clears targeted folder directories before copying the CAR file contents to those directories. The <code>pkg_import</code> utility clears only those directories that have representative resources in the CAR file, whether or not the <code>-overwrite</code> option is specified.
<code>-overrideLocks</code>	Optional	If the import user is not the locked resource

Package Import Parameters	Optional/ Required	Comments
		owner, and if this option to override is specified, then it is overwritten, else an exception is reported.
-messagesonly	Optional	Displays the messages generated in a package import without actually performing the import.
-genopt	Optional	By giving this option and an opt file name, the parameters will be stored in the opt file with the password encrypted. While doing an export or import operation the password saved will not be shown as a clear text. Note: An optfilePassword is required for the generated optfile. TDV versions prior to 8.4 do not have this optfile password feature.
-verbose	Optional	Reports problems encountered during the import. If neither verbose nor quiet is mentioned, verbose is the default behavior.
-quiet	Optional	Command information is not reported if this option is mentioned.
-ignoreEncryption	Optional	If this option is used, then all backup data will be imported regardless of whether a valid encryption key was provided. This means that the import will not fail. This option can be used to allow partially importing any backed up data. However, the import process will only import data that is not encrypted or can be decrypted using the provided encryption key. All encrypted portions of the backup data that cannot be decrypted will be imported as empty values and the import will otherwise succeed. This affects all encrypted values in the backup

Package Import Parameters	Optional/ Required	Comments
		data, which includes, but is not limited to data source and LDAP domain connection passwords.

Example 1

This example changes the password property of a data source:

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

```
pkg_import.sh -pkgfile C:/Store/EnterpriseArchive/mycar.car
-server localhost -user ProdAdmin -password AdminPassW0rd
-set /shared/myDataSource DATA_SOURCE password MyNewPassword
```

Example 2

This example changes the resource attributes:

```
./pkg_import.sh -pkgfile MyCustJavaDB.car
-optfile C:\X.opt -optfilepassword password
-set /shared/datasources/DBCustCustom_Java DATA_SOURCE
classpath "C:\Program Files\JavaDev;D:\My Documents\JavaJim"
```

Example 3

This example imports a WSDL file, where the host attribute is used:

```
pkg_import.bat -pkgfile ..\TEST.car -server localhost -port 9420 -user admin -password
sam -domain composite
-set /shared/test/DataSources/testWebService DATA_SOURCE
```

```
host http://localhost:9430/services/TEST/testWebService?wsdl
-set /shared/test/DataSources/testWebService DATA_SOURCE user admin
-set /shared/test/DataSources/testWebService DATA_SOURCE password admin
```

The TDV Package Export Utility

The `pkg_export` command-line utility let you export specified TDV Server directories and resources to a single CAR file, which provides these advantages:

- All metadata files and resources are aggregated.
- Multiple resource files are packaged in a compressed, portable archive.

The `pkg_export` utility is available in `<TDV_install_dir>/bin`. Different versions are provided for use in different computing environments.

Platform	Utility
Windows	<code>pkg_export.bat</code>
UNIX	<code>pkg_export.sh</code>

Note: For details about using Studio to export selected resources, see the *TDV User Guide*.

The System user of the Operating System where the TDV Server is running on, should have write access to `TDV_INSTALL/conf/server` directory

The `pkg_export` utility exports each of the listed resources (using a namespace path such as `/users/composite/manager/sources`). It does not include any domains, users, groups, or server settings.

Any user can execute this command. If you do not have Read privilege on any of the specified resources, the export fails. If you have the Read privilege for all resources and their children, all children are also included.

Note: Child resources for which the user does not have the Read privilege are omitted from the export package, without notice.

To use the package export utility

1. Open a command prompt window.
2. Navigate to `<TDV_install_dir>/bin`.
3. Run the `pkg_export` command:

```
./pkg_export -pkgfile <filename>
```

```
-encryptionPassword <encryptionPassword>
```

```
-server <hostname> [ -port <port> ] [ -encrypt ]
```

```
[ -user <username> ] [ -password <password> ] [ -domain <domain> ]
```

```
[ -sso ] [ -sspi ] [ -spn <spn> ] [ -krb5Conf <krb5Conf> ]
```

```
[ -pkgname <name> ] [ -description <text> ]
```

```
[ -optfile <filename> ] ...
```

```
[ -optfilePassword <optpassword> ]
```

```
[ -newOptfilePassword <new optpassword> ]
```

```
[ -rebindable <path> <description> ] ...
```

```
[ -includeaccess ] [ -includecaching ]
```

```
[ -nosourceinfo ] [ -includejars ]
```

```
[ -includeAllUsers ] [ -includeUser <domain> <user> ] ...
```

```
[ -includeGroup <domain> <group> ] ... [ -includeDomain <domain> ]
```

```
[ -includeRequiredUsers ] [ -includeDependencies ]
```

```
[-includeStatistics ]
```

```
[-genopt <filename> ] [-optfilePassword <optpassword>]
```

```
[-verbose ] [ -quiet ]
```

```
[-includeParentResources]
```

The table below describes the pkg_export parameters, in alphabetical order.

Package Export Parameter	Optional / Required	Comments
-pkgfile <file_name>	Required	The new CAR file name.
-encryptionPassword	Required	The password used to encrypt /decrypt sensitive data in your car file, for example, data source password, etc.
-server <host_name>	Required	TDV Server host to which the utility connect.
-port <port_number>	Optional	Specifies Web Services base port (HTTP) to use to communicate with the TDV Server. The default value is 9400.
-encrypt	Optional	Encrypts communication between the command line and TDV using SSL sent over the dedicated HTTPS port. The HTTPS port is the base port number plus 2.
-user <user_name>	Required	TDV system administrative user name. 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>	Required	Password for user profile used to export

Package Export Parameter	Optional / Required	Comments
		package. Note: When using an optfile, if the password contains a double quote character, it is necessary to escape the character with a double-quote.
-domain <domain>	Optional	User domain. The default value is composite. Specify a value if the exporting user's domain is not composite.
-sso	Optional	Enables SSO authentication. Must be used with the -spn option. You do not need to input a user name or password.
-sspi	Optional	For Windows environments configured with SSPI, you can use this parameter after -sso and -spn.
-spn <spn>	Optional	Use this parameter after specifying -sso, to indicate what the service principal name is. Use one of the following formats depending on the protocol you are using: - sspi is <ServiceName/Full_ComputerName@Realm> - JGSS is <ServiceName@Full_ComputerName>
-krb5Conf <krb5Conf>	Optional	For environments configured with multiple Kerberos authentication files, you can use this parameter after -sso and -spn to specify which authentication file to use. <krb5Conf> must specify the full path and filename to use.

Package Export Parameter	Optional / Required	Comments
		If this parameter is not specified, the default Kerberos system file is used.
-pkgname "<name>"	Optional	Package name can be set. Spaces are allowed, but punctuation is not.
-description "<text>"	Optional	Package description of the archive file. Notation appears when Studio is used to import the CAR.

Package Export Parameter	Optional / Required	Comments

-optfile <file_name>	Optional	Specifies an options file to pass options without using the command line. The options file is useful for hiding password information. For example:

```
pkg_export
-server localhost
-user test
-password password
-pkgfile sample.car
```

is the same as:

```
pkg_export
-optfile sample.opt -optfilepassword password
```

where sample.opt contains:

```
-server localhost -user test
-password password
-pkgfile sample.car
```

Notes:

- When the command-line utility is executed, the opt file will be modified and the password strings in the file will be encrypted.
- An optfilePassword is required to access the optfile.

Package Export Parameter	Optional / Required	Comments
-optfilePassword	Required if optfile is used.	Password required to access the option file.
-newOptfilePassword	Optional	Use this option to change the optfile password and re-encrypt the password strings in the option file.
-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 <code>-rebind</code> option must be specified on import for that action to take place on import. That message is also displayed in Studio to prompt designation of a new resource (path) as the resource dependency. Rebinding must be done after the import unless this option specifies the new resource for rebinding during import.
-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 can occur if this option is used and the exporting user is not a member of the admin group (which has the Read All Resources right).
-includeCaching	Optional	Includes the details of caching on views and procedures in the export file. This option must be specified to include cached data from materialized views, or configurations that

Package Export Parameter	Optional / Required	Comments
		include scheduling for cache refreshes.
-nosourceinfo	Optional	Data source connection details (such as user name, password, host name, and port) are included by default. Specify -nosourceinfo to exclude these. If passwords are included, they are encrypted.
-includeJars	Optional	Exports any included custom Java procedure data source's JAR.
-includeallusers	Optional	Exports all domains, groups, and users to the export file. Requires the Read All Users right.
-includeuser <domain_name> <user_name>	Optional	Includes the specified user in the export file. This option can be repeated to export multiple users. Repeat the option keyword -includeUser with arguments for the new domain and user as many times as necessary.
-includegroup <domain_name> <group_name>	Optional	Exports group information about the specified group in the export file.
-includedomain <domain_name>	Optional	Exports the specified domain metadata to the export file.
-includeRequiredUsers	Optional	Includes the information about the required users in the export file.
-includeDependencies	Optional	Gathers and includes all dependent resources for the resources you choose to export.
-includeStatistics	Optional	Includes any resource statistics known about the table or column boundaries.

Package Export Parameter	Optional / Required	Comments
-genopt	Optional	By giving this option and an opt file name, the parameters will be stored in the opt file with the password encrypted. While doing an export or import operation the password saved will not be shown as a clear text. Note: An optfilePassword is required for the generated optfile. TDV versions prior to 8.4 do not have this optfile password feature.
-verbose	Optional	Reports problems encountered during the export. If neither verbose nor quiet is mentioned, verbose is the default behavior.
-quiet	Optional	Command information is not reported when this option is set.
-includeParentResources	Optional	By enabling this option, the parent resources are included while exporting the data source metadata.

Example

In this example, myParameterizedQuery is exported with dependencies that include the products table from the orders 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.

```
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.
```

Example 2

In this example, shared/sources is backed up to Sources_Backup.car.

```
pkg_export -pkgfile Sources_Backup.car
```

```
shared/sources
```

```
-optfile C:/BackupScripts/Sources/weekly.opt -optfilepassword password
```

```
-includeDependencies
```

```
-nosourceinfo
```

The options file weekly.opt must contain any required arguments that were missing in the original command. For example, the following options with some value might be required:

```
-server localhost
```

```
-user DBASecure1
```

```
-password Password
```

```
-domain EnterpriseLDAP
```

An options file can be dynamically generated to specify options and arguments, including the user name, password, and domain. This makes it possible to set up

programmatic backups, while preventing the DBA login from being displayed in the application window or in the file prior to running the scheduled script.

The TDV Server Utility Program

Typically, you are only asked to use the TDV server utility program (server_util) by the TDV Support team. You can use this utility to:

- Retrieve server performance profile reports
- Reset the system namespace
- Get or set the server name
- Deploy or undeploy a package or Pluggable Authentication Module (PAM)
- Generate log files
- Calculate and save object memory sizes
- Regenerate files that are based on configuration settings

Instructions for how to perform certain tasks using the server_util script are contained throughout this guide, and in the *TDV User Guide and TDV Installation and Upgrade Guide*.

The server_util program is described in these sections:

- [Using the Server Utility](#)
- [Server_util.sh Examples](#)

Server Utility Performance Profile Report

The table below describes a few of the metrics contained in the server performance profile report.

Metric	Description
components.archive	Import and export detail.
components.operator	Query engine processing time for SQL that could not be pushed to the underlying data sources during.

Metric	Description
ds	Aggregated time required to communicate with external data sources.
internal.repository	Repository response time for metadata information gathering.
request.data	Aggregated amount of time required to send data to the requesting client after SQL processing was underway.
request.setup.sql	Time to construct the SQL sent to outside data sources.

Using the Server Utility

The `server_util` program is invoked with options that specify the user and environment, and the action to take.

To use the `server_util` program

1. Open a command prompt window.
2. Navigate to `<TDV_install_dir>/bin`.
3. Enter the `server_util` syntax:

```
./server_util -server <host_name> [ -port <port> ] [ -encrypt ]
```

```
-user <user_name> -password <password> [ -domain <domain> ]
```

```
<command> [-verbose]
```

Server Utility Command Options

The table below describes the server utility command options, in alphabetical order.

Command Options	Comments
-clearProfile	Clears all existing profiling data.
-createMemorySizeFile	Calculates and saves object memory sizes.
-deploy -package <package_ name> [-checksum <algorithm>]	Deploys a package file, which can be a Pluggable Authentication Module (PAM), a CAR file, or a JAR file. Provide a checksum if one is furnished with the file.
-getServerName	Retrieves the server name.
-profile	Displays current server performance profile data, which is by default collected by a background process. See Server Utility Performance Profile Report .
-regenerateFiles	Regenerates files that are based on configuration settings. Restarting TDV or using the Configuration window Apply button also regenerates such files.
-resetNamespace	Reset the system namespace. Resets the <i>server</i> namespace to show changes to the <i>system</i> namespace—for example, a system table change after application of a patch.
-rollbackRepository [-toVersion x.y.z]	Rolls back any repository structure changes a patch created, so that the patch can be rolled back cleanly, and then shuts down the server. Refer to the release notes for your server version to see what target versions are available. For certain target versions, rollbackRepository uses a -toVersion argument.
-saveLogs [-port <port_number>] [-folder <file_path>]	Saves log files. See System Event and Log Monitoring . The <file_group> can be any combination of logs, conf, and sysinfo.

Command Options	Comments
<code>[-exclude <file_group>]</code>	
<code>-setServerName</code> <code>-serverName <server_name></code>	Sets the server name, which must be a unique display name.
<code>-undeploy</code> <code>-name <package_name></code> <code>-version <version_number></code>	Undeploys a package, which can be a Pluggable Authentication Module (PAM), a CAR file, or a JAR file. The version number for an extension adapter is always 1. The package name typically is specified without the dot-suffix

Server Utility Parameters

The table below describes the `server_util` command parameters (not command parameters).

Server Utility Parameters	Comments
Required	
<code>-password <password></code>	Password of the administrative user who is performing the command.
<code>-server <host_name></code>	Target TDV server to which the utility is to connect.
<code><command></code>	Exactly one subcommand is required for the <code>server_util</code> command.
<code>-user <user_name></code>	User name of the TDV system administrator.
Optional	
<code>-domain <domain></code>	User domain. The default value is composite.
<code>-encrypt</code>	Encrypts communication between the command line and TDV

Server Utility Parameters	Comments
	using SSL sent over the dedicated HTTPS port.
-port <port_number>	Specifies the Web Services base port (HTTP) used to communicate with the TDV Server. Default is 9400.
-verbose	Generates output describing the process and its progress in the command-line window.

Server_util.sh Examples

The -profile subcommand lists the server profile for a default installation:

```
server_util -server host_name -user admin -password admin -profile
```

The -clearProfile subcommand clears the server profile information. Aggregate and average statistics data start again from zero.

```
server_util -server localhost -user admin -password admin -clearProfile
```

The -resetNamespace subcommand resets the system namespace.

```
server_util -server localhost -user admin -password admin -resetNamespace
```

The -deploy subcommand deploys a package, adding it to the Studio resource tree.

```
server_util -server localhost -user admin -password admin -deploy myadapter.jar -verbose
```

The -undeploy subcommand removes a package from the Studio resource tree.

```
server_util -server localhost -user admin -password admin -undeploy myadapter -verbose
```

Using the TDV Server Heap Dump Utility Program

You can use the `composite_dumpHeap` utility to generate a heap dump in a HPROF format. HPROF format can be used to track down and isolate performance problems involving memory usage and inefficient code.

The `composite_dumpHeap` utility is not supported for AIX platforms.

To use the `composite_dumpHeap` program

1. Open a command prompt window.
2. Navigate to `<TDV_install_dir>/bin`.
3. Enter the `composite_dumpHeap` syntax:

```
./composite_dumpHeap.[bat, sh] - argument 1 <dumpFileNamePath>
```

If this `<dumpFileNamePath>` exists then it will be deleted and the file will be regenerated.

Deployment Manager

This topic describes Deployment Manager, a web-based tool you can use to manage and streamline the development life cycle of TIBCO® Data Virtualization (TDV) resources. Deployment Manager enables you to build repeatable deployment plans to seamlessly promote resources across environments. Typically, the migration is from development machines to test machines to production machines.

- [About Deployment Manager](#)
- [Starting Deployment Manager](#)
- [Defining Sites](#)
- [Defining Deployment Plans](#)
- [Executing Deployment Plans](#)
- [Backing Up and Restoring the Deployment Manager Server](#)

About Deployment Manager

Most organizations use a software development life cycle where new features for mission-critical applications are carefully tested prior to moving them into production. TDV lets you import and export resources between environments using CAR files, but for large-scale migration, the work to reconfigure those resources and rebind the resources can be challenging.

Deployment Manager simplifies the development life cycle by helping you define TDV sites and their resources and deployment plans so that you can easily migrate resources from one site to another. Site resources include data source definitions, tables, views, procedures, and so on as well as domains, users, and groups. Deployment Manager helps you manage the definitions of these resources on the target and source sites and the relationships between them. You bundle resources together and create a deployment plan that defines a sequence of migration and other operations to:

- Migrate resources and users.
- Rebuild or create new cache databases.
- Remove resources on the target site.


- Execute procedures to customize a deployment.
- Deploy the plan from your browser or the command line.
- Export the deployment plan to execute at a later time.

See these topics for more overview information about Deployment Manager:

- [Deployment Manager Architecture](#)
- [Basic Deployment Manager Concepts and Definitions](#)
- [User Roles and Workflows](#)

Deployment Manager Architecture

Typically, TDV developers and administrators use a development life cycle which is a sequence of environments (for example, development, test, and production environments) to build and test new features and functions. As testing and validation are performed, the code needs to be migrated to the next environment. The architecture of Deployment Manager supports the migration of resources from one environment to another as shown here:



Deployment Manager is structured around a “site”, which is an instance of TDV. Deployment Manager provides the tools for you to easily define a TDV site, create bundles of resources, retain resource relationships, user and group permissions, and so on which you can then migrate to another site. The specific resources you can migrate include:


- Views, tables, procedures, transformations, and other resources defined in TDV
- Data source connectivity, schema, and catalog definitions

- Users and groups
- Domains
- Cache database settings

You can also define how the data sources and principals (users, groups, and domains) are mapped to the target site. You can use the source site data source connection properties and principal names by default, or map them to something different on the target site.

A single instance of Deployment Manager can connect to all TDV source and target sites with the correct credentials. If a site is not available or accessible (perhaps it is a remote production site, offline, or on a different network), you can define the site as offline and create a deployment plan that includes the offline site. A deployment plan can be exported and executed on a remote site at a later time.

An example of the resources that you might define for a source site is illustrated here:



After you've defined the source and target sites and the source site resource and principal bundles and their mappings, you can create a deployment plan. The deployment plan specifies the resource and principal bundles you want to migrate, resources that you want to remove from the target site (if any), and any procedure calls to execute that further customize the deployment.

See these sections for more overview information about Deployment Manager:

- [Limitations](#)
- [Basic Deployment Manager Concepts and Definitions](#)

- [User Roles and Workflows](#)

Limitations

None.

Basic Deployment Manager Concepts and Definitions

Deployment Manager uses these concepts and definitions:

Site	A TDV instance.
Resource	A view, table, procedure, transformation, or other resource defined in TDV.
Principal	A domain, user, or group.
Bundle	One or more resources of one of these types: resource or principal.
Resource Bundle	A collection of TDV resources that you want to deploy as a group. A resource bundle can contain views, tables, procedures, transformations, and so on that are defined in TDV.
Principal Bundle	A collection of TDV principals that you want to deploy as a group. A principal bundle can contain users, groups, and domains.
Mappings	Data source definitions which includes connection properties, catalogs, and schemas.
Deployment Plan	The source site, the target site, and the resource bundles to migrate and how they should be migrated. Can also include operations to remove target resources or execute procedure calls.

See these topics for more information about Deployment Manager:

- [Deployment Manager Architecture](#)
- [User Roles and Workflows](#)

User Roles and Workflows


Deployment Manager is a flexible tool that can be used by one person or many people. However, there are two main user roles in Deployment Manager:

- **TDV Developers**—Use TDV to create and publish the resources for the deployment. TDV Developers can then use Deployment Manager to create the Development site, add resources to it, and define a deployment plan to migrate resources to a Test environment to validate them.
- **System Administrators**—Execute the deployment plans to migrate changes across environments.


You might have a single person who does all of these tasks, or perhaps many developers and administrators who are responsible for pieces of a deployment project.

Note: Deployment Manager users must be user 'admin' or members of the 'admin' group.

The Deployment Manager workflow can be summarized in four basic steps: define the source and target sites, define the source site resources and mappings, define a deployment plan, and execute the deployment plan. The tasks in each of these steps is illustrated below:



The process for defining a deployment plan is illustrated in more detail here:



Starting Deployment Manager

When you start Deployment Manager, you need to know this information about the TDV instance connection information.

Deployment Manager runs on the browsers supported by TDV as listed in the *TDV Installation and Upgrade Guide*.

The TDV and Business Directory servers require a secure connection. So when you first connect a browser to any TDV web-based application, you might get a warning about connecting to an untrusted site.

Depending on your browser:

You might be asked to allow the connection process to continue.

You might want to configure the browser to trust the site so that warning messages no longer appear. For some site configurations this might require configuration of SSL connections for your entire TDV environment.

TDV provides a configuration option that can be tuned to override the “Select a certificate” popup. In TDV Studio, choose Administration > Configuration > Server > Communications > Want Client Authentication.

Default value for this setting is True. When True, Server will send client certificate request for SSL authentication.

You can set this configuration to false to avoid the popup "Select a Certificate" in Web Manager. TDV server must be restarted to take effect.

Note: If you have mutual authentication configured in any of your published web services, setting this configuration to false will disable it.

To start Deployment Manager

1. Start TDV, if necessary.
2. Use one of these methods to start Deployment Manager:

In Studio, select Launch Deployment Manager (Web) from the Administration menu. Use a supported browser to access a URL similar to one of the following.

```
http://<cishostname>:<portnumber>/deploy/
```

For example:

```
http://localhost:9400/deploy
```

The port you specify when using http: is automatically redirected to a secure port by adding 2 to the port number. So, you could directly type a secure URL as in this example for port 9400:

```
https://localhost:9402/deploy
```

3. Enter your user credentials and information for the TDV server to which you are connecting. For example:

Field	Notes	Example Value
domain	"composite" is the default domain. Click this value to change it to another domain.	composite
username	Deployment Manager users must be user 'admin' or members of the	admin

Field	Notes	Example Value
	'admin' group.	
	Note: All users share access to the same defined sites, plans, and so on in this version of Deployment Manager.	
password	Click the default value and enter a password.	@#\$\$%!-#

4. Click Login.

After starting Deployment Manager, see these topics:

- [Defining Sites](#)
- [Defining Deployment Plans](#)

Switching the Language of Deployment Manager UI

After logging into Deployment Manager, you can change the language of your user interface. To do that,

1. Choose the language of your choice in the Language drop down found in the top menu of the HOME page.
2. A dialog will be displayed to confirm reloading of the application. Click Yes to confirm your language selection or No if you do not wish to continue.
3. Once you confirm the language selection, the UI will be displayed in the language of your choice.

Defining Sites

A site in Deployment Manager is an instance of TDV. By adding sites, you are defining the instances of TDV that you want to manage with Deployment Manager. For example, you might have a TDV site where you develop TDV resources, a TDV site where you test those resources, and a TDV site that is the target production environment where the resources are accessed by your client applications.

While Deployment Manager can manage deployment for many TDV sites, it typically is run from a single TDV instance that is a development or test instance and not a production instance where performance might be impacted.

You must define both a source site and target site. Your target site does not need to be running at the time of its definition. This can be handy if it's not accessible or available. However, the target site needs to be running when accessing resources (during mapping, for example) or when executing a deployment plan via the PLANS user interface Execute Plan command.

Note: You can export a plan and then execute it on the offline site from a command line as described in [Exporting a Deployment Plan](#) and [Importing and Executing a Deployment Plan](#).

Topics that describe working with sites include:


- [Accessing Sites](#)
- [Adding a New Site](#)
- [Editing Site Properties](#)
- [Refreshing Site Resources](#)
- [Deleting a Site](#)
- [Defining Resource Bundles](#)
- [Defining Principal Bundles](#)
- [Defining Mappings](#)

Accessing Sites

You can access all sites defined in Deployment Manager in multiple ways.

To access sites

1. If necessary, start Deployment Manager.
See [Starting Deployment Manager](#).
2. Click the links shown below to access sites:



Adding a New Site

When you add a new site in Deployment Manager, you need to have the TDV connection information.

To add a new site

1. Click **SITES** at the top of the left pane.
2. Click **Add** at the bottom of the SITES panel on the left.
3. Type values for the following fields:

Field	Description of Value to Enter
Host	Enter the name of the machine where the TDV Server is installed. The name can contain letters, numbers, and hyphens.
Port	Enter the port number for the TDV Server repository. Default: 9400
Name	Type a name for this site within Deployment Manager. The name can contain letters, numbers, and underscores. Default: <host>_<port>

Field	Description of Value to Enter
Domain	Enter a domain name that already exists in TDV. See TDV Manager for this TDV instance for possible domains.
User	Enter the user name of a user that can access the information stored under the domain that you have specified.
Password	Enter the password for this user.
Annotation	Enter a description of the site that can be used to help you identify it given a number of similarly named sites. Optional.
Server Offline	Check this box if the TDV server defined for this site is not currently running or is a remote production site that is not accessible on the same network as the development or test site. If the server becomes available, you can change this setting by editing the site on the General tab. See Editing Site Properties .

4. Click **Save** to create the new site, or **Cancel** to quit without saving the information.

For a source site, you are ready to define the resources and principal bundles you want to migrate to another site, and how the data sources for those resources are mapped. After you have created a source site, see these sections for how to perform these tasks:

- [Defining Resource Bundles](#)
- [Defining Principal Bundles](#)
- [Defining Mappings](#)

Editing Site Properties

You can edit the properties for a TDV site as defined in Deployment Manager.

To edit site properties

1. Click **SITES**.
2. Select a site.
3. Click the **General** tab.

4. Click **Edit**.

The screenshot shows the 'localhost_9410' configuration window in Deployment Manager. The window has a green header bar with the title 'localhost_9410' and a red note: '* Modifications made in this view are immediately saved to the server.' Below the header is a tabbed interface with five tabs: 'Resource Bundles', 'Principal Bundles', 'Resource Mappings', 'Principal Mappings', and 'General (...)' (which is selected). The 'General' tab contains the following fields:

- Host: localhost
- Port: 9410
- Domain: composite
- User: admin
- Password: ••••••••
- Annotation: QA's test site
- Server Offline: ☐

At the bottom of the window is a green bar with two buttons: 'Save' (with a checkmark icon) and 'Cancel' (with a close icon).

5. Change any of the properties.

If you navigate away from the General tab before saving, your changes will be lost.

6. Click **Save**.7. If the server is online, click **Test Connection** to make sure that the site is accessible.


Refreshing Site Resources

You can refresh the TDV site resources in Deployment Manager. Deployment Manager updates the list with any updates to sites that might have been made by another user or in another browser.

Note: A site does not need to be ONLINE or running to be refreshed.

To refresh all existing site resources

1. If necessary, start Deployment Manager.
2. Click **SITES**.



Note: The green or gray indicators to the right of a site only indicate if the site is configured to be ONLINE or OFFLINE in Deployment Manager. This is set using the Server Offline option on the General tab. It does not indicate if the server is actually running or not.

3. Beneath the SITES panel, click the **Refresh** button.

All sites are refreshed. The site does not need to be online for the information Deployment Manager stores about it to be refreshed.

Deleting a Site

You can delete a TDV site from Deployment Manager. Deleting a site removes everything defined for the site in Deployment Manager including sites, bundles, mappings, and deployment plans. However, the TDV instance itself is not affected; only the definitions within the Deployment Manager are affected.

To delete a site

1. In Deployment Manager, select **SITES**.
2. Select a site.
3. At the bottom of the site list panel, click **Delete**.

4. Click **Yes** to confirm that you want to delete this site and all of its definitions.

Defining Resource Bundles

Deployment Manager lets you organize site resources into one or more resource bundles and then deploy the resources in a bundle as a unified set. For each resource or resource bundle, you can also choose to include dependencies. These topics describe how to work with resource bundles:

- [Creating a Resource Bundle](#)
- [Adding Resources to a Bundle](#)
- [Previewing the Resources in a Bundle](#)
- [Excluding and Including Resources in a Bundle](#)
- [Viewing the Details of a Resource](#)
- [Setting Resource Dependencies](#)
- [Copying a Resource Bundle](#)
- [Removing Resources from a Resource Bundle](#)
- [Deleting a Resource Bundle](#)

Creating a Resource Bundle

A resource bundle contains a collection of resources like views, procedures, tables, data sources, and their containers. Resource bundles allow you to migrate resources as a group to another TDV when you execute a deployment. You need to first add a resource bundle and then add the site resources to it.

You can add entire containers (folders) of resources or individual resources to a bundle, but there are some limitations.

Limitations for bundle resources:

- A resource bundle can contain resources from only one TDV site.
- You can migrate only one resource bundle per deployment plan.
- Resources do not need to be locked.
- You can add an entire data source but you cannot add only a single table or file from a data source.

To add a resource bundle

1. Click **SITES**.
2. Select a site.
3. On the **Resource Bundles** tab, click the **+** symbol at the bottom of the Bundle Definitions pane to Add Resource Bundle.
4. Type a name for your new bundle. Bundle names can contain letters, numbers, and underscores.
5. Optionally, enter a description of the bundle contents in the Annotation field.
6. Click **Add Bundle**.

You should see your newly added bundle in the Bundle Definitions panel.

You can now add resources to the bundle as described in [Adding Resources to a Bundle](#).

Adding Resources to a Bundle

To add resources to a bundle


7. Open the Resource Bundles tab.
8. If necessary, click Refresh at the bottom of the Available Resources panel to make sure the Available Resources list is up-to-date.
9. Add resources to a bundle. The Available Resources pane contains all resources in the site that can be added to the resource bundles. You can:

From Available Resources...	Description
Select and drag a container	All of the objects beneath it are moved to the bundle.
Select and drag a specific resource	Only the resource is moved to the bundle.
Select an available resource or container	Click Add Resources to, then choose the bundle to which you want to add the resources.

Notes:

- Do not add duplicate resources. If you add duplicate resources, the bundle will become impacted.

This example shows two resource bundles; one has a container with many resources in it and the other has a single resource:



Note: The number of resources in the bundle is shown next to the bundle name. This number includes both explicitly added resources and dependent resources. Dependent resources are not displayed under the bundle but are displayed in the preview. See [Previewing the Resources in a Bundle](#).

10. Optionally, perform these operations to review and customize the bundles for your deployment:

- [Excluding and Including Resources in a Bundle](#)
- [Previewing the Resources in a Bundle](#)
- [Viewing the Details of a Resource](#)
- [Setting Resource Dependencies](#)
- [Copying a Resource Bundle](#)
- [Removing Resources from a Resource Bundle](#)
- [Deleting a Resource Bundle](#)

Excluding and Including Resources in a Bundle

By default, when you explicitly add a resource or resource container to a bundle, all resources are included. You might want to exclude some of them, or include some that are contained in a resource that you excluded. Excluding a resource does not mean that

the resource is deleted from the bundle; it means that it will not be migrated with the bundle in a plan.

These rules apply when excluding resources:

You can exclude...	You cannot exclude...
Resources in a folder including:	A root folder in the bundle
- data sources	A resource at the root of the bundle
- views	A schema, table, or element within a data source.
- procedures	
- models	
...and so on	

If you exclude a resource on which another included resource is dependent, that resource is still included. The exception to this is if you set the dependency option to No. See [Setting Resource Dependencies](#).

While editing the resource bundles, you can also perform these tasks:

- [Previewing the Resources in a Bundle](#).
- [Viewing the Details of a Resource](#) .
- [Setting Resource Dependencies](#).

To exclude and include resources in a bundle

11. Open the **Resource Bundles** tab.

12. If necessary, add resources to your bundle as described in [Adding Resources to a Bundle](#).

13. In Bundle Definitions, expand the folder that contains the resources to be excluded.

14. Select one or more resources.

Note: You can use Ctrl-click and Shift-click to select multiple resources at one time to exclude.

15. Click **Exclude**.

Deployment Manager displays an excluded resource in dimmed text as shown in this example.




16. Optionally, to include a resource that has been excluded, select the excluded resource and click **Include**.

After you have finished excluding resources, you can preview the bundle to see what resources would be migrated with this bundle (see [Previewing the Resources in a Bundle](#)). Only included resources and their dependencies (depending on the Inherit setting) are included in the preview.

Previewing the Resources in a Bundle

The Bundle Definitions pane shows the bundle definition: the bundles, the resources or resource containers that have been directly added to the bundle, and the resources that are excluded from the bundle, if any. A bundle preview lets you see only the resources that will be migrated with the bundle at plan execution time. The bundle preview shows the included resources and their dependencies in blue text as shown in this example:



Inclusion of the dependent resources with a bundle is determined by the bundle and resource Include Dependencies properties. To see how you can choose whether or not to include the dependent resources with a bundle, see [Setting Resource Dependencies](#).

To preview the resources in a bundle

17. Open the **Resource Bundles** tab.
18. Under Bundle Definitions, click **Refresh** to make sure that the resources are up-to-date.
19. Select a bundle.
20. Click **Preview**.
21. If necessary, expand the directories to view the bundle resources.
22. Click **OK** to close the preview.

Viewing the Details of a Resource

All of the objects displayed for available resources and resource bundles have properties that can be viewed and in some cases, edited.

To view details of the available and resource bundle resources

23. Click **SITES**.
24. Select a site.
25. On the **Resource Bundles** tab, select any object under Available Resources or Bundle Definitions.
26. Under Property View, review the details in the fields for the selected item. The properties vary depending on the object type selected:

Available Resource Properties	Bundle Definition Object	Properties
Name Path Type Orig Creation Date Owner Last Modified Date Last Modified User	Bundle	Name Annotation Include Dependencies Orig Creation Date Owner Last Modified Date Last Modified User
	Root folder	Name

Available Resource Properties	Bundle Definition Object	Properties
		Path Type Include Dependencies
	Folders and resources	Notice if excluded Name Path Type

Note: Properties in bold can be edited.

27. Optionally, edit the Name, Annotation, or dependencies setting for the object.

Setting Resource Dependencies

Dependencies can be specified on a bundle level or the level just below the bundle which can be a container or a resource. By default, all resources added directly to a bundle, which can be a container or a resource, have the same dependencies as the bundle. However, you can override the dependencies settings for a bundle or any container or resource added at the root level in the bundle.

To set resource dependencies

28. Click **SITES**.

29. Select a site.

30. Select the **Resource Bundles** tab.

31. In Bundle Definitions, select a resource bundle or a container or resource at the root level of a bundle.

32. In the Property View column, select the Include Dependencies setting:

Option	Description
Yes	Include dependencies for this resource. This is the default for bundles.

Option	Description
No	Do not include dependencies for this resource.
Inherit (resource or folder only)	Defer to the bundle's dependency setting. This is the default for containers and resources in bundles.

Copying a Resource Bundle

You can create a duplicate resource bundle that you can use as a starting point for another resource bundle, for example.

To copy a resource bundle

33. Click **SITES**.
34. Select a site.
35. Select the **Resource Bundles** tab.
36. In Bundle Definitions, select a resource bundle.
37. Click the **Copy Bundle** button.
Deployment Manager creates a duplicate bundle and appends “_copy” to the bundle name.
38. Optionally, under Property View, rename the bundle.

Removing Resources from a Resource Bundle

You can remove a resource container (including its resources) or a resource at the root of a bundle using the steps below. The resources are still listed under Available Resources and can be reinserted.

To remove a resource from a resource bundle

39. Click **SITES**.
40. Select a site.
41. Select the **Resource Bundles** tab.
42. Select a container or a resource in a resource bundle.

43. Click **Remove**.

Deleting a Resource Bundle

To delete a resource bundle

44. Click **SITES**.
45. Select a site.
46. Select the **Resource Bundles** tab.
47. Select a resource bundle.
48. Click the **Delete Bundle** button.
49. Click **Yes** to confirm that you want to delete this bundle.

Defining Principal Bundles

Principal bundles are containers that you can use to migrate domains, users, and groups from one TDV site to another.

- [Creating a Principal Bundle](#)
- [Deploying Privileges](#)
- [Adding Principals to a Bundle](#)
- [Viewing the Details of a Principal Resource](#)
- [Copying a Principal Bundle](#)
- [Removing a Principal from a Principal Bundle](#)
- [Deleting a Principal Bundle](#)

Creating a Principal Bundle

A principal bundle contains a collection of domains, groups, or users (that is, principals) that you can migrate to another TDV as a unit when you execute a deployment. You can create a principal bundle and then add domains, groups and users to the principal bundle. Because a single deployment plan can contain multiple principal bundles, you can organize principals (domains, groups, and users) into multiple principal bundles to suit your needs.

When you add a group to a bundle, you are adding the definition of that group but not the users it contains. You must include the group and the group's users in the same principal bundle.

Note: The domains, users, and groups that are defined for TDV by default for the composite domain do not appear as available principals in Deployment Manager. These include the composite and admin domains, the admin and all groups, and the admin, anonymous, monitor, nobody, and system users. Only domains, users, and groups that have been added to a TDV instance appear under Available Principals.

To create a principal bundle

1. Click **SITES**.
2. Select a site.
3. Select the **Principal Bundles** tab.
4. Under Bundle Definitions, click the **+** symbol to add a principal bundle.
5. Type a name for your new bundle. Optionally, enter a description of the bundle contents in the Annotation field.
6. Click **Add Bundle**.

You should see your newly added bundle in the Bundle Definitions column.

Deploying Privileges

To deploy privileges, the principal bundle must have a principal mapping. Setting a value for the principal bundle filter is also recommended.

Adding Principals to a Bundle

To add principals to a bundle


7. Open the Principal Bundles tab.
8. If necessary, under Available Principals, click **Refresh** to make sure the principals list is up-to-date.
9. Select and drag the domain, group, and user principals from Available Principals to the bundle.

The Available Principals column contains all principals in the site that can be added to the principal bundles. To add principals to a bundle:

You can...	Description
Select and drag a container	All of the objects beneath it are moved to the bundle.
Select and drag a specific principal	Only the principal is moved to the bundle.
Select an available principal or container	Click Add Principle to..., then choose the bundle to which you want to add the principals.

Note: When adding a group, be sure to also explicitly add the group's users to the bundle.

Deployment Manager displays principals in each bundle and bundle container.



10. Optionally, drill into the bundles and their containers to view the resources, rename the bundles or their annotations, or delete bundles or containers to customize the bundles for your deployment.

Viewing the Details of a Principal Resource

All of the objects displayed for available principals and principal bundles have properties that can be viewed and in some cases, edited.

To view details of the available and principal resources

11. Click **SITES**.
12. Select a site.
13. Select the **Principal Bundles** tab.
14. Select any object under Available Principals or Bundle Definitions.
15. Under *Property View*, review the details in the fields for the selected item. The properties vary depending on the object type selected:

Principal Properties	Principal Bundle Properties
Name	Name (editable)
Path	Annotation (editable)
Type (DOMAIN, GROUP, or USER)	Original Creation Date
	Owner
	Last Modified Date
	Last Modified User

16. Optionally, edit the Name or Annotation for bundles.

Copying a Principal Bundle

You can create a duplicate principal bundle that you can use as a starting point for another bundle, for example.

To copy a principal bundle

17. Click **SITES**.
18. Select a site.
19. Select the **Principal Bundles** tab.
20. Under Bundle Definitions, select a bundle.
21. Click the Copy Bundle button.
Deployment Manager creates a duplicate bundle and appends “_copy” to the bundle name.

22. Optionally, under *Property View*, rename the bundle.

Removing a Principal from a Principal Bundle

You can remove a principal container and its resources or a principal in a bundle using the steps below. The principals are still listed under Available Principals.

To remove a principal in a principal bundle

23. Click **SITES**.
24. Select a site.
25. Select the **Principal Bundles** tab.
26. Under Bundle Definitions, select a container or a principal in a principal bundle.
27. Click **Remove**.

Deleting a Principal Bundle

To delete a principal bundle

28. Click **SITES**.
29. Select a site.
30. Select the **Principal Bundles** tab.
31. Under Bundle Definitions, select a principal bundle.
32. Click the **Delete Bundle** button.
33. Click **Yes** to confirm that you want to delete this bundle.

Defining Mappings

A mapping defines how data source connection properties and principal definitions are mapped to the target site when related bundles are migrated. Deployment Manager uses the data source connection properties and principal definitions on the source site by default. However, you can define the mapping properties from scratch or use the properties from any defined site as the starting point and then edit those properties.

For data source mappings, you can define these properties:

- Data source properties (host, IP addresses, etc.)
- Data source logins (encrypted)
- Data source schema and catalogs (for some data sources like Oracle)

For principal mappings, you define the mapped name to define how this mapping should be named on the target site.

Requirements

If you have a database that supports enabled database links, you must disable database links prior to defining mappings in Deployment Manager. See [Disabling Database Links for Mapping](#) for more information.

Note: Accessing resources during mapping requires that the source and target TDV instances are online.

See these topics for how to define the mappings:

- [Editing the Resource Mappings for a Site](#)
- [Changing the Mapping Target Values Site](#)
- [Editing the Principal Mappings for a Site](#)
- [Removing Mapping Definitions](#)

Editing the Resource Mappings for a Site


The resource mappings define the connection properties for a data sources on your target site. Defining or editing resource mappings is optional. By default, the data source resources are mapped to the same properties that they have on the source site. You only need to define or edit the resource mappings if the data sources have different connection properties on the target site.

Note: If you are using caches, see [Caching and Deployment Manager](#) for information about mapping the cache data sources.

To edit the resource mappings for a site

1. Click **SITES**.
2. Select a site.
3. Click the **Resource Mappings** tab.

Deployment Manager finds all data sources on the source site that contain resources that might be migrated to the target server. In this example, the source site has three data sources:



You can optionally map the connection properties on the source site to something different on the target site. If you do not edit the mappings, the current settings on the source site are used.

4. Click **Choose Target Site** and select a target site for the mapping.

The target site is the site to which you plan to migrate the bundle. Choosing a target site now makes it easy for you to populate the Target Values column.

Note: When you create a plan, you can migrate the bundle to a different site but you might want to adjust the mappings prior to plan execution.

5. Click a data source to expand it and see the three property groups: Basic, Advanced, and for some data sources, Containers.


The fields displayed in the Basic Properties, Advanced Properties, and optional Container groups depend on the data source type (such as relational data sources or XML data sources) and vendor (such as PostgreSQL, Oracle, SQL Server, etc.). See the sections below for information about the displayed fields:

- [Relational Data Sources](#)
- [XML Data Sources](#)

For complete documentation on all of the data source properties for each data source type and vendor supported by TDV, see the *TDV User Guide*.

Note: You can view the properties for a data source by opening it in TDV Studio. The location of the properties on the Basic and Advanced tabs and the property names displayed in Deployment Manager might be slightly different than in Studio. Also, some fields are editable in Deployment Manager but not in Studio. Where they differ, Studio is the best authority.

6. Click a group to review its properties as shown below for the Basic group.



7. Optionally, populate the Target Values from either the source site or target site:

- Click **Populate from Source** to fill in the properties from the source site which you can then edit. These are the default properties that are used if you do not edit the mappings.
- Click **Populate from Target** to populate the fields using values from another TDV site which you can then edit. Deployment Manager displays a dialog for you to choose a site defined in Deployment Manager, then choose the data source with the properties to populate the fields. You should choose a compatible data source. That is, if your target data source is a relational data source, you should choose a relational data source to populate the fields.

Note: The target site must be online to Populate from Target. You can define mappings for offline sites, but you need to enter them manually or Populate from Source and edit the values.

8. Edit a property:

Double-click the row.

Type text directly into the field or choose an option.

Click **Save** to save the new value.

9. Repeat editing the mapping information for each data source, as necessary.

Relational Data Sources

Relational data sources have three possible property groups: Basic, Advanced, and Container (for data sources that support catalog and schema definitions). The properties shown in Deployment Manager for an Oracle data source are listed below.

Basic Properties (Oracle Example)

Property	Value Type	Notes
urlIP	text	
urlPort	text	
urlDatabaseName	text	
login	text	
password	text	Appears as asterisks.
connProperties	name/value	You can specify property name-value pairs to pass to the JDBC data source. You can add whatever connection property/value pairs are appropriate for your data source. For example, the Oracle data source might have these connection properties: <pre>processEscapes: true disableDefinecolumnType: true</pre>

SetBigStringTryClob: true
AccumulateBatchResult:
true

TDV does not validate
property names. Some data
source adapters ignore
invalid property names or
values; others return an
error.

connPoolMinSize	numeric text
connPoolMaxSize	numeric text
connPoolTimeout	numeric text
connStaleTimeout	numeric text
execTimeout	numeric text
maxSourceSideCardinalityForSJ	text
minTargetToSourceRatioForSJ	text
maxNumberForOrSyntax	numeric text

Advanced Properties (Oracle Example)

Property	Type	Notes
persistPassword	true false none	
isPassThrough	Enabled Disab led	
txnIsolationLevel	Read Committed Serializable	

urlPatternStr	text	Use the format: jdbc:postgresql://<HOST>:<PORT>/ <DATABASE_NAME>
connValidateQuery	text	Example: SELECT 1
commitOnFetchDone	true false none	
connCheckOutProcedure	text	
supportsStarSchema	true false none	
nativeDataLoadingEnabled	true false none	
collationSensitive	true false none	
enablePassThroughPrepared Statements	true false none	
introspectUsingDBA_Views	true false none	
introspectProceduresEnabled	true false none	
authentication	BASIC KERBE ROS	
ticketCache	text	
includeInvalidObjects	true false none	
useLoginCertEncryption	true false none	
supportsDataship	true false none	
supportsDatashipAsTarget	true false none	
LowerBoundForDataShip	numeric text	
UpperBoundForDataShip	numeric text	

destinationSchema	text	
dataShipTempTablePrefix	text	
supportsDBLink	true false none	If set to true, you must meet the following requirements for running data source mapping across sites successfully: • Valid connection credentials for the target site. For example, you must be able to connect to the data source using the supplied port number from the mapping definition. • The DBLink information must be valid. • The Database Links defined including the Database Link Name and the Path of data source must exist in the target site before executing the mapping action. Otherwise, disable this option before mapping across sites. See Disabling Database Links for Mapping .
DBLinkList	name/value	Add a row for each database link, then enter the database link name and the path of the data source.

Container Properties (Oracle Example)

Containers are the names of catalogs and schema definitions. Some data sources do not have the Container group if they do not support catalogs and schema.

Catalog or Schema	Type	Notes

Name	
RQAN1	text
QAN	text
SYSTEM	text
DBSNMP	text
RQAN	text

XML Data Sources

Basic Properties (XML Example)

Property	Type	Notes
root	text	
url	text	Example format: file:///C:/engineering/test.xml
schemaLocation	text	Example format: http://www.compositesw.com/services/webservices/system/admin/resource file:///C:test.xsd
noNamespaceSchema Location	text	

Advanced Properties (XML Example)

Property	Type	Notes

local	true false none	
charset	text drop-down	<auto-detect> - Default Cp1250-Cp1257 iso-8859-1 us-ascii utf-8, 16, 16be, 16le windows-1250 to windows-1257
filters	text	File name filters to restrict the files included, such as *.xml to select only files of type XML. Separate multiple filters with commas.

Caching and Deployment Manager

If you are using caching, keep these things in mind when defining mappings:

- By default, the source site caching definitions for target files or data sources are used when you migrate resources to a target site.
- If you are using single-file or multi-file caching and have specified file or data source caching information on the source site, these data sources can be mapped to a different location on the target site.
- If you are using file-based caching and want to map the cache to a different location on the target site, you must create the folder and cache files on the target source prior to executing the deployment plan.

When you initially create a file-based cache data source, you need to configure the caching status table (cache_status) and tracking table (cache_tracking). The cache_status and cache_tracking files are generated in the storage directory of the file cache data source.

You can map the file caching storage directory to a different location on the target server, but you need to copy all folders and files from the source directory to the target directory manually. Otherwise, the mapping will fail.

For example, if you have the following cache storage directories:

- Source cache storage directory: D:\dm\ds\filecacheds
- Target cache storage directory: D:\temp

You need to copy cache_status and cache_tracking files to D:\temp manually before doing the Deployment Manager mapping operation.

When you migrate resources in a deployment plan, you can specify how caches for new or existing objects should be handled. You can choose to create or re-create the cache tables or not. These options apply to data source caches only; they are not applicable to file-based caches. See [Defining the Resource Bundle to Migrate](#).

Changing the Mapping Target Values Site

Changing the resource mappings for a specific environment can be done quickly from the Resource Mapping tab.

To change the target value for your resource mapping

10. From the Resource Mapping tab, click on the site name next to the Target Value column label.

Resource Bundles	Principal Bundles	Resource Mappings	Principal Mappings	General
Resource Name		Source Value	Target Value	localhost_9400
/shared/examples/ds_inventory *				localhost_9500
				localhost_9410
Basic				
urlIP		localhost	localhost	

11. Select a different site from the list.

Disabling Database Links for Mapping

For data sources that support database links, you might need to disable database links support prior to mapping the data source in Deployment Manager. See the supportsDBLink Advanced Property in the [Relational Data Sources](#) section for the requirements.

To disable database links

12. Open TDV Studio.
13. Open the data source.
14. Click the **Advanced** tab.
15. If necessary, clear the database link check box. The database link option is labeled differently for example:

Data Source Example	Database Link Option Example
Oracle	Enable Oracle Database Link
PostgreSQL	Enable PostgreSQL dblink
Vertica	Enable Export To Another Vertica Database
Sybase IQ	Enable Sybase IQ SQL Location

Editing the Principal Mappings for a Site

The principal mappings define the names you want to give the selected principals on the target site. Editing principal mappings is optional; by default, the principals are mapped to the same names as they have on the source site, so you only need to edit principal mappings if the principals have different names on the target site.


To edit the principal mappings for a site

16. Click **SITES**.

17. Select a site.

18. Click the **Principal Mappings** tab.

A list of the existing principals—groups and users—for the TDV site is displayed as shown in this example:



19. Click **Choose Target Site** and select a target site for the mapping, then click **OK**.

20. Double-click a row to enter a new target value.

Deployment Manager displays a text box.

21. Type the target name you want to give this principal on the target site.

The target name must include the domain and group or user as a prefix to the target name using this syntax:

`<domain>/<group|user>/<target_name>`

as in:


`composite/group/developers`

For example, if you want to migrate this principal—`composite/group/developers`—to a different domain on the target server and also give it a new name, you might specify:

`my_domain/group/Product_Developers`

22. Click **Save**.

In this example, the source site name for the group is “developers”, but the name for this principal group will be “Product_Developers” on the target site. Also, two of the users are each mapped to a different name. All are mapped to a different domain.



Removing Mapping Definitions

You can easily remove any individual resource or principal mapping definitions that you have made. For resource mappings, you can delete all mapping definitions for a data source.

To remove an individual mapping definition

23. Click **SITES**.
24. Select a site.
25. Select the **Resource Mappings** or **Principal Mappings** tab.
26. Double-click the row containing the target value you want to remove.
27. Click **Delete**.

To remove all resource mapping definitions for a data source

28. Click **SITES**.
29. Select a site.
30. Select the **Resource Mappings** tab.
31. Select the data source.
32. Click **Delete Mapping** (the “-” button to the far right of the data source name).

Defining Deployment Plans

A deployment plan defines the source and target site and the sequence of operations that occur during deployment.

A deployment plan can specify one or more of these operation types:

- Migrate resources (only one resource bundle can be migrated per plan)
- Migrate principals
- Remove resources or principals from the target
- Procedure calls to run precoded SQL scripts

The migrate operations specify the resources and principals to migrate from one site to the other and how those resources should be migrated. For example, you can specify to

overwrite any existing resources, rebuild or create new caches, or configure the permissions for the resources. You can also remove resources from the target site and specify procedure calls to perform specific procedures when the deployment plan is executed. Collectively, the migrate, remove, and procedure call operations defined in a plan are performed when you execute the plan.

Note: Each deployment plan involves one source site and one target site. You can create multiple plans if you want to deploy resources from multiple source sites to the same target site, for example.

Topics that describe working with deployment plans include:

- [Accessing Plans](#)
- [Creating a New Deployment Plan](#)
- [Defining the Resource Bundle to Migrate](#)
- [Defining the Principal Bundles to Migrate](#)
- [Removing Resources from a Target Site](#)
- [Removing Principals from a Target Site](#)
- [Extending Deployments with Procedure Call Operations](#)
- [Previewing a Deployment Plan](#)
- [Editing Deployment Plan Properties](#)
- [Refreshing Deployment Plans](#)
- [Exporting a Deployment Plan](#)
- [Deleting a Deployment Plan](#)

Accessing Plans

You can access all plans defined in Deployment Manager in multiple ways.

To access plans

- At login, click **MANAGE PLANS** or **CREATE NEW PLAN**.
or
- Click **PLANS** at the top of the leftmost pane.

Creating a New Deployment Plan

A new deployment plan defines the source and target sites.

Note: The target site does not need to be online during plan creation, but it does need to be online at plan execution time.

To create a new deployment plan

1. Make sure that the source and target sites are defined in Deployment Manager and that the resources you want to migrate in the source site are defined.
2. Click **PLANS**.
3. At the bottom of the PLANS panel, click **Add**.
4. Type a name for your plan.
5. Optionally, enter a description of the plan.
6. Select the source site.
7. Select the target site.
8. Click **OK**.

After you create the plan, you can perform the following tasks:

- [Defining the Resource Bundle to Migrate](#)
- [Defining the Principal Bundles to Migrate](#)
- [Removing Resources from a Target Site](#)
- [Removing Principals from a Target Site](#)
- [Extending Deployments with Procedure Call Operations](#)

Defining the Resource Bundle to Migrate

Each deployment plan can contain one operation that migrates a resource bundle to a target site.

You control how you want the migration to happen. For example, you can specify things like resource retention policies, how to filter resources, and how to handle caches for the resources being migrated. This section describes how to define the operation to migrate a resource bundle.

Note: Each deployment plan can include only one resource bundle. If you want to migrate multiple resources in a single plan, you need to add them to one resource bundle. See [Defining Resource Bundles](#).

To define the resources to migrate in a deployment plan


1. Click **PLANS**.
2. Select a plan.
If you need to create a new plan, see [Creating a New Deployment Plan](#).
3. Select the **Operations** tab.
4. Click **Edit**.
5. Click **Add Operation**.
6. Select **Migrate Resources Operation** to add a resource bundle to the deployment plan.
7. In the **Select Bundle from <source_site>** dialog, select the resource bundle you want to include in the plan.
Note: Only one resource bundle can be selected per deployment plan.
Deployment Manager displays the bundle and its resource tree.
8. Expand the bundle resource tree to see the first level of resources for this bundle.
9. Optionally, click **Preview Bundle** to view a complete list of all resources that are included and excluded in this bundle. All resources in the main bundle resource are included unless they shown as dimmed text.
Click **OK** to close the preview.
10. Click **OK**.
Deployment Manager adds the operation to the list of operations.
11. Optionally, edit the deployment plan options for this bundle as needed and click **Save** after changing an option.

Option	To edit, double-click the option row, then...	Required?
Bundle	Review the current bundle resources. Optionally, select a different bundle to migrate.	Yes

Option	To edit, double-click the option row, then...	Required?
	If a resource has been excluded from the bundle, but another object in the bundle has a dependency on the excluded resource, the excluded resource is migrated anyway.	
Retention Policy	Select the resource retention policy from the following list of values: • Keep Non-Colliding (default)—If the target site folder to which the resources are being migrated contains any resources, they are kept. Only resources by the same name are overwritten. • Keep None—If the target site folder to which the resources are being migrated contains any resources, they are deleted from the target site.	Yes
Target Location	If the target TDV is online, double-click the row and browse for the desired location. If the target TDV is offline, type a location in the target site where you want the bundle to be migrated. If you do not specify a target location, the default target location specified on the General tab is used. If no default target location is specified on the General tab, then the resources are migrated to the same location they had on the source site. Format: /<directory>/<directory>... Example: /shared/examples	Yes
Target Default Owner	Enter the user on the target site who will own these resources. This user must exist on the target site or must be migrated to the target site in a principal bundle in the same deployment plan. Also, this user will have the same resource privileges as they have on the source site. If you do not specify an owner, the owner will be the same as the owner on the source site.	No

Option	To edit, double-click the option row, then...	Required?
	Format: <domain>/<username> Example: composite/Mary	
Principal Bundle Filter	Select a principal bundle to use as a filter in the dialog box. The resource privileges for the principals in that bundle are migrated to the same principals in the target. If the principals do not exist on the target site, you should also migrate a principal bundle containing them. To clear a previous selection for the Principal Bundle Filter, select <clear selection>.	No
Caches for New Objects	Select the cache retention policy from the following list of values: • Create Cache Tables (default) • Do Not Create Cache Tables This option applies to data source caches only; it is not applicable to file-based caches. See Caching and Deployment Manager for more information about caches and Deployment Manager.	Yes
Caches for Existing Objects	Select the cache retention policy from the following list of values: • Re-Create Cache Tables (default) • Do Not Re-Create Cache Tables This option applies to data source caches only; it is not applicable to file-based caches. See Caching and Deployment Manager for more information about caches and Deployment Manager.	Yes

An example Migrate Resources operation is shown here:



12. Click **Save** at the bottom of the Operations tab when you've completed editing the plan properties.

Defining the Principal Bundles to Migrate

Each deployment plan can contain one or more operations that migrate principal bundles from the source to the target site. This section describes how to define principal migration operations.

Note: A deployment plan can define multiple principal bundles to migrate.

To define the principal resources to migrate in a deployment plan

1. From **PLANS**, select a plan.
If you need to create a new plan, see [Creating a New Deployment Plan](#).
2. Select the **Operations** tab.
3. Click **Edit**.
4. Click **Add Operation**.
5. Select **Migrate Principals Operation** to add a principal bundle to the deployment plan.
6. Select the principal bundle to migrate with this plan.

7. Click **OK**.

Deployment Manager adds the operation to the list of operations. An example Migrate Principals operation is shown here:

Migrate Principals	
Principal Bundle	Developers

8. Click **Save** on the Operations tab.

9. Repeat steps 3-8 above for each principal bundle that you want to add to the deployment plan.

Removing Resources from a Target Site

You can use a deployment plan to remove specified resources from the target site. Removing resources from the target site does not remove them from the source site.

A couple of notes:

- The resources must exist on the target site.
- You cannot remove resources unless the target server is online.

To remove resources from a target location

1. From **PLANS**, select a plan.
2. Click **Edit**.
3. Click **Add Operation**.
4. Select **Remove Resources Operation**.

Deployment Manager adds a Remove Resources operation to the Operations tab.

5. Edit the following options as needed.

Property	Do this...
Target Location	6. Double-click the row and in the dialog box, navigate to the

Property	Do this...
	location within the target site that contains the resource you want to remove. 7. Select the resource and click OK.
Target Type	• Double click the row and select the resource type: – CONTAINER – DATA_SOURCE – DEFINITION_SET – EXTENSION – LINK – MODEL – POLICY – PROCEDURE – TABLE – TREE – TRIGGER EXTENSION is listed but is not supported. Note: Choosing a Target Type is optional unless there is more than one resource with the same name but a different type at the Target Location.

8. Click **Save** on the Operations tab.

Removing Principals from a Target Site

You can use a deployment plan to remove existing principals from the target site. Removing them on the target site does not remove them on the source site.

A couple of notes:

- The principals must exist on the target site.
- You cannot remove principals unless the target server is online.

To remove principals from a target location

1. Click **PLANS** and select a plan.
2. Click **Edit**.
3. Click **Add Operation**.
4. Select **Remove Principals Operation**.

Deployment Manager adds a Remove Principals operation to the Operations tab.

5. Double-click the Principal Bundle row and select the principal bundle on the target site to remove.
6. Click **OK**.
7. Click **Save** on the Operations tab.

Extending Deployments with Procedure Call Operations

You can include one or more precoded procedures to customize your deployment. For example, you might want to configure the target server in some way such as enabling or disabling triggers or changing the server configuration. You do this by adding a procedure call operation to the deployment plan.

A procedure call operation requires a precoded, existing SQL script that can be run as part of the deployment plan. The procedure can contain any API call and other code that TDV supports but it cannot require any parameters. The procedure must reside on the TDV server where the Deployment Manager server is running.

When you add a procedure to a deployment plan, you need to specify the procedure name and where you want the procedure to execute: the source server, the target server, or the Deployment Manager server.

Note: Plan operations can't be reordered, so be sure that any necessary procedure call operation is added before other operations in the plan that might require it.

To specify a procedure call operation with a deployment plan

1. Click **PLANS** and select a plan.
If you need to create a new plan, see [Creating a New Deployment Plan](#).
2. Click **Edit**.

3. On the Operations tab, click **Add Operation**.
4. Select **Procedure Call Operation** from the menu.
5. In the dialog that lists all procedures found on the TDV source, select a procedure and click **OK**.

Deployment Manager adds the Procedure Call operation to the Operations list with these properties:

Property	Description
Procedure Path	Displays the path to this procedure in the source server.
Execution Environment	• DM—Deployment Manager server • SOURCE—Source site server • TARGET—Target site server (default)

6. Edit the Execution Environment if necessary. To edit this option, double-click the Execution Environment row, then select a different value.
7. Click **Save**.


Previewing a Deployment Plan

You can preview a deployment plan to review the source and target sites, the number of operations that will be executed, the type of each operation and its configuration options, and for resource operations, the full path name of the source resources.

To preview a plan

1. Click **PLANS**.
2. Select a plan.
3. Select the **Preview** tab.
4. Review the information that displays.

The following is an example of the preview of a deployment plan:



- Optionally, click **SHOW DELTA PREVIEW** to display a list of changes on the source site since this plan was last executed. See [Viewing Source Site Updates Since the Last Plan Execution](#) for more information.

Note: SHOW DELTA PREVIEW is only available when the Migrate Resources Retention Policy property is set to Keep Non-Colliding.

Editing Deployment Plan Properties


You can edit two of the general properties for a deployment plan:

- The default target location that will be used if no specific location is specified for the resources.
- The annotation for the deployment plan which can give information about the plan like its purpose, contents, or use.

To edit deployment plan properties

- Click **PLANS**.
- Select a plan.

3. Click the **General** tab.
4. Click **Edit**.



5. Click the button to browse for a target location for the resources in this plan.
6. Optionally, add or edit the text in the Annotation field.

If you navigate away from the General tab before saving, your changes will be lost.

7. Click **Save**.

Refreshing Deployment Plans

You can refresh the list of deployment plans. Deployment Manager updates the list with any new or updates to plans that might have been made by another user or in another browser.

To refresh the deployment plans

1. In Deployment Manager, select **PLANS**.
2. Below the plan list, click the **Refresh** button.

Exporting a Deployment Plan

You can export a deployment plan in order to execute a deployment plan for a target server at some later time. All plan resources (bundles, site info, mappings, operations, and so on) are exported to a package file. This can be useful if:

- You develop a deployment plan on a different server than the one where you want to execute the plan.
- The target server is unavailable or offline. For example, you have a remote production server that is inaccessible to the development or test network.
- You want to import and execute the plan on the target server at a later time (see [Importing and Executing a Deployment Plan](#)).

You can export a deployment plan using the Deployment Manager user interface or using the command line as described in the two sections below.

To export a deployment plan from the user interface

1. Click **PLANS**.
2. Select a plan.
3. Above the plan definition, click **Export Plan**.
4. In the Export Plan dialog, click **Download**.

At the bottom of the browser window, the prompt to open or save the plan package is displayed.

5. Choose an option:

Option	Do this
Open	Opens the plan in the appropriate application, or prompts you to select an application.
Save	Saves the file in your download directory.
Save as	Choose a location in the file system.
Save and open	Saves the file in your download directory and opens the plan in the appropriate application, or prompts you to select an application.

To export a deployment plan from the command line

6. Using a command line tool, export the plan using this command:

```
curl -X GET -u <username:password> "http://<hostname>:<port>/rest/deploy/export_plan_package?plan=/<target_site>/<source_site>/<plan_name>&encryptionPassword=<password>" -o plan.pkg -H "Content-Type:application/binary"
```

For example:

```
curl -X GET -u admin:admin "http://localhost:9400/rest/deploy/export_plan_package?plan=/targetSite/sourceSite/planName&encryptionPassword=123456" -o plan.pkg -H "Content-Type:application/binary"
```

See [Importing and Executing a Deployment Plan](#) when you're ready to import a plan that you have exported.

Deleting a Deployment Plan

To delete a plan

1. Click **PLANS**.
2. Select a plan.
3. Below the plan list panel, click **Delete**.
4. Confirm that you want to delete this plan.

Executing Deployment Plans

After you have created a deployment plan and previewed it, you can execute the deployment plan. Deployment Manager saves a log of the execution, as well as keeps a history of all plan executions. See these sections for more information:

- [Executing a Deployment Plan](#)
- [Executing a Deployment Plan Remotely](#)
- [Viewing Source Site Updates Since the Last Plan Execution](#)

- [Viewing the Execution Log Results](#)
- [Purging the Execution Logs](#)
- [Importing and Executing a Deployment Plan](#)

Executing a Deployment Plan

When you execute a plan, Deployment Manager executes all of the operations defined on the plan Operations tab in sequence. Prior to executing a plan, you might want to preview the changes that have occurred on the source site (see [Viewing Source Site Updates Since the Last Plan Execution](#)). After executing the plan, you can check the Execution Log tab for details about the execution.

Requirements

- Executing a deployment plan requires that both the source and target TDV instances are online. See [Editing Site Properties](#) to change a site to online mode and test the connection.
- If you are using file-based caches, you must create the folder and cache files on the target source prior to executing the deployment plan. See [Caching and Deployment Manager](#) for more information about caches and Deployment Manager.


To execute a deployment plan from Deployment Manager

1. Click **SITES** and make sure that the indicator is green for both the source and target sites in the plan.
2. Click **PLANS**.
3. Select a plan.
4. Above the plan, click **Execute Plan**.

Deployment Manager shows a progress bar and then displays the results on the Execution Log tab.

5. Click the **Execution Log** tab to see the results.
6. Click the latest execution log row.

Deployment Manager displays the plan execution result as shown in this example:



Executing a Deployment Plan Remotely

You can execute a deployment plan created using Deployment Manager using a command line tool. Both the source and target sites must be running and online in Deployment Manager to follow the instructions in this section.

Note: If you want to execute a deployment plan that you have previously exported as described in [Exporting a Deployment Plan](#), follow the instructions in [Importing and Executing a Deployment Plan](#).

To run a deployment plan from the command line

1. Click **PLANS**.
2. Select a plan.
3. Click the **General** tab.
4. Click **Copy To Clipboard**.

Deployment Manager copies a plan execution command like the following to the clipboard:

```
curl -d "/targetSiteName/sourceSiteName/planName" -u "username:password" -X POST
"http://localhost:9400/rest/deploy/executePlan" -H "Content-Type:text/plain"
```

5. Open a command line tool.
6. Paste the text that you copied from Deployment Manager into your command line tool.
7. In the pasted text, replace “username” and “password” with your user name and password.
8. Run the command.
9. In Deployment Manager, click the Execution Log for the plan to verify that instructions from the deployment plan were implemented as expected.

Viewing Source Site Updates Since the Last Plan Execution

As part of previewing a plan, you can review changes that have occurred on the source site since the last execution of the plan. This feature is only available when a) there is a migrate resources operation and b) the Retention Policy is set to Keep Non-Colliding.

To view source site updates since this plan was last executed

1. Click **PLANS**.
2. Select a plan.
3. Make sure that there is a Migrate Resources operation and that its Retention Policy is set to **Keep Non-Colliding**.
4. Select the **Preview** tab.
5. Click **SHOW DELTA PREVIEW**.
6. In the new dialog, expand and collapse the following containers to review the changes since the last plan execution. The following containers are shown:
 - New Resources
 - Deleted Resources
 - Updated Resources
 - Renamed/Moved Resources
7. Click **OK**.

Viewing the Execution Log Results

You can view a list of the executions of a plan, the dates and times it was executed, and display a log that describes details about the execution including source and target sites, and number of operations.

To view the execution log files

1. Click **PLANS**.
2. Select a plan.
3. Select the **Execution Log** tab to see the list of execution logs.
4. Select the log for the execution you want to review.

2015-07-10 14:58:47 SUCCEEDED	
SOURCE	localhost_9410
TARGET	localhost_9420
EXECUTION START TIME	2015-07-10 14:58:46
EXECUTION END TIME	2015-07-10 14:58:47
NUMBER OF OPERATIONS	1
OPERATION	Migrate Principal Operation
EXECUTION START TIME	2015-07-10 14:58:47.373
EXECUTION START TIME	2015-07-10 14:58:47.411
OPERATION EXECUTION STATUS	SUCCEEDED

Purging the Execution Logs

You can purge all execution logs of this plan. For example, if you have changed the plan, you might want to start a fresh history.

To purge the execution logs

1. Click **PLANS**.
2. Select a plan.
3. Select the **Execution Log** tab to see the list of execution logs.
4. Click **CLEAR EXECUTION HISTORY**.
5. Confirm that you want to clear the history.

Deployment Manager clears the Execution Log tab.

Importing and Executing a Deployment Plan

You can import and execute a deployment plan that you have previously exported as described in [Exporting a Deployment Plan](#) . When you want to execute the deployment plan that you have exported, the target server must be running and set to online in Deployment Manager. The source site where the plan originated does not need to be running.

At this time, importing and executing a deployment plan can only be done using a command line tool; there is no feature for this in the user interface.

To import and execute a deployment plan

1. Locate the exported plan package file and copy/paste it to your local machine.

Note: The plan package name cannot include special characters like (and). Because Deployment Manager adds (1), (2), etc. to the plan package name by default if you export it multiple times to the same name, you might need to rename the plan package before executing it.

2. Start the target server and make sure that it is not set to offline mode:

Run Deployment Manager.

Click **SITES**.

Select the target site.

Click the **General** tab and make sure that **Server Offline** is unchecked.

3. Using a command line tool, execute the plan using this command:

```
dm_apply_plan.[bat|sh] -password [password] -package [path/plan package] -
server <hostname> -port <port_number> -user <user_name> -domain
<domain_name> -encryptionPassword <encryption password> -
ignoreEncryption -verbose
```

For example:

```
./dm_apply_plan.sh -password admin -package
C:\\users\\jsmith\\downloads\\plan_package.pkg -server localhost -port 9410 -
user admin -domain composite -encryptionPassword 123456 -verbose
```

where C:\\users\\jsmith\\downloads\\plan_package.pkg is the path and filename of the execution plan package. Notice that a double-backslash is required on a Windows machine.

Example with ignoreEncryption option:

```
./dm_apply_plan.sh -password admin -package C:\\users\\jsmith\\downloads\\plan_  
package.pkg -server localhost -port 9410 -user admin -domain composite -ignoreEncryption -  
verbose
```

where `C:\\users\\jsmith\\downloads\\plan_package.pkg` is the path and filename of the execution plan package. Notice that a double-backslash is required on a Windows machine.

Note: If `-ignoreEncryption` option is used, then all backup data will be imported regardless of whether a valid encryption key was provided. This means that the import will not fail. This option can be used to allow partially importing any backed up data. However, the import process will only import data that is not encrypted or can be decrypted using the provided encryption key. All encrypted portions of the backup data that cannot be decrypted will be imported as empty values and the import will otherwise succeed.

This affects all encrypted values in the backup data, which includes, but is not limited to data source and LDAP domain connection passwords.

4. Verify that the resources in the plan have been migrated to your target server.

Backing Up and Restoring the Deployment Manager Server

You can back up the entire Deployment Manager server in a CAR file that contains all Deployment Manager-defined resources and metadata. You can then restore the backup file to another TDV Deployment Manager server.

- [Backing Up the Deployment Manager](#)
- [Restoring the Deployment Manager Server](#)

Backing Up the Deployment Manager

You can back up all defined Deployment Manager server metadata and resources in a CAR file. The backup file can be restored as described in [Restoring the Deployment Manager Server](#).

By default, the backup file name is `deployment_metadata.car`. Or, you can specify the name of your choice for the backup file. When a backup file of the same name exists, it

is not overwritten; instead, a number is added to the subsequent backup filenames as in deployment_metadata (1).car, deployment_metadata (2).car, and so on.

To back up the Deployment Manager Server

1. From the Admin menu, select **Backup**.

Deployment Manager asks you to confirm that you want to create the CAR file and download it.

Type and confirm a password to encrypt the backup CAR file. You will need to communicate the password to anyone that needs to import the file.

2. Click **Download** to confirm.
3. If your browser is configured to download to a particular location, Deployment Manager downloads the backup CAR file to that location. Otherwise, select a location for the backup file.

To back up the Deployment Management Server from command line

Using a command line tool, backup the Deployment Management server using this command:

```
curl -X GET -u <username:password> "http://<hostname>:<port>/rest/deploy/export_dm_metadata?encryptionPassword=<encryption password>" -H "Content-Type:application/binary" -o <filename>
```

For example:

```
curl -X GET -u admin:admin "http://localhost:9400/rest/deploy/export_dm_metadata?encryptionPassword=password" -H "Content-Type:application/binary" -o exportedDM.car
```

Restoring the Deployment Manager Server

If you have backed up the Deployment Manager server as described in [Backing Up and Restoring the Deployment Manager Server](#), you can restore the backup file. This restores all Deployment Manager resource metadata from the backup CAR file to a Deployment Manager server, with the option of overwriting any existing resource metadata by the same name.

Note: Sites and plans that exist on the Deployment Server but not in the CAR file are not affected by the restore operation and are retained.

To restore the Deployment Manager Server

1. If necessary, start Deployment Manager.
See [Starting Deployment Manager](#).
2. From the Admin menu, select **Restore**.
3. In the Upload dialog, click **Browse** and select the CAR file from the file system.
4. Optionally, check the **Overwrite** check box if you want to overwrite the existing metadata with the contents of the CAR file that is selected in the previous step.
5. Type the password that was used to encrypt the CAR file.
6. Optionally, check the Ignore Encryption Errors check box to ignore any encryption errors.

Note: If this option is used, then all backup data will be imported regardless of whether a valid encryption key was provided. This means that the import will not fail. This option can be used to allow partially importing any backed up data. However, the import process will only import data that is not encrypted or can be decrypted using the provided encryption key. All encrypted portions of the backup data that cannot be decrypted will be imported as empty values and the import will otherwise succeed.

This affects all encrypted values in the backup data, which includes, but is not limited to data source and LDAP domain connection passwords.

7. Click **Upload**.
Deployment Manager displays “Result - Imported Successfully” on a successful import and you are returned to the Deployment Manager home page which shows updated values for sites and plans.
8. Refresh your browser to see the imported resources.

To restore the Deployment Management Server from command line

Using a command line tool, backup the Deployment Management server using this command:

```
curl -X POST -u <username:password> "http://<hostname>:<port>/rest/deploy/import_dm_metadata? -H "Content-Type:multipart/form-data" -v -F 'file=@localfilename' -F 'encryptionPassword=<encryption password>' -F 'overwrite=true/false'
```

For example:

```
curl -X POST -u admin:admin "http://localhost:9400/rest/deploy/import_dm_metadata" -  
H "Content-Type:multipart/form-data" -v -F 'file=@localfilename' -F  
'encryptionPassword=password' -F 'overwrite=true/false'
```

with ignoreEncryption option:

```
$ curl -X POST -u admin:admin "http://localhost:9400/rest/deploy/import_dm_metadata"  
-H "Content-Type:multipart/form-data" -v -F 'file=@exportedDM.car' -F  
'ignoreEncryption=true' -F 'overwrite=true'
```

Note: If -ignoreEncryption option is used, then all backup data will be imported regardless of whether a valid encryption key was provided. This means that the import will not fail. This option can be used to allow partially importing any backed up data. However, the import process will only import data that is not encrypted or can be decrypted using the provided encryption key. All encrypted portions of the backup data that cannot be decrypted will be imported as empty values and the import will otherwise succeed.

This affects all encrypted values in the backup data, which includes, but is not limited to data source and LDAP domain connection passwords.

Using TDV Workload Management

Workload management is a process for determining the proper workload distributions to provide optimal performance. It provides control over each work request. It is helpful when you need to manage high-resource usage requests and to control the amount of resources or returns through you TDV systems. Slow downs in processing of requests can be caused by a few high-resource requests. The ability to control request processing by cluster group, user groups, date, and time can help increase customer satisfaction.

For example, using TDV workload management, you can define a rule that limits requests from Marketing to maximum 30% of memory use. Furthermore, a trigger can be used to enforce the workload rule only for certain times. For more information, refer [Example - Creating a trigger to limit memory use using Workload Management](#)

Resource throttling refers to the cutting down or lowering of the amount of resources or returns in a system.

Key benefits include:

- Important workloads get prioritized.
- Intelligent allocation of resources.
- Maximizes performance by preventing bottlenecks.
- Maximizes uptime by eliminating problematic requests.

The following topics are covered:

- [Limitations](#)
- [About Rule Precedence](#)
- [Setting Up and Configuring Workload Management](#)
- [Viewing Your Workload Management Rules in Studio](#)
- [Viewing Rule Precedence](#)
- [Testing Workload Rules](#)
- [Configuring Email Alerts for Workload Management](#)

Limitations

- System resources are not allowed to be added to resource assignments in rules.
- You cannot add a dre user, dre_manager, an admin user in the admin group or with admin privileges to member assignments in rules.
- Only one exception rule can be executed per request.
- Each request is given an initial 2 MB of memory. After the 2 MB is exhausted, then any workload rules with memory limits are applied. This allows normal requests, such as log-in and system requests, to proceed and Studio to continue functioning.
- In workload management, the memory limit rule does not restrict a request from consuming ANY memory. It only restricts memory consumption for query engine operators such as DISTINCT, JOIN, GROUP BY, etc.
- A user without WLM permission will not be able to view
 - the WLM icon on resources.
 - rule in lineage panel of a resource.
- When a user without WLM permission exports a car file with dependencies, assuming WLM rule exists on a resource, the exported car file will not contain the WLM rule in dependencies.

About Rule Precedence

Rules can conflict with one another and rules of the same type can co-exist. For example:

- The TDV Server will choose to implement the most specific rule over a more restrictive rule.
- In the absence of a specific rule on a resource, more restrictive rule is picked among rules of same rule type but different filter limits.

When certain 3rd party clients execute WLM, messages and sometimes actions are captured in the data result set. (e.g.: JDBC, ODBC, ADO.NET, SOAP and REST.)

For JDBC clients, we use SQLWarning.

For SOAP, warnings are returned via a SOAP fault.

For REST, can be returned via a "warnings" json element.

When queries are run against multiple published resources that have WLM rules defined for them, conflicts between the rules defined for each resource can occur. For example, `emp_view` has a full table scan rule defined and `salary_view` has a row limit rule defined, when a query is run that selects data from both `emp_view` and `salary_view`, the different rules cause conflicts in how the query would normally run. Some of these conflicts can be avoided by understanding the order in which rules are given precedence during conflicts. Conflicts arise when two rules apply to the same resource for the same user, and have different filter definitions, or different action types.

- Full table scans and Cross joins are the first rules to be checked.
- Rules with exception actions are executed last. When an Exception is encountered, any other rule that is defined will not be processed.
- When `maxRowLimit` is conflicting the more restrictive rule is used.
- When full table scans are set to true in only one of the rules, the rule in which the full table scan is true is used.
- When cross joins are set to true in only one of the rules, the rule in which the cross join is true is used.
- When `maxRequestTime` and `maxRequestTimeUnit` conflict in two rules, the more restrictive rule is used.
- When different resources have different `memoryLimitPercentage` rules for the same user, the most restrictive rule is used.
- Multiple rules can co-exist on resources and its descendants. In that case, the rule on the descendants is considered a more specific rule compared to a rule on the parent resource.
- When global rules (with null member assignments) co-exist with specific rules on users `u1`, `u2`: For `u1`, `u2`, the specific rule is picked and for all other users, the global rule is used.
- A user can have multiple group memberships. If two rules exist on different groups, but applicable to the same user `u1`, the most restrictive rule is picked for the user on the specific resource.
- When action types, filters, resources or member assignments are different, then the rules are allowed to be created and the most effective, restrictive, or specific rule is chosen at runtime during rule violation checks.
- When action type differs for two rules, with the same filters and same resource/member assignments, and one of those rules already exists, the second

rule will not be allowed to be created. An error will be encountered, indicating the conflict.

- Duplicate rules cannot be created. Similarly updating an existing rule to look identical to another existing rule is not allowed.
- TDV Manager Workload Management page has an option to view effective rules for any user. Users with WLM permission can perform this operation.
- The WLM rule types given below are applicable when the following resources are published as webservices
 - /services/webservices/publishTable - All rules are applicable.
 - /services/webservices/publishView - All rules are applicable.
 - /services/webservices/publishTransformation - Only “Row Limit, Request life time and Memory Limit” are applicable.
 - /services/webservices/published_script - Only “Row Limit, Request life time and Memory Limit” are applicable.
 - /services/webservices/Published_package query - Only “Row Limit, Request life time and Memory Limit” are applicable.

Setting Up and Configuring Workload Management

Workload management rules that you define are applied to Published resources only.

Users with WLM permissions are allowed to:

- create, update, delete and view workload management rules.
- get effective rules of a user
- get effective rules of a user for a resource
- enable or disable workload management feature
- enable or disable workload management rules

About Global Rules

If a parent container, like SCHEMA, CATALOG or DATA_SOURCE, is chosen, all the children are included in rule.

For member assignments, if a group is chosen, all users in the group are included.

Global member rules are rules without any member assignments and are applied to all users in the system, except the default admin user, users in admin group, dre user and dre_manager. Global resource rules are rules without any resource assignments and are applied to all resources except system resources.

To avoid accidental creation of a global rule, some rules might be disabled when resource assignments or user/group assignments get deleted. For example, assuming a rule has only one user in its user assignments, when the user is deleted, TDV automatic impact analysis disables the rule. If you define a rule that depends on a specific resource and if that resource gets deleted, the rule is automatically disabled. This is to ensure that a rule does not become a global rule accidentally.

To set up and configure workload management

1. Open the TDV Web Manager.
2. From the CONFIGURATION menu, choose Workload Management.
3. Click Add Rule.
4. Type a name and any annotation text that you want to have for this rule.
5. Make sure Enable is selected.
6. Select a rule type:

Rule Type	Description
Full Table Scan	Allows you to define a rule when a full table scan occurs.
Memory Limit	Allows you to define a rule based on percentage of memory limit that has been reached.
Product Join	Allows you to define a rule for a SQL cross join.
Request Lifetime Limit	Allows you to define a rule based on how long a request has been running.
Row Limit	Allows you to define a rule based on maximum number of rows being reached.

Rule Type	Description
Request Limit	Allows you to restrict the parent requests created by a user. group or claim in the TDV server.

7. Select one or actions:

Action	Description
Client Warning	Allows you to create a custom warning message that is sent to the client application.
Email	Allows you to create an email notification as part of your workload management rule. Some of the rows are returned in the result set and email is sent to email addresses mentioned in the to definition.
Exception	Allows you to create a custom exception message that will be sent to the client application. An exception message is sent to clients and no result sets are returned.
Log Server Event	Allows you to create a custom message that will be sent to the TDV Server log files. Some of the rows are returned and messages are written to the cs_server_events.log.

8. Optionally under Users/Groups, select Add.

Global member rules are rules without any member assignments and are applied to all users in the system, except the default admin user, dre user, dre_manager and users in admin group.

- Type or select the TDV Domain for which you are defining the rule.
- Select User or Group.
- Select the Name from the list of values.

9. Optionally under Resource, select Add.

- Use up arrow or double click the published resource to select one or more of the TDV published resources.

Global resource rules are rules without any resource assignments and are applied to all resources except system resources.

If a parent container, like SCHEMA, CATALOG or DATA_SOURCE, is chosen, all the children are included in rule.

10. Click OK.

11. Click Save.

12. For email rules, see the WLM instructions under [Configuring Email Alerts for Workload Management](#).

When a rule mentions a specific resource that is present in the TDV Server, a WLM icon is shown against the published resources in Studio. Also, a tooltip is available for that resource indicating which rule is applied to the resource. For global rules, an icon is not shown.

For adding a New User and managing User Rights see the see the Group and User Rights Template under [Understanding TDV User Templates and Rights](#)

Viewing Your Workload Management Rules in Studio

You can use Studio to view the rules defined for workload management.

To view rules

1. Open Studio.

The rules are automatically created in the Studio resource tree under:

```
.. > Policy > Workload
```

2. You can open and review the rule code.
3. Creating, editing or deleting rules must be done from within Web Manager.

Viewing Rule Precedence

In TDV manager, effective rules for a user can be viewed by selecting a specific user on the Workload Management page. A green star is displayed against the rule that is effective for the user selected. This is used to view effective rules for a user when lots of rules exist with same rule types and different filter limits.

To view rule precedence

1. Open Web Manager and navigate to the Workload Management page.
2. Add rules if none exist.
3. Above the list of rules to the right side of the screen, select the domain and user for which you want to view rule precedence.

Manager displays a green star next to the rule that will take precedence for that user.

Testing Workload Rules

This section includes several recommendations for how to test the WLM rules that you create. It is not an exhaustive set of tests. TIBCO recommends that you design your own tests in addition to the ideas presented here.

From Studio SQL Scratchpad, WLM rule violations can be observed.

Occasionally, you can execute a view and use Show Contents to view rule violation. However, using Show Contents to test is not as reliable as using SQL Scratchpad to test rules.

To test your WLM rules

1. From Studio SQL Scratchpad, run the queries that you expect will touch the resources for which you have WLM rules defined.
2. If rule violations occur, review the errors.
3. Determine which rules are causing the violations.
4. Determine if the best course of action is to modify the:
 - WLM rules
 - query that uses those resources

5. Re-test with SQL Scratchpad.
6. Repeat tests with SQL Scratchpad until your queries run without unwanted rule violations.
7. Design and run your own set of tests that use your client applications to access data through TDV.
8. Review TDV log files and the data that is returned to your client application.

For example, your client application queries data through TDV where several resources have WLM rules defined and returns the result set as a REST packet to your client application. If there are any violations to the WLM rules, they will be included in the REST packet. Depending on your client application, you might want those messages displayed or you might want them hidden.

Configuring Email Alerts for Workload Management

You can trigger email notifications for TDV actions or events.

To enable email alerts

1. Open and log in to Studio.
2. From the Administration menu, choose Configuration.
3. Navigate to Server > Configuration > E-Mail.
4. Set values for the following:

Configuration Parameter	Description of Value	Example
From Address	Email address that you want to appear in the From line for alerts.	meg@queenbeesknees.net
SMTP Authentication required	A boolean field and indicates whether an authentication is required or not.	False

Configuration Parameter	Description of Value	Example
SMTP Host Name	Name of the email server host.	javamail.queenbeesknees.com
SMTP Port	Port number of the SMTP server that the server uses when sending out e-mail.	25
SMTP Authentication User Name	User name for connecting to the SMTP server for sending e-mails.	
SMTP Authentication User Password	Password for connecting to the SMTP server for sending e-mails.	

5. Save and exit the Configuration window.
6. In the TDV Manager, when a workload management rule is created or updated, an Email action type can be chosen to receive email notifications about workload rule violations. The following details should be provided. The fields with an “*” next to it are Required:
 - From *
 - To *
 - Cc
 - Bcc
 - Subject *
 - Reply-to
 - Email body

Rule violations will be sent to the email addresses mentioned in the To, Cc and Bcc fields. If notification to multiple addresses is required, a mailing list should be created and used.

Request Limit Rule using Workload Management

Administrators can restrict the parent requests created by a user/group/claim in the TDV server by creating Workload management Concurrent Request limit rules.

The rule type for Concurrent Request Limit WLM rule is defined as "CONCURRENT_REQUEST_LIMIT" and it accepts a number > 0 as part of the parameter "concurrentRequestLimit" in the REST API /rest/workload/v1/rules.

Considerations for Setting up Concurrent Request Limit Rules

Following are the considerations when setting up the rules to restrict the number of requests for a user/group/claim:

- When applied to a user, the Concurrent Request Limit rule will limit the number of parent requests created by the user.
- When applied to a group, the rule will limit the number of parent requests created by users in the group.
- When there are multiple Concurrent Request Limit rules for same user or group, then the most restrictive limit is chosen. This applies to rules applicable to a user coming from a group.
- When there are Concurrent Request Rules on existing requests, then the current requests will remain unaffected by concurrent request limit rules.
- Error will be thrown when concurrent request limit is 0 or less or when resource assignments are specified. Only when workload management and the rule is enabled, the concurrent request limit will apply.
- A SQL warning, an Exception, a Server Event or an Email trigger are the various actions that are defined as rule implications.
- When no member information is specified for a Concurrent Request Limit Rule, it will apply to all users except the admin user, users in admin group, system and inbuilt users. This is called a global rule.
- As a general rule, the "Most Restrictive" rule will be picked up in case of multiple rules existing for the same member (user/group/claim).

To set the Request Limit for a user in Workload Management:

1. Login to Manager as an Admin user.
2. Open Workload Management.
3. Click on the Add Rule button.
4. Provide a Rule name.
5. Select "RequestLimit" from the Rule Type dropdown.
6. Enter the Maximum Request Count.
7. Select the Action you want to trigger when the limit is reached.
8. In the User/Groups/Claims pane, select the domain and user you want to apply this rule to.
9. Save the WLM Rule. The saved rule will be active from then on and will not affect any previous requests created.

To set the Request Limit for a Group/Claim in Workload Management:

1. Login to Manager as an Admin user.
2. Open Workload Management.
3. Click on the Add Rule button.
4. Provide a Rule name.
5. Select "RequestLimit" from the Rule Type dropdown.
6. Enter the Maximum Request Count.
7. Select the Action you want to trigger when the limit is reached.
8. In the User/Groups/Claims pane, select the Domain and Group/Claim you want to apply this rule to.
9. Save the WLM Rule. The saved rule will be active from then on and will not affect any previous requests created.

Note: The rule is not equally shared between the users. It will be allocated on a first-come basis.

Example 1 - Request Limit for Users not part of a Group

1. Login to TDV Manager as an admin user.
2. Go to Configuration->Workload Management.

3. Create a rule with the maximum number of requests (for example, 1) you want to allow and choose an action in case of violation of the rule.
4. Choose the users from the User/Group/Claim panel. Every user chosen will be assigned the maximum request limit of 1 as entered in the previous step.

If the action chosen is an Exception, then when the maximum number of requests for the user exceeds the number defined, then an exception will be displayed in the TDV application that the user logs in.

If the action chosen is a Server Event, then the violation is logged in the Server event log file.

Note: In case of a long running request, you can release it by canceling it (either using the API *cancelInvoke* or via TDV Studio)

Example 2 - Request Limit for a Group/Claim

1. Login to TDV Manager as an admin user.
2. Go to Configuration->Workload Management.
3. Create a rule with the maximum number of requests (for example, 2) you want to allow and choose an action in case of violation of the rule.
4. Choose the group from the User/Group/Claim panel. The number of allocated requests will be shared between all users in the group/claim.

For example, Group/Claim A has a Workload management rule of 2 Concurrent Requests and Server Event is defined as the action type. Group/Claim A has the following users and requests:

User 1 - Executes a query (uses 1 request limit)

User 2 - Executes a query (uses 1 request limit)

User 1 - Attempts to execute another query.

When User 1 tries to execute a second query when the first request is still active, a Server Event will be logged in the server_event log file for exceeding the request limit.

If the action chosen is an Exception, then when the maximum number of active requests has exceeded the number defined, then an exception will be displayed in the TDV application that the group user 3 logs in.

Example 3 - Request Limit for 2 Groups (with a common user)

1. Login to TDV Manager as an admin user.
2. Go to Configuration->Workload Management.
3. Create a rule with the number of requests you want to allow for the two groups and choose an action in case of violation of the rule. The number of allocated requests will be shared between all users in the group/claim.

For example, Group A and Group B have a Workload management rule of 2 Concurrent requests and an Exception is defined as the action type.

Group A has the following users and requests:

User 1 - Execute a query (this uses 1 request limit).

Group B has the following users and requests:

User 1 - Execute a query (this uses 1 request limit).

User 2 - Execute a query (this uses 1 request limit).

User 3 - Execute a query

When User 3 tries to execute a query, an Exception will be displayed.

If the action chosen is a Server Event, then when the maximum number of active request has exceeded the number defined, then a Sever Event will be displayed in the event log.

Example 5 - Global Request Limit rules

To create a global request limit rule,

1. Login to TDV Manager as an admin user.
2. Go to Configuration->Workload Management.
3. Create a rule with the maximum number of requests (for example, 1) that you want to allow and choose an action in case of violation of the rule.
4. Do not assign a user, group or claim.

All users/groups/claims will each have a request limit of 1.

Example - Creating a trigger to limit memory use using Workload Management

Following are the steps to create a trigger and enable/disable it using Workload Management:

1. Create a REST data source.
2. Enter login and password credentials.
Base URL = `http://localhost:9400/rest/workload/v1`
3. Check "JSON Format".
4. Click on the green plus sign to add an operation. Set the following values:
Name = "enable"
HTTP Verb = "PUT".
Operation URL = enable
5. Add a Header/Body Parameter. Enter the following values:
Name = "[rawdata]"
Data type = CHAR
In/Out = "IN"
6. Create a Trigger with the condition as 9am and Action = "Execute Procedure". Point it to the "enable" procedure you just created. Set the "rawdata" parameter to 1.
7. Create a 2nd Trigger with the condition as 1pm and Action = "Execute Procedure". Point it to the "enable" procedure you just created. Set the "rawdata" parameter to 0.

Configuring NTLM Authentication

TDV supports NTLM (NT LAN Manager) authentication, a Microsoft authentication protocol.

NTLM authentication for TDV can be implemented on TDV on Windows or UNIX platforms and can be configured for Studio.

This section includes the following:

- [NTLM Authentication and TDV](#)
- [Implementing NTLM Authentication for Windows](#)
- [Implementing NTLM Authentication for UNIX](#)

NTLM Authentication and TDV

NTLM authentication uses a challenge-response sequence which allows clients to prove their identities without sending a password to the server. It consists of three messages, commonly referred to as Type 1 (request), Type 2 (challenge) and Type 3 (authentication). It works like this:

1. The client sends a Type 1 message to the server. This contains a list of features supported by the client and requested of the server.
2. The server responds with a Type 2 message. This contains a list of features supported and agreed on by the server. Most importantly, it contains a challenge generated by the server.
3. The client replies to the challenge with a Type 3 message. This contains several pieces of information about the client, including the domain and user name of the client user. It also contains one or more responses to the Type 2 challenge.

The responses in the Type 3 message are the most critical piece, because they prove to the server that the client user has knowledge of the account password.

Limitations When Using NTLM with TDV

There are a few limitations when using NTLM with TDV:

Category	Description
Domain Support	NTLM cannot authenticate users from the composite domain. For clients to authenticate successfully when accessing TDV through NTLM, an LDAP domain must be configured for the Windows domain being used, and the TDV must be pointed at this LDAP domain.
Proxy Support	Because NTLM is connection-oriented, it cannot support proxies. TDV is unable to support proxies when using NTLM authentication.
Pass-Through Authentication	Pass-through authentication (delegation) is not possible, because the user does not provide a password, which can be used to construct an NTLM Type 3 message.

Note: According to the Microsoft website 'All versions of NTLM, including LANMAN, NTLMv1, and NTLMv2, are no longer under active feature development and are deprecated. Use of NTLM will continue to work in the next release of Windows Server and the next annual release of Windows. Calls to NTLM should be replaced by calls to Negotiate, which tries to authenticate with Kerberos and only falls back to NTLM when necessary.'

Implementing NTLM Authentication for Windows

There are two scenarios in which NTLM authentication is used with TDV:

Scenario	Description
TDV is the server	Using NTLM to authenticate clients requesting a published web service through Studio. After receiving a client request, TDV replies with the NTLM challenge, to which the client must respond with an authentication message. After TDV is satisfied with the response, it returns the results to the client.
TDV is the client	Introspecting or consuming a WSDL or XML/HTTP data source through a Web service such as Microsoft SharePoint. To retrieve the data, TDV

Scenario	Description
	makes a request to which the Web service replies with the NTLM challenge. TDV must then respond with the correct authentication message. After the challenge is resolved, TDV gets the results from the Web services.

The configuration process depends on which scenario you are configuring NTLM authentication for:

- [Configuring TDV as the Server](#)
- [Configuring SQL Server in the TDV Server to use Windows Authentication](#)
- [Configuring TDV as the Client](#)

Configuring TDV as the Server

If you want TDV to use NTLM to authenticate requests for Web services, you must configure TDV and the TDV resources.

NTLM authentication is generally stronger than Basic authentication. However, because NTLM authentication provides no password information to the server, pass-through authentication might fail for sessions authenticated through NTLM.

To implement NTLM authentication where TDV is the server

1. Install the latest version and patches for TDV.
2. To verify that the necessary libraries and files have been installed, make sure that Common_WindowsSSPI_JNI.dll exists in the following directory:

```
<TDV_install_dir>\apps\common\lib\win64
```

3. Configure an LDAP domain.

Open Manager in your Web browser.

Choose SECURITY > Domain Management to open the DOMAIN MANAGEMENT page.

Add a new LDAP domain that specifies an LDAP domain and password.

Add the groups and users to the new LDAP domain who need to consume resources using NTLM authentication.

For more information about configuring an LDAP domain, see [LDAP Domain Administration](#).

4. Using Studio, set the NTLM authentication configuration parameters:

Choose Administration > Configuration to access the TDV Configuration window.

Expand the TDV Server > Configuration > Security > Authentication configuration parameters:.

Change parameters as shown in the table.

Parameter	Description of Change to Make
Allow NTLM Authentication	Change this value to True.
Tolerate Unused HTTP Authentication Schemes	Keep the default: WARN (Valid values are WARN, IGNORE, and ERROR.)
Windows Domain Mapping	Enter a key-value pair that maps the Windows domain of an authenticated user to the name of the corresponding external domain as it is defined in the TDV Server (the name of the LDAP domain you created). The values entered are case-sensitive.

5. Continue with [Verifying NTLM for a Web Service](#) .

Verifying NTLM for a Web Service

You can verify that the NTLM authentication worked for a REST Web service using the steps below.

Note: You cannot verify a SOAP Web service using a browser.

To verify a REST Web service

6. Create a new REST or SOAP Web service.
7. Publish resources in the new REST or SOAP Web service.

8. Open a browser (for example, Internet Explorer or Firefox) and enter the URL defined for the data source.
9. Enter the username and password when prompted.
10. The results of the Web Service Operation should be displayed.

Configuring SQL Server in the TDV Server to use Windows Authentication

The TDV installation provides a startup script for Windows that you can customize for your own purposes.

To maintain customizations made to this script across hotfix or patch updates, you must activate an environment variable.

The following solution works for the credentials of the user that runs the TDV Server service being used to access the data source. If you require access for another account, use of the JTDS driver is required.

To configure the TDV Server for Windows Authentication

1. Stop the TDV Server.
2. In Windows Explorer, navigate to <TDV_install_dir>\conf.
3. Configure your Microsoft data source using the JDBC driver sqljdbc4.jar. (not JTDS)
4. Make sure that sqljdbc_auth.dll is in a location on the PATH of the CIS host. For example, in %SYSTEMROOT%\System32 and <CIS>\apps\common\lib\win64
5. On the Advanced tab of the data source in Studio, edit the JDBC connection string information. You may need to exclude the port number.

For example:

```
jdbc:sqlserver://<HOST>;IntegratedSecurity=true;DatabaseName=<DATABASE NAME>
```

<HOST> should match the value you used for the hostname\instance name on the Advanced tab for the data source. For example, the generated JDBC string was:

```
jdbc:sqlserver://myHost\myInstance;IntegratedSecurity=true;DatabaseName=myDatabase
```

If you are connecting to a host running a single SQL Server instance with the default port of 1433, the JDBC string would be:

```
jdbc:sqlserver://<HOST>:<PORT>;IntegratedSecurity=true;DatabaseName=<DATABASE NAME>
```

If connecting to a SQL Server instance which is not running the default port, ensure the correct JDBC string is being generated. To do this, the user name and password fields are left blank.

Change the TDV Server service so that it runs as the logged on user, and not the Local System user. The correct credentials are then picked up to access the SQL Server instance.

To ensure that the user name is being entered correctly when providing the credentials for the TDV Server service, use the domain\username.

Troubleshooting the use of Windows Authentication

- Connection Refused from Studio.

Symptom: PostgreSQL errors in the Monitor log.

Due to startup problems when setting the TDV Service to use the Service Account while the Repository Service used the local account. Edit the service restarts until you get the order right.

- TDV Adds the 1433 Port Number into the Connection String

Hard-code the entire connection string in the SQL Server data source Advanced tab.

Configuring TDV as the Client

If you are creating a new REST, SOAP, WSDL, or XML/HTTP data source that needs to use NTLM authentication, follow these steps.

To implement NTLM authentication where TDV is the client

1. Install the latest version and patches for TDV.

2. To verify that the necessary libraries and files have been installed, make sure that Common_WindowsSSPI_JNI.dll exists in one of the following directories:

```
<TDV_install_dir>\apps\server\lib\win64  
<TDV_install_dir>\apps\common\lib\win64
```

3. Configure an LDAP domain.

Open Manager in your Web browser.

Choose SECURITY > Domain Management to open the DOMAIN MANAGEMENT page.

Add a new LDAP domain that specifies an LDAP domain and password.

Add the groups and users to the new LDAP domain who need to consume resources using NTLM authentication.

For more information about configuring an LDAP domain, see [LDAP Domain Administration](#).

4. Using Studio, set the NTLM authentication configuration parameters:

Choose Administration > Configuration to access the TDV Configuration window.

Expand the TDV Server > Configuration > Security > Authentication configuration parameters:.

Change parameters as shown in the table.

Parameter	Description of Change to Make
Allow NTLM Authentication	Change this value to True.
Tolerate Unused HTTP Authentication Schemes	Keep the default: WARN. (Valid values are WARN, IGNORE, ERROR.)
Windows Domain Mapping	Enter a key-value pair that maps the Windows domain of an authenticated user to the name of the corresponding external domain as it is defined in the TDV Server (the name of the LDAP domain you created). The values entered are case-sensitive.

5. Create a new Web service for the REST, SOAP, WSDL, or XML/HTTP data source and publish a resource to the new Web service.

For information about publishing Web resources, see information on publishing in the *TDV User Guide*.


6. For a REST Web service, follow these steps:

Open the REST TDV Web service that you want to configure for NTLM authentication.

Select the REST tab.

Set these Service properties to configure for NTLM:

- Enabled: true
- Enable HTTP NTLM: true




7. For a SOAP or WSDL TDV Web service, follow these steps:

Open the SOAP or WSDL TDV Web service that you want to configure for NTLM authentication.

Select the SOAP tab.

Set these properties to configure for NTLM:

- Enabled: true
- Security Policy: /policy/security/system/Http-NTLM-Authentication.xml



8. In Studio, create a new REST, SOAP, WSDL, or XML/HTTP data source, specifying the following parameters on the Basic tab:

REST connection parameters:

Connection Type	Parameters to Specify
REST	Base URL: URL to access this REST data source using the syntax: Login: <LDAP login for this domain> Password: <LDAP password for this domain> Pass-through Login: Disabled Authentication: NTLM Domain: <LDAP domain name> Method: For the XML/HTTP protocol, under Operations, the specification for HTTP Verb must be POST or GET.

REST example:

SOAP connection parameters:

Connection Type	Parameters to Specify
SOAP	URL: <URL to access this SOAP data source> Login: <LDAP login for this domain> Password: <LDAP password for this domain> Pass-through Login: Disabled Authentication: NTLM Domain: <LDAP domain name>

SOAP example:

Connection Information

Basic

Advanced

Adapter Name:

SOAP

URL:

http://qa-ntlm:9400/soap11/rest_test?wsdl

(Eg. file:///Z:/test.wsdl)

Login:

qa1

Password:

☒ Save Password

Pass-through Login:

Disabled

Authentication:

NTLM

Domain:

SUPPORT

Service Principal Name:

SAML Handler Class:

☒ Enable Data Source

Add/Remove Resources...

Test Connection

WSDL connection parameters:

Connection Type	Parameters to Specify
WSDL Connection Information	URL: <URL to access this WSDL> Login: <LDAP login for this domain> Password: <LDAP password for this domain> Pass-through Login: Disabled Authentication: NTLM Domain: <LDAP domain name>

WSDL example:

The screenshot shows the 'Connection Information' dialog box with the 'Basic' tab selected. The fields are filled as follows:

- Adapter Name: WSDL
- URL: http://qa-ntlm:9400/soap11/rest_test?wsdl (with a hint: (Eg. file:///Z:/test.wsdl))
- Login: qa1
- Password: masked with asterisks
- Save Password: checked
- Pass-through Login: Disabled
- Authentication: NTLM
- Domain: SUPPORT
- Service Principal Name: (empty)

At the bottom, there is a checkbox for 'Enable Data Source' (checked), and buttons for 'Add/Remove Resources...' and 'Test Connection'.

XML/HTTP connection parameters:

Connection Type	Parameters to Specify
XML/HTTP Connection Information	URL: <URL to access this WSDL> Login: <LDAP login for this domain> Password: <LDAP password for this domain> Pass-through Login: Disabled Authentication: NTLM Domain: <LDAP domain name> Method: For the XML/HTTP protocol, under Operations, the specification for HTTP Verb must be POST or GET.

XML/HTTP example:

The screenshot shows the 'Connection Information' dialog box with the 'Basic' tab active. The configuration is as follows:

- Adapter Name:** XML/HTTP
- URL:** http://qa-ntlm:9400/xml/rest_test/customers (with a hint: (Eg. http://qa-ds-lin-1/cgi-bin/GetVoters.cgi))
- Method:** GET
- Login:** qa1
- Password:** [masked]
- Save Password:** ☒
- Pass-through Login:** Disabled
- Authentication:** NTLM
- Domain:** SUPPORT
- Service Principal Name:** [empty]
- Input Options:**
 - ☒ No Input
 - ☐ Input In URL
 - ☐ Input Document Definition [empty] Browse...
- Output Document Definition:** [empty] Browse...
- Buttons:** ☒ Enable Data Source, Add/Remove Resources..., Test Connection

9. Verify that the connection works:

Introspect the REST, SOAP, or WSDL data source.

Open the Web service operation and run it.

Implementing NTLM Authentication for UNIX

In Studio, you can configure NTLM authentication to control access to a WSDL, REST, SOAP, or OData data service. The process to configure for NTLM authentication requires the steps in this section.

Note: Have your IT group review the settings for your UNIX configuration files.

To implement NTLM authentication for UNIX

1. Make sure Samba is installed.
2. Make sure Winbind is installed.
3. Locate and edit the `../etc/samba/smb.conf` file to include the following:

```
[global]
workgroup = SUPPORT          # Domain or workgroup name
server string = NTLM Test Machine
winbind uid =10000-20000    # Range big enough for all domain users
winbind gid =10000-20000
winbind enum users = yes
winbind enum groups = yes
winbind use default domain = yes
winbind separator = +
netbios name = qa-ntlm     # Machine name to report to windows network
encrypt passwords = yes
socket options = TCP_NODELAY SO_RCVBUF=8192 SO_SNDBUF=8192
local master = no
domain master = no
preferred master = no
wins server = 10.1.1.3     # Address of the WINS server
dns proxy = no
```

```
security = domain    # Make Samba machine a member of windows domain
password server = qaad.support.net # Name of domain controller
```

4. Locate `../etc/nsswitch.conf` and edit it as follows:

```
passwd:  files winbind
```

```
shadow:  files
```

```
group:   files winbind
```

5. Test the configuration using the following command:

```
$ testparm
```

6. Start `nmbd`, `smbd` and `winbindd` services.

7. Join the machine to the domain:

```
$ net rpc join -Uroot%<password>
```

8. Test the configuration using a command like the following, replacing the

`authenticate user` value with your user name and password:

```
wbinfo --authenticate=<your user>%<your password>
```

9. Configure an LDAP domain.

Open Manager in your Web browser.

Choose SECURITY > Domain Management to open the DOMAIN MANAGEMENT page.

Add a new LDAP domain that specifies an LDAP domain and password.

Add the groups and users to the new LDAP domain who need to consume resources using NTLM authentication.

For more information about configuring an LDAP domain, see [LDAP Domain Administration](#).

10. Using Studio, set the NTLM authentication configuration parameters:

Choose Administration > Configuration to access the Configuration window.

Expand the TDV Server > Configuration > Security > Authentication configuration parameters:

Change parameters as shown in the table.

Parameter	Description of Change to Make
Allow NTLM Authentication	Change this value to True.
NTLM External Domain	Enter the name of the LDAP domain you configured. This name is only required for UNIX hosts.

11. Verify the Web service by following the steps for the type of Web or Data service:

- [Verifying NTLM for a Web Service](#)
- [Verifying NTLM for an OData Data Service](#).

12. Verify the NTLM configuration with these steps:


Introspect the REST, SOAP, or WSDL data source.

Open the Web Service Operation and run it.

Verifying NTLM for an OData Data Service

To verify NTLM for an OData data service

1. Configure NTLM as described in [Implementing NTLM Authentication for UNIX](#).
2. Publish a table with primary key to a TDV Database in Data Services/Databases.
For example, publish the /shared/examples/ds_inventory/products table to a TDV Database such as Data Services/Databases/examples.
3. Open the TDV Database that contains the resource you published.
4. On the OData tab, check the NTLM check box as shown here.



5. Use the `curl` command line tool to verify the user/password as shown in these examples:

```
curl --ntlm --user ntlmuser1:password
http://DBntlm.comp.com:9410/odata/examples/products (TDV on local linux)
```

```
curl --ntlm --user qa:password http://mega-lt.comp.com:9400/odata/examples/products
(TDV on remote win 7)
```

Note: You should get the result back if NTLM authentication passes.

For negative case: (wrong password)

```
curl --ntlm --user ntlmuser1:password1  
http://DBntlm.comp.com:9410/odata/examples/products
```

Result:

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

```
<html>
```

```
<head/>
```

```
<body>
```

```
<div style="font-family: sans-serif; color: #990000; margin-top: 5px; margin-bottom:  
5px; text-align: center">TDVCO INFORMATION SERVER
```

```
<hr style="border-style: groove;"/>
```

```
</div>
```

```
<div style="font-family: sans-serif;">
```

```
<b>Error: </b>401 Unauthorized
```

```
</div>
```

```
<div style="font-family: sans-serif;">
```

```
<b>Description: </b>Authentication failed.
```

```
</div>
```

```
</body>
```

Pluggable Authentication Modules

A Pluggable Authentication Module (PAM) is a Java-based security mechanism. PAM provides an optional mechanism for positively identifying valid users. TDV supports it as a way for custom implementation modules to participate in the TDV logon processing.

Note: PAM implementation and management changed in TDV 7.0.3. This new PAM implementation is described here.

PAMs are tightly integrated with the TDV Server in the TDV extensions framework. Within this framework they can:

- Implement authentication against one or more Kerberos realms
- Store credentials in the user session to be applied concurrently to related data sources
- Implement custom authentication against external security access providers
- Implement ACLs (access control lists) to control access to lists of users based on a schedule or other criteria
- Perform real-time auditing and notification of user logon activity
- Enable TDV logging directly into the cs_server.log
- Generate a detailed dump of PAM configuration state & options
- Generate a dump of internal security objects like subject and principal objects

The following topics are covered:

- [About Pluggable Authentication Modules](#)
- [Working with TDV and PAM](#)
- [What Happens at Deployment and Run Time](#)
- [Undeploying Pluggable Authentication Modules](#)
- [Example](#)

About Pluggable Authentication Modules

Pluggable authentication modules let you enforce a secure authentication regime to regulate access to resources accessed through TDV. Each active authentication module, when consulted, must do one of the following.

Action at sign-on	If credentials are
Abort	
Approve	
Disqualify itself	

The module can also add information to the security context in the session.

After using TDV to define user and group access profiles, you can begin to layer PAM security protocols. You can use one or more overlapping PAM implementations on the same server to achieve the desired level of user identification.

Login modules that implement PAM determine authentication based on the data in headers, properties, certificates, and on the user name and password provided.

Authentication Location	Description
HTTP Headers	Incoming HTML headers are passed to authentication modules.
JMS Properties	Properties associated with an incoming message object are passed to authentication modules.
SOAP Headers	Each distinct element in the SOAP Header element of an incoming SOAP envelope is added to the list of supplied properties, keyed by the QName of the element. When present, the header value is represented by an instance of <code>org.apache.axiom.om.OMElement</code> . This applies to the <code>AuthenticationFilter</code> and the <code>WsapiServlet</code> entry-points.
JDBC/ODBC/ADO.NET	To pass into TDV, values must be encoded into a single,

Authentication Location	Description
Properties	fixed, known property name. ADO.NET and ODBC uses semicolons as property separators; JDBC uses ampersands. For user-legibility and compatibility with RFC-2396, security tokens in URL form are passed through using: "user_tokens=("NAME"="VALUE (","NAME"="VALUE)*")". Nonalphanumeric characters within a NAME or VALUE must be URL-encoded. GUI support deletes the parenthesis characters and removes all whitespace characters prior to processing. If a value for user-tokens is specified through the ODBC or ADO.NET UIs and is overridden by a user-specified value, the entire user-token value is replaced.
Incoming SSL certificates	If the user connects to TDV through mutually-authenticated SSL, the connection's public certificate is added to the security context for use by PAM modules.

Minimum Elements of a PAM

The minimum elements that comprise a TDV PAM are:

- A manifest file with all required, TDV-specific entries.
- A Java class that implements the `javax.security.auth.spi.LoginModule` interface and applies a required TDV-specific `@CisExtensionProvider` class-level annotation.
- A primary deployment JAR file (referenced during PAM deployment) that includes the required components listed above.
 - The top-level folder structure of this file's elements must be: `com > tibco > cis > security > auth > module > utils`.
 - In parallel with the `com` folder are folders named `config` (optional; contains `pam_instance.properties`), `lib` (optional; contains utility classes and custom or third-party JAR files), and `META-INF` (contains the manifest file).

You can add optional elements to this JAR file, such as other custom Java classes directly related to the PAM implementation. You can also add utility classes in separate JAR files (in a lib folder within the primary deployment JAR folder) that are referenced as part of the implementation.

You deploy the JAR file to a TDV Server using the TDV utility server_util script, adjusting parameters in deploy.bat to target the installed TDV instance on your system, JAR file name, and so on. After that, you use the TDV Web Manager to add, configure, and enable the login module as an active PAM.

A generic implementation would include these elements:

- The manifest.mf file required for PAM deployment.
- A primary TDV class that implements the required `javax.security.auth.spi.LoginModule` interface and applies the required `@CisExtensionProvider` annotation. This annotation identifies the PAM implementation class as a TDV extension provider.
- An interface that implements global TDV constants.
- Helper classes that provide debug logging of internal TDV session security content, PAM-specific email notifications, and other helper methods.
- A JRE system library that is compatible with JDK version 1.7.x.
- Required TDV libraries, other referenced libraries; custom self-logging exception classes; Java mail helper classes.

Working with TDV and PAM

Pluggable authentication modules work with TDV to direct or augment TDV user security policies.

- [The Manifest File](#)
- [Creating a Principal Authentication Module](#)
- [Enabling PAM Features for TDV](#)
- [Adding a Module](#)
- [Ordering Module Execution Sequence](#)
- [Assigning Users to TDV Groups or Identities](#)
- [Creating a Principal Authentication Module](#)

The Manifest File

The manifest.mf file is a simple text file that contains at least the five entries that TDV requires for PAM deployment and undeployment, and ending with a blank line. The manifest file is required for PAM

The possible name-value pairs are listed in the table.

Name	Req.	Description of Value
Manifest-Version	YES	Always set to 1.0 for TDV.
cisext-name	YES	Unique name of the PAM packaged extension applied when deployed to TDV. By convention, the name equals the primary deployment JAR file name, minus the JAR suffix.
cisext-version	YES	Integer version number of the PAM relative to the custom implementation. Defaults to 1 if missing. The version number is important: it is applied to the deployed PAM, and must be supplied as a parameter during PAM undeployment.
cisext-annotation	YES	Description of the TDV extension package. This annotation is visible in Studio for the deployed PAM. It is useful to include the word "PAM" in the object's name.
cisext-introspectAll	YES	TDV extension introspection action. Defaults to TRUE. Must be set to TRUE for the PAM to appear in the TDV Web Manager's Add Module Instance configuration dialog box.
cisext-dependencies	NO	TDV extension dependencies. Multiple dependencies are listed on indented lines, followed by an empty line. Unversioned dependencies are not accepted. An empty line is required between general properties and file-specific properties.
name	NO	Full path and class file name of the JAAS class that is

Name	Req.	Description of Value
		included in the JAR file for the package that describes the PAM. For example: test/PamTest1.class
cisext-introspect	NO	For PAM, set the value to TRUE, which indicates that TDV should read any included class files.

Creating a Principal Authentication Module

This section provides guidelines for creating a JAAS-based PAM.

The TDV uses authentication modules that JAAS (Java Authentication and Authorization Services), which contain zero-argument constructors.

Note: It is recommended that user-supplied LoginModule implementations avoid storing state in class variables.

To create the necessary JAR file

1. Create a manifest.mf file with the following name-value pairs:

```
Manifest-Version: 1.0
```

```
cisext-name: <name>
```

```
cisext-annotation: <description>
```

```
cisext-version: <integer_version_number>
```

```
cisext-introspectAll: true
```

```
cisext-dependencies: <dependencyName>:<version_number>
```

```
name: <fullpath_and_class_file_name>
```

```
cisext-introspect: true
```

2. Make sure the last line of the manifest file is empty.
3. Save the manifest file.
4. Create your JAAS login module class files using the following recommendations.

Value	Recommendation
Initialization	Initialization method. Whenever a user signs in to TDV, each registered LoginModule class is instantiated and its initialization method is called.
subject	A <code>javax.security.auth.Subject</code> instance in which the LoginModule might store principals, certificates or other security-related credentials.
callbackHandler	An instance of <code>javax.security.auth.callback.CallbackHandler</code> that might be used to retrieve the user's name and password, through the <code>NameCallback</code> and <code>PasswordCallback</code> classes. Passing any other <code>Callback</code> instances to the handler results in an <code>UnsupportedCallbackException</code>.
sharedState	All operating LoginModules are passed a copy of a <code>Map<String,?></code> that has been initialized with all of the channel properties for that connection (HTML headers and others). If you want an authentication module to validate a particular channel type, populate the extra Transport-Type channel property with one of these constants: <code>http</code>, <code>jms</code> or <code>db</code> (for JDBC/ODBC/ADO.NET).
options	Each LoginModule instance is passed a <code>Map<String,?></code> containing its configuration parameters.
Login	The method where the module decides login status. The three possible outcomes are: • Pass—The login method returns <code>TRUE</code>.

Value	Recommendation
	• Fail—The method throws a <code>LoginException</code> • Neither—The login module returns <code>FALSE</code> to disqualify itself from the sign-on process. For example, a <code>LoginModule</code> designed to validate JDBC connections where the user is authenticating through a Web service might then return <code>FALSE</code>.
Commit	If no exceptions occur, the sign-on is considered successful, and each module can add whatever security credentials it wants to the Subject instance identified in the initialize method.
Abort	If any <code>LoginModule</code> fails, that module's abort method is called to allow the module to release any resources it might be holding.
Logout	The method to call when the user signs out.

5. Save the class files and the manifest.mf file in a single JAR file.
6. Use `server_util.sh` to deploy the package that contains your PAM authentication. For instructions, see [Deploying Pluggable Authentication Modules](#).

Deploying Pluggable Authentication Modules

Use of PAM is optional. PAM implementations can coexist with standard user name and password authentication on the same server to provide overlapping degrees of access and privilege to use data resources.

To deploy a pluggable authentication module

1. Open a command prompt window.
2. Navigate to `<TDV_install_dir>/bin`.
3. Deploy the PAM security project bundle using the following `server_util` syntax:

```
server_util -server <hostname> [ -port <port> ] [ -encrypt ]
-user <username> -password <password> [ -domain <domain> ]
-deploy -package <package file in file system> [ -verbose ]
```

For details about using `server_util`, see [The TDV Server Utility Program](#).

4. To validate deployment of the module, see [Verifying that the PAM Deployed](#).

Note: You can also implement the module to use the ip address. A sample implementation is provided in the location `<TDV_InstallDir>\docs\examples\SamplePAMAuthenticationModule`

What Happens at Deployment and Run Time

At deployment, the embedded JAR files are automatically extracted to disk, and at run time the PAM's packaged extension class loader will include the JAR files when resolving Java classes.

The optional `lib` folder in the primary deployment JAR file is a TDV-packaged extension framework mechanism that lets you bundle implementation-dependent run-time JAR file libraries—for example, custom or third-party JAR files—with the PAM.

Verifying that the PAM Deployed

The first level of verification is to make sure that TDV has recognized the package that you deployed. Optionally, you can use the TDV-supplied checksum function to verify the deployment of PAM packages supplied by third parties.

To verify that TDV recognized the PAM module

1. Open Studio.
2. Expand `<localhost>/Packages` in the Studio resource tree.
3. Verify that the name of your package is displayed under the `<hostname>/packages/<package_name>` folder.

This is the value you specified in `cisext-name` in the `manifest.mf` file, plus the `.jar` suffix. The `cisext-version` value is used to identify the instance. The `cisext-annotation` is also visible in Studio.

4. If the PAM fails to deploy properly or does not work, try the techniques listed in [Troubleshooting PAM Deployment](#).

To use the checksum function

5. Open a command prompt window.
6. Navigate to <TDV_install_dir>/bin.
7. Run the checksum validation against the PAM security project bundle using the following server_util syntax:

```
server_util -server <hostname> [ -port <port> ] [ -encrypt ]
```

```
-user <username> -password <password> [ -domain <domain> ]
```

```
-checksum <algorithm/hex-bytes> [ -verbose ]
```

For details about using server_util, see [The TDV Server Utility Program](#).

8. If the PAM fails to deploy properly or does not work, try the techniques listed in [Troubleshooting PAM Deployment](#).

Troubleshooting PAM Deployment

Although the deployment script may indicate success, there are several conditions under which the PAM may fail to deploy properly or simply not work:

- The PAM hierarchy of objects is incomplete in Studio (e.g. 4 objects are expected but only the first two objects appear for a PAM). This was observed to occur when the MANIFEST.MF properties were not properly set. Review the minimal set of PAM manifest entries as documented.
- The PAM may not execute as expected when a user logs in. The PAM code may not be correct. It is important to test the PAM for the range of valid use cases. You might also run sanity checks immediately after you add and activate the PAM. You can manage any enabled PAM module from the Security tab Login Modules page of the TDV Web Manager for the target TDV instance.
- You can run a tail command on cs_server.log to observe what happens when you log in to TDV as a non-admin user.

Note: TDV admin user logon bypasses PAM handling.

Enabling PAM Features for TDV

You can use the Manager Web UI to enable your pluggable authentication modules.

To enable pluggable authentication

1. From Studio, select Administration > Launch Manager.
2. Log on as the admin user.
3. Select Security > Login Modules.



4. Use the Change Enabling buttons to enable or disable any of the following options.


Property	Description
Run Login Modules	PAM is disabled by default. This button must be switched to ENABLED for any authentication module to run.
Modules May Deny TDV Users	Enables PAM to prohibit the user from signing on, even if the user passes TDV or LDAP security.
Log Performance Data	Tracks the time consumed from the moment of login submission to return from the authentication module. It is recommended that you disable this after you have determined that authentication works as expected.
Log Authentication Failures	When enabled, all authentication failures are logged.

Adding a Module

Modules must be added to make them available for TDV to execute them at login. List security modules in the order in which you want them executed. This is useful especially if authentication modules should sequentially update and evaluate the security context of a user attempting to log in.

To add a login security module

1. From the Login Module page, select Add Module.



2. Select a module from the drop-down list to make it available for execution and to set its status.
3. Optionally, select Disabled or Enabled.
4. For Group Mapping, start typing a composite domain name, and select it from the list that appears.

Assigns the authentication module to the specified domain and group. This is a generic assignment for every user successfully authenticated using the security module. If the login module is not used to identify the user, the group mapping is not assigned to the user.

5. Optionally, type an annotation that helps describe the login module that you are adding.
6. Optionally, select the Properties tab.

These are the configuration properties defined in the Java code for the module. You can add or remove properties on this tab, but if the Java implementation of the login module has properties that are already set, those properties should be loaded from that module.

7. Optionally, add properties for your module, in the form of name-value pairs.

8. Click OK.

Ordering Module Execution Sequence

From the Login Modules page of Manager, you can order module execution by ordering their names in a list. The table of modules lists the modules present in all project bundles that are currently enabled.

To define the order of module execution

1. Open TDV Manager.
2. Navigate to the Login Modules page.
3. Make sure that all of the modules you need are listed on the page. If any are missing, follow the instructions in [Deploying Pluggable Authentication Modules](#) to add them.
4. Select a module from the list.
5. Use the Move Up and Move Down buttons to position the module in the list.
6. If the list contains a module that you want to keep but not run, use the Change Enabling button to disable the module.
7. If the list contains a module that you no longer want, use Remove Module.

Assigning Users to TDV Groups or Identities

The TDV Manager Principal Mapping page allows you to link PAM authentication with TDV-defined users and groups.

If a user is validated, the security context values present in the Subject instance or any channel properties for that connection can be used to authenticate a specific group or user for that session. Group membership can be assigned to the user session based on any of the following:


- TDV-wide static group list
- Groups associated with approving authentication modules
- Groups associated with principals contained in the security context

If the user is from the dynamic domain, the permission check bases decisions on the session's security context.

It is recommended that dynamically-defined users have a restricted set of privileges and rights to access information in the public domain, unless you require PAM authentication for all users.

To enable use of principal mapping

1. Navigate to the TDV Manager Principal Mapping page, and from there select Security > Principal Mapping.



2. If the text to the left of Principal to Group Mapping says Disabled, click Change Enabling.

Enable mapping to use values from the security context of the authenticated user to assign that user to a TDV group.

3. Select Add Mapping.
4. Select a name for the mapping.
5. Select the Assignment Types to map values from the user's security context.

Assignment Type	Description
X.500 Distinguished Name	This is a hierarchical string expression that allows domain, organization, and group granularity to be set in the Name field. Users defined within the more general container are mapped to the specified group for data and resource privileges.

Assignment Type	Description
Name Match (exact)	The name in the security context must match the specified string exactly for the group assignment to occur.
Name Match (regex)	Wild cards and special characters can be used to enable members of the same domain to have group privileges. Refers to the Java docs describing Java Class Pattern for regular expressions in <code>java.util.regex.Pattern</code> .
Kerberos Realm	If a Kerberos principal has been added to the security context by a login module, a group mapping can be assigned based on the principal realm. This adds to the mapping that was set for the module instance.

- Specify the domain and group rights and privileges to assign to the users who match the given name.

The domain specification can be omitted if you want to specify the TDV domain, which is the default behavior.

Undeploying Pluggable Authentication Modules

Pluggable authentication modules can be decommissioned (undeployed) from TDV.

To undeploy a PAM

- From Studio, select Administration > Launch Manager.
- Log on as the admin user.
- Select Security > Login Modules.
- Select the PAM module.
- Click Remove Module.

Add Module and Remove Module do not affect the deployed or undeployed state of a PAM.

- Open a command prompt window.

7. Navigate to <TDV_install_dir>/bin.
8. Undeploy the PAM security project bundle using the following server_util syntax:

```
server_util -server <hostname> [ -port <port> ] [ -encrypt ]
```

```
-user <username> -password <password> [ -domain <domain> ]
```

```
-undeploy -name <package name> -version <version number>[ -verbose ]
```

For details about using server_util, see [The TDV Server Utility Program](#).

Example

This is an example of a PAM module that performs a callback.

```
Manifest-Version: 1.0
```

```
cisext-name: example
```

```
cisext-annotation: disqualification or callback
```

```
cisext-version: 2
```

```
cisext-introspectAll: true
```

```
package com.tibco.cis.pam.example;
```

```
import java.util.Map;
```

```
import java.io.IOException;
```

```
import com.compositesw.extension.sdk.annotations.CisExtensionProvider;
```

```
import javax.security.auth.Subject;
```

```
import javax.security.auth.callback.Callback;
```

```
import javax.security.auth.callback.CallbackHandler;
```

```
import javax.security.auth.callback.NameCallback;
```

```
import javax.security.auth.callback.PasswordCallback;
```

```
import javax.security.auth.callback.UnsupportedCallbackException;
```

```
import javax.security.auth.login.LoginException;
```

```
import javax.security.auth.spi.LoginModule;
```

```
import com.compositesw.extension.ds.Logger;
```

```
import com.compositesw.extension.ds.impl.LoggerImpl;
```

```
@CisExtensionProvider(  
    name = "TDV7CallbackExamplePAM",  
    annotation = "PAM Module that performs a callback")  
public class TDV7CallbackExamplePAM implements LoginModule {  
    protected static Logger logger = LoggerImpl.getLogger  
        (TDV7CallbackExamplePAM.class);  
}
```

```
private CallbackHandler handler;
```

```
private String user;
```

```
private String pass;
```

```
public void initialize(Subject subject, CallbackHandler callbackHandler,
```

```
Map<String, ?> sharedState, Map<String, ?> options) {
```

```
    logger.info("Method: " + this.getClass() + ".initialize() called...");
```

```
    handler = callbackHandler;
```

```
}
```

```
public boolean abort() throws LoginException {
```

```
    logger.info("Method: " + this.getClass() + ".abort() called...");
```

```
    return true;
```

```
}
```

```
public boolean commit() throws LoginException {
```

```
    logger.info("Method: " + this.getClass() + ".commit() called...");
```

```
    logger.info("User " + user + " signed on using password: *****");
```

```
return true;
```

```
}
```

```
public boolean login() throws LoginException {
```

```
    logger.info("Method: " + this.getClass() + ".login() called...");
```

```
    NameCallback nameCallback = new NameCallback(" ");
```

```
    PasswordCallback passwordCallback = new PasswordCallback(" ", false);
```

```
    Callback[] callbacks = new Callback[] { nameCallback, passwordCallback };
```

```
    try {
```

```
        handler.handle(callbacks);
```

```
        user = nameCallback.getName();
```

```
        pass = String.copyValueOf(passwordCallback.getPassword());
```

```
        pass = "*****"; // mask password
```

```
    } catch (IOException | UnsupportedCallbackException e) {
```

```
        logger.error("Error during PAM login of user: " + user, e);
```

```
        throw new LoginException("Error during PAM login of user: " + user);
```

```
}
```

```
return true;
```

```
}
```

```
public boolean logout() throws LoginException {
```

```
    logger.info("Method: " + this.getClass() + ".logout() called...");
```

```
return true;
```

```
}
```

```
}
```

Collecting TDV and Data Usage Metrics

Key performance indicators (KPI) around data usage are important to making sure that a company is devoting precious time to the most critical business needs.

The following topics are covered:

- [About Data Usage Metrics](#)
- [Setting Up and Configuring Metrics Collection](#)
- [Configuring Email Alerts for Metrics Notification](#)
- [Publishing and Reporting on TDV Metrics Data](#)
- [About Using MAXMEMORY in Your Reports](#)

About Data Usage Metrics

Using TDV, you can collect data on key performance indicators (KPI) such as which TDV objects are accessed the most, the least, how frequently those objects are accessed, and who is accessing them. After the data is collected, you can publish the tables it is collected in and use your favorite reporting tools to provide your team with these important data points.

Metrics Life Cycle Phase	Description
Capturing	Metrics are incrementally captured in a database of your choice.
Retention	The retention of the metrics are configurable by time period.
Consumption of Metrics	Publish the tables it is collected in and use your favorite reporting tools to provide your team with these important data points.

Supported Data Source Types for Metrics Storage

Data collected for the metrics needs to be collected outside of the TDV Repository. The supported databases for metrics table storage are:

- PostgreSQL
- Oracle
- SQL Server

Limitations

- The tables used or created for the metrics database cannot have a dash in their names.

Setting Up and Configuring Metrics Collection

To set up and configure metrics collection

1. Perform the steps in one of the following sections:
 - [Pre-Creating the External Database and Tables for Metrics Data Storage](#)
 - [Using Studio to Create the Database and Tables for Metrics Data Storage](#)
2. Review and manage your Studio configuration parameters as described in:
 - [Configuring TDV Metrics Collection](#)

Pre-Creating the External Database and Tables for Metrics Data Storage

If you need to create the database and table that are used to store your usage metrics without using Studio, you can use the following instructions.

To create the metrics database and required tables

1. Using the administration tools or command line calls, create the following tables on a PostgreSQL, Oracle, or SQL Server database. For Oracle there are a few syntax differences, such as BIGINT needs to be number (10,0).

Table Name	Create Table Syntax	Create Table Syntax for Oracle	Create Table for SQL Server
metrics_sessions	<pre>CREATE OR REPLACE TABLE tutorial.metrics_ sessions ("cluster" VARCHAR(255), nodehost VARCHAR(255) NOT NULL, nodeport INTEGER NOT NULL, sessionid BIGINT NOT NULL, sessiontype VARCHAR(40) NOT NULL, clienthost VARCHAR(255), "type" VARCHAR (20) NOT NULL, logintime TIMESTAMP NOT NULL, logouttime TIMESTAMP, "status" VARCHAR (20), totalduration BIGINT, totalRequests BIGINT, bytestoclient BIGINT, bytesfromclient BIGINT, "user" VARCHAR (255), "domain" VARCHAR(255), "group" VARCHAR (255));</pre>	<pre>CREATE OR REPLACE TABLE tutorial.metrics_ sessions ("cluster" VARCHAR2(255), nodehost VARCHAR2(255) NOT NULL, nodeport INTEGER NOT NULL, sessionid NUMBER(10,0) NOT NULL, sessiontype VARCHAR2(40) NOT NULL, clienthost VARCHAR2(255), "type" VARCHAR2 (20) NOT NULL, logintime TIMESTAMP NOT NULL, logouttime TIMESTAMP, "status" VARCHAR2(20), totalduration number(10, 0), totalRequests number(10, 0), bytestoclient number(10, 0), bytesfromclient number(10, 0), "user" VARCHAR2 (255), "domain" VARCHAR2(255), "group"</pre>	<pre>CREATE TABLE [Northwind].[guest]. [metrics_sessions] (["cluster"] varchar (255), [nodehost] varchar (255) NOT NULL, [nodeport] int NOT NULL, [sessionid] bigint NOT NULL, [sessiontype] varchar(40) NOT NULL, [clienthost] varchar (255), ["type"] varchar(20) NOT NULL, [logintime] datetime NOT NULL, [logouttime] datetime, ["status"] varchar (20), [totalduration] bigint, [totalRequests] bigint, [bytestoclient] bigint, [bytesfromclient] bigint, ["user"] varchar (255), ["domain"] varchar (255), ["group"] varchar (255))</pre>

Table Name	Create Table Syntax	Create Table Syntax for Oracle	Create Table for SQL Server
		(VARCHAR2(255));	
metrics_requests	<pre>CREATE OR REPLACE TABLE tutorial.metrics_ requests ("cluster" VARCHAR(255), nodehost VARCHAR(255) NOT NULL, nodeport INTEGER NOT NULL, requestid BIGINT NOT NULL, parentid BIGINT, sessionid BIGINT NOT NULL, requesttype VARCHAR(255) NOT NULL, description VARCHAR (10000000), starttime TIMESTAMP NOT NULL, endtime TIMESTAMP, totalduration BIGINT, serverduration BIGINT, rowsAffected BIGINT, maxmemory BIGINT,</pre>	<pre>CREATE OR REPLACE TABLE tutorial.metrics_ requests ("cluster" VARCHAR2(255), nodehost VARCHAR2(255) NOT NULL, nodeport INTEGER NOT NULL, requestid NUMBER (10,0)NOT NULL, parentid NUMBER (10,0), sessionid NUMBER(10,0) NOT NULL, requesttype VARCHAR2(255) NOT NULL, description CLOB, starttime timestamp(9) NOT NULL, endtime timestamp (9), totalduration number(10, 0), serverduration number(10, 0), rowsAffected number(10, 0), maxmemory</pre>	<pre>CREATE TABLE [Northwind].[guest]. [metrics_requests] (["cluster"] varchar (255), [nodehost] varchar (255) NOT NULL, [nodeport] int NOT NULL, [requestid] bigint NOT NULL, [parentid] bigint, [sessionid] bigint NOT NULL, [requesttype] varchar(255) NOT NULL, [description] text, [starttime] datetime NOT NULL, [endtime] datetime, [totalduration] bigint, [serverduration] bigint, [rowsAffected] bigint, [maxmemory] bigint, [maxdisk] bigint, [message] text, ["status"] varchar (20), ["user"] VARCHAR (255), ["domain"]</pre>

Table Name	Create Table Syntax	Create Table Syntax for Oracle	Create Table for SQL Server
	<pre> maxdisk BIGINT, message VARCHAR (65535), "status" VARCHAR (20), "user" VARCHAR (255), "domain" VARCHAR(255), "group" VARCHAR (255)) "queryplan" VARCHAR(255); </pre>	<pre> number(10, 0), maxdisk number (10, 0), message CLOB, "status" VARCHAR2(20), "user" VARCHAR2 (255), "domain" VARCHAR2(255), "group" VARCHAR2(255)) "queryplan" VARCHAR2(255); </pre>	<pre> VARCHAR(255), ["group"] VARCHAR(255) ["queryplan"] VARCHAR(255)) </pre>
metrics_ resources_ usage	<pre> CREATE OR REPLACE TABLE tutorial.metrics_ resources_usage ("cluster" VARCHAR(255), nodehost VARCHAR(255) NOT NULL, nodeport INTEGER NOT NULL, sessionid BIGINT NOT NULL, "user" VARCHAR (255), "domain" VARCHAR(255), "group" VARCHAR (255), requestid BIGINT NOT NULL, parentid BIGINT, </pre>	<pre> CREATE OR REPLACE TABLE tutorial.metrics_ resources_usage ("cluster" VARCHAR2(255), nodehost VARCHAR2(255) NOT NULL, nodeport INTEGER NOT NULL, sessionid number (10, 0) NOT NULL, "user" VARCHAR2 (255), "domain" VARCHAR2(255), "group" VARCHAR2(255), requestid number (10, 0) NOT NULL, parentid number </pre>	<pre> CREATE TABLE [Northwind].[guest]. [metrics_ resources_usage] (["cluster"] varchar (255), [nodehost] varchar (255) NOT NULL, [nodeport] int NOT NULL, [sessionid] bigint NOT NULL, ["user"] varchar (255), ["domain"] varchar (255), ["group"] varchar (255), [requestid] bigint NOT NULL, [parentid] bigint, [datasourcepath] varchar(255), </pre>

Table Name	Create Table Syntax	Create Table Syntax for Oracle	Create Table for SQL Server
	<pre>datasourcepath VARCHAR(255), datasourcetype VARCHAR(255), resourcepath VARCHAR(255), resourcetype VARCHAR(40), resourceguid VARCHAR(40), resourcekind VARCHAR(20), starttime TIMESTAMP NOT NULL, endtime TIMESTAMP);</pre>	<pre>(10, 0), datasourcepath VARCHAR2(255), datasourcetype VARCHAR2(255), resourcepath VARCHAR2(255), resourcetype VARCHAR2(40), resourceguid VARCHAR2(40), resourcekind VARCHAR2(40), starttime timestamp(9) NOT NULL, endtime timestamp (9));</pre>	<pre>[datasourcetype] varchar(255), [resourcepath] varchar(255), [resourcetype] varchar(40), [resourceguid] varchar(40), [resourcekind] varchar(20), [starttime] datetime NOT NULL, [endtime] datetime)</pre>

2. Make sure that permissions on the database and tables allow for their modification from TDV Studio.
3. Note the connection information for the database, so that you have the necessary information to add it as a data source in Studio.
4. Open Studio.
5. Create a new data source that connects to the database and tables that you have created to hold metrics storage data.
6. In the Studio resource tree, select localhost > policy > metrics.
7. Right click and select Open.
8. Browse to the data source where you created the tables for the storage of metrics data. See [Supported Data Source Types for Metrics Storage](#).
9. Bind the tables that you created for the storage of metrics data.
10. Browse to the schema location of the tables that you created.
11. Bind each of the following:

Table Name	Description of Steps to Create the Table
metrics_sessions	- Click Browse. - Navigate to the data source or schema where you created the table. - Click OK.
metrics_requests	- Click Browse. - Navigate to the data source or schema where you created the table. - Click OK
metrics_resources_usage	- Click Browse. - Navigate to the data source or schema where you created the table. - Click OK

12.Click OK.

13.Click Save.

14.Optionally, make selections for the following:

Option	Description
Memory Threshold	Number of megabytes to set aside as the buffer space to store records. When the buffer reaches the threshold, the results are posted to the metrics tables.
Request Count Threshold	Number of records or rows to retain in the buffer. When the buffer reaches the threshold, the results are posted to the metrics tables.
Never Delete Metrics Data	If you choose this option, then the following two fields "How long do you want to keep the metrics data?" and "How often do you want to run the truncate process on expired data?" will not be effective. When this field is not selected, then the following 2 fields will be effective.

Option	Description
How long do you want to keep the metrics data?	Use to indicate how long you want your metrics data retained. The default is 30 days and 0 Hours. If you want to retain metrics data for 1 hour, indicate it as 0 Day(s) and 1 Hour(s).
How often do you want to run the truncate process on expired data?	Use to indicate the time frequency with which you want the metrics collection database data to be truncated. The default is 1 hour.

15. In order to control the traffic in the Metrics collection, you can optionally filter (exclude) the following sessions:

Option	Description
Studio Sessions	All Studio sessions of the users mentioned in the Users field will be filtered.
Internal Sessions	All Internal sessions of the users mentioned in the Users field will be filtered.
Cache & Target Trigger Task Sessions	All Trigger and Cache sessions of the users mentioned in the Users field will be filtered.
Built-in Procedures	All Built-in Procedure sessions of the users mentioned in the Users field will be filtered.
Users	Mention the list of users qualified by the domain name (user@domain) for which you want the sessions to be filtered.

Note:

- The above principal filters apply only to new sessions or requests that are received after the filter has been enabled.
- The KPI module is expected to track only new user sessions that are created after the module has been enabled.

16. Click Enable.

Using Studio to Create the Database and Tables for Metrics Data Storage

If you need to create the database and table that are used to store your usage metrics without using Studio, you can use the following instructions.

If metrics are enabled, you cannot edit the values on the page. To modify any of the values, make sure to clear the Enable check box.

To use Studio to create the metrics database and required tables

1. In the Studio resource tree, select localhost > policy > metrics.
2. Right click and select Open.
3. Browse to an existing data source that supports the storage of metrics data. See [Supported Data Source Types for Metrics Storage](#).
4. Select a schema location to hold the data.
5. Create or bind to the following tables:

Table Name	Description of Steps to Create the Table
metrics_sessions	• Click Browse. • Navigate to the data source or schema where you want to create the table.
metrics_requests	• Click Browse. • Navigate to the data source or schema where you want to create the table.
metrics_resources_usage	• Click Browse. • Navigate to the data source or schema where you want to create the table.

6. Click Execute DDL.

7. Click OK.
8. Click Save.
9. Optionally, make selections for the following:

Option	Description
Memory Threshold	Number of megabytes to set aside as the buffer space to store records. When the buffer reaches the threshold, the results are posted to the metrics tables.
Request Count Threshold	Number of records or rows to retain in the buffer. When the buffer reaches the threshold, the results are posted to the metrics tables.
How long do you want to keep the metrics data?	Use to indicate the number of days you want your metrics data retained. The default is 30 days.
How often do you want to run the truncate process on expired data?	Use to indicate the time frequency with which you want the metrics collection database data to be truncated. The default is 1 hour.

10. Click Enable.

Configuring TDV Metrics Collection

There are several configuration parameters that can be used to determine what data is collected and stored.

To configure metrics collection

1. Log into Studio as the admin user.
2. From the Administration menu, choose Configuration.
3. Locate the following configuration parameters and modify their values according to the instructions in the Description column:

Configuration Parameter	Description
Enable Metrics Events	Enables the collection of metrics. Defaults to TRUE.
Metrics Backup End	Use a comma to separate multiple choices. Multiple choices cannot include ALL or NONE. Currently supported event filters: • DB: event sent to database only. • LOG: event sent to log file only. • SNMP: event sent to SNMP processor only. • CUSTOM: event sent to custom event handler only. • ALL: event sent to database, log, SNMP processor, and custom event handler. • NONE: event ignored.
Metrics Backup Failed	
Metrics Backup Start	
Metrics Persistent End	
Metrics Persistent Failed	
Metrics Persistent Start	
Metrics Restore End	
Metrics Restore Failed	
Metrics Restore Start	
Metrics Truncation End	
Metrics Truncation Fail	
Metrics Truncation Start	
Metrics Data Lost	

4. Click Apply.

5. Click OK.

Restart of the TDV Server is not necessary for these configuration parameters to be applied to the TDV Server.

Configuring Email Alerts for Metrics Notification

You can trigger email notifications for TDV actions or events.

Tip from an expert: The following configuration parameters are left over from functionality that does not send email alerts: Email Addresses for CC, Enable Email Events, Email Addresses.

To enable email alerts

1. Open and log in to Studio.
2. From the Administration menu, choose Configuration.
3. Navigate to Server > Configuration > E-Mail.
4. Set values for the following:

Configuration Parameter	Description of Value	Example
From Address	Email address that you want to appear in the From line for alerts.	meg@queenbeesknees.net
SMTP Host Name	Name of the email server host.	javamail.queenbeesknees.com
Maximum number of rows included in email attachment.	If set to 0, there is no restriction on the size of the email attachment. If set to a value greater than 0, the value is used as the maximum number of rows allowed for the attachment.	0

5. Save and exit the Configuration window.

6. From the Studio resource tree, right-click and select New Trigger.
7. Name and enable it.
8. Set the type of event that you want to trigger the alert. For example, a system event such as a cache refresh or data source going down. For information on how to set the different types of triggers, see the *TDV User Guide*. Choose System Event to collect information for typical TDV events including, Metrics collection, caching actions, and request spikes.

Typical Event Areas	System Event Name
Metrics Alerts	MetricsPersistentFailure
	MetricsTruncationFailure
	MetricsBackupFailure
	MetricsRestoreFailure
	StatisticsGatheringFailure
Caching	CacheRefreshFailure
	CacheRefreshSuccess
Cluster Management	ClusterServerJoined
	ClusterServerConnected
	ClusterServerDisconnected
	ClusterServerShunned
Data Source Management	DataSourceDown
	DataSourceUp
Request Management	RequestFailure
	RequestInactive
	RequestRunForTooLong
	RequestsSpike
	TransactionFailure

Typical Event Areas	System Event Name
Resource Management	ResourceLock
	ResourceUnlock
Errors and Login Management	ErrorsSpike
	FailedLoginSpike
Server Management	ServerStart
	ServerStop
Trigger Management	TriggerStart
	TriggerEnd
	TriggerFail

9. Select the Action Type of Send E-mail.
10. Specify a Resource path. For example, /shared/examples/ds_orders/tutorial/customers.
11. Type the email addresses for which to send the email alerts. For example, meg@queenbeesknees.net.
12. Type a meaningful Message Subject.
13. Type a meaningful Message Body.
14. Save the trigger.

Publishing and Reporting on TDV Metrics Data

Publish data through TDV is a process that should be familiar to all TDV users. Reporting on the data will vary depending on what tools you have available to you for the forming and presentation of table data.

The published metrics tables are used to collect and store metrics on the other published TDV objects.

To publish and report on TDV metrics data

1. Select the Enable check box on the Metrics page.

The metrics tables are automatically created in the Studio resource tree under:

```
Desktop > Composite Data Services > Databases > system > metrics
```

2. Optionally, using Studio, open and execute the metrics tables to view the result set.
3. If necessary, connect your favorite reporting tool to TDV following the instructions in the *TDV Client Interfaces Guide*.
4. Review the data and design your reports.

About Using MAXMEMORY in Your Reports

The value of MAXMEMORY is coming from the SYS_REQUESTS table.

Each new request gets a 2 MB chunk for it to use. When that runs out, TDV requests more space, sometimes in 512 KB chunks. When TDV processes a SQL statement, it can use a lot of memory. TDV fetches data and sometimes does several calculations on it before the end result is delivered. If TDV detects that it is exceeding the maximum memory available, it can store the data onto the disk, to free up memory.

SNMP Trap Message Reference

The TDV system supports SNMP v3 traps. This topic provides a list of events and their corresponding SNMP traps. TDV Server generates traps for monitoring the events that occur in the server.

For a MIB definition of the SNMP traps supported in TDV, see the MIB file available in the product installation directory at:

<TDV_install_dir>\apps\server\CompositeSoftware-MIB.mib

For the procedure to set up traps, see [Enabling SNMP Traps in TDV](#).

SNMP details are grouped into tables by category:

- [SNMP Details for Monitor Events](#)
- [SNMP Details for Server Events](#)
- [SNMP Details for Requests](#)
- [SNMP Details for Transactions](#)
- [SNMP Details for Cached Resources](#)
- [SNMP Details for Triggers](#)
- [SNMP Details for Data Sources](#)
- [SNMP Details for Sessions](#)
- [SNMP Details for Resources](#)
- [SNMP Details for Storage](#)
- [SNMP Details for Server Events](#)
- [SNMP Details for Security](#)
- [SNMP Details for Workload](#)
- [SNMP Details for KPI](#)
- [SNMP Details for Audit Startup and Shutdown](#)

SNMP Details for Monitor Events

SNMP ID	Event	Variables	Description
10000	csMonitorStart	{ trapTime, trapServerHostName, trapServerPort }	A TDV Server Monitor is started.
10001	csMonitorStop	{ trapTime, trapServerHostName, trapServerPort }	A CTDV Server Monitor is stopped.
10002	csMonitorFail	{ trapTime, trapServerHostName, trapServerPort }	A TDV Server Monitor fails.
10003	csServerStopUnplanned	{ trapTime, trapServerHostName, trapServerPort }	A TDV Server has an unplanned stop.
10004	csServerStopPlanned	{ trapTime, trapServerHostName, trapServerPort }	A TDV Server has a planned stop.
10005	csServerRestart	{ trapTime, trapServerHostName, trapServerPort }	A TDV Server is restarted.
10006	csServerRestartFail	{ trapTime, trapServerHostName, trapServerPort }	A TDV Server has a restart failure.
10007	csRepositoryUp	{ trapTime, trapServerHostName, trapServerPort }	A TDV Server Repository is started.
10008	csRepositoryDown	{ trapTime, trapServerHostName, trapServerPort }	A TDV Server Repository is stopped.

SNMP Details for Server Events

SNMP ID	Event	Variable	Description
20000	csServerStart	{ trapTime, trapServerHostName, trapServerPort }	A TDV Server is started.
20001	csServerStop	{ trapTime, trapServerHostName, trapServerPort }	A TDV Server is stopped.
20002	csUserCreate	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName }	A user is created in a domain.
20003	csGroupCreate	{ trapTime, trapServerHostName, trapServerPort, trapGroupName, trapDomainName }	A group is created in a domain.
20004	csUserDelete	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName }	A user is deleted from a domain.
20005	csGroupDelete	{ trapTime, trapServerHostName, trapServerPort, trapGroupName, trapDomainName }	A group is deleted from a domain.
20006	csUserAddTo Group	{ trapTime, trapServerHostName, trapServerPort,	A user is added to a group.

SNMP ID	Event	Variable	Description
		trapUserName, trapUserDomainName, trapGroupName, trapGroupDomainName }	
20007	csUserRemove FromGroup	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapUserDomainName, trapGroupName, trapGroupDomainName }	A user is removed from a group.
20008	csDomainCreate	{ trapTime, trapServerHostName, trapServerPort, trapDomainName }	A domain is created.
20009	csDomainDelete	{ trapTime, trapServerHostName, trapServerPort, trapDomainName }	A domain is deleted.
20010	csUserPassword Modify	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName }	A user password is modified.

SNMP Details for Requests

SNMP ID	Event	Variables	Description
20100	csRequestStart	{ trapTime, trapServerHostName, trapServerPort, trapRequestId, trapOptionalRequestParameter1, trapOptionalRequestParameter2, trapOptionalRequestParameter3, trapOptionalRequestParameter4 }	A request is started.
20101	csRequestWait	{ trapTime, trapServerHostName, trapServerPort, trapRequestId, trapTransactionId, trapSessionId }	A request is waiting to run.
20102	csRequestEnd	{ trapTime, trapServerHostName, trapServerPort, trapRequestId, trapTransactionId, trapSessionId }	A request is completed.
20103	csRequestFail	{ trapTime, trapServerHostName, trapServerPort, trapRequestId, trapOptionalRequestParameter1, trapOptionalRequestParameter2, trapOptionalRequestParameter3 }	A request has failed.
20104	csRequest Cancel	{ trapTime, trapServerHostName, trapServerPort, trapRequestId, trapTransactionId, trapSessionId }	A request is cancelled.
20105	csRequestWaitQueue ThresholdPass	{ trapTime, trapServerHostName, trapServerPort, trapRequestId, trapTransactionId, trapSessionId }	A request passes the wait queue threshold.
20106	csRequestWait QueueThreshold Reset	{ trapTime, trapServerHostName, trapServerPort, trapRequestId, trapTransactionId, trapSessionId }	A request reset the wait queue threshold.

SNMP ID	Event	Variables	Description
20107	csPrepared Statement Success	{ trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapOptionalRequestParameter1, trapOptionalRequestParameter2 }	A prepared statement is successfully executed.
20108	csPrepared StatementFail	{ trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapSqlQuery, trapOptionalRequestParameter1, trapOptionalRequestParameter2 }	A prepared statement has failed during execution.

SNMP Details for Transactions

SNMP ID	Event	Description	Description
20200	csTransactionStart	{ trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapSessionId }	A transaction is started.
20201	csTransactionCommit	{ trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapSessionId }	A transaction is committed.
20202	csTransactionFail	{ trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapMessage, trapStackTrace, trapSessionId }	A transaction has failed.

SNMP ID	Event	Description	Description
20203	csTransactionRollBack	{ trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapSessionId }	A transaction is rolled back.
20204	csTransaction Compensate	{ trapTime, trapServerHostName, trapServerPort, trapTransactionId, trapMessage, trapSessionId }	A transaction is compensated for.

SNMP Details for Cached Resources

SNMP ID	Event	Variables	Description
20300	csCacheEnable	{ trapTime, trapServerHostName, trapServerPort, trapCacheName }	A cache is enabled.
20301	csCacheDisable	{ trapTime, trapServerHostName, trapServerPort, trapCacheName }	A cache is disabled.
20302	csCacheClear	{ trapTime, trapServerHostName, trapServerPort, trapCacheName, trapCacheParameters }	A cache is cleared.
20303	csCacheRefreshStart	{ trapTime, trapServerHostName,	A cache refresh is started.

SNMP ID	Event	Variables	Description
		trapServerPort, trapCacheName, trapCacheParameters }	
20304	csCacheRefreshEnd	{ trapTime, trapServerHostName, trapServerPort, trapCacheName, trapCacheParameters }	A cache refresh is completed.
20305	csCacheRefreshFail	{ trapTime, trapServerHostName, trapServerPort, trapCacheName, trapCacheParameters, trapOptionalMessage }	A cache refresh has failed.

SNMP Details for Triggers

SNMP ID	Event	Variables	Description
20400	csTriggerStart	{ trapTime, trapServerHostName, trapServerPort, trapTriggerName, trapTriggerType, trapTriggerAction }	A trigger is started.
20401	csTriggerEnd	{ trapTime, trapServerHostName, trapServerPort, trapTriggerName, trapTriggerType, trapTriggerAction }	A trigger is completed.

SNMP ID	Event	Variables	Description
20402	csTriggerFail	{ trapTime, trapServerHostName, trapServerPort, trapTriggerName, trapTriggerType, trapTriggerAction, trapOptionalMessage }	A trigger has failed.

SNMP Details for Data Sources

SNMP ID	Event	Variables	Description
20500	csDataSourceOn	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType }	A data source is enabled.
20501	csDataSourceOff	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType }	A data source is disabled.
20502	csDataSourceUp	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType }	A data source is started.
20503	csDataSourceDown	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName,	A data source is stopped.

SNMP ID	Event	Variables	Description
		trapDataSourceType }	
20504	csDataSourceModify	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType }	A data source is modified.
20505	csIntrospectStart	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType }	A data source introspection is started.
20506	csIntrospectEnd	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType, trapDataSourceReport }	A data source introspection has completed.
20507	csIntrospectCancel	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType }	A data source introspection is cancelled.
20508	csIntrospectFail	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType, trapMessage }	A data source introspection has failed.
20509	csTestStart	{ trapTime, trapServerHostName, trapServerPort,	A data source test is started.

SNMP ID	Event	Variables	Description
		trapDataSourceName, trapDataSourceType }	
20510	csTestSuccess	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType }	A data source test is successful.
20511	csTestFail	{ trapTime, trapServerHostName, trapServerPort, trapDataSourceName, trapDataSourceType }	A data source test has failed.
20512	csConnPoolSizeIncrease	{ trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId }	The size of a connection pool has increased.
20513	csConnPoolSizeDecrease	{ trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId }	The size of a connection pool has decreased.
20514	csConnCheckOut	{ trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId }	A connection is checked out a connection pool.
20515	csConnCheckIn	{ trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId }	A connection is checked into a connection pool.
20516	csConnInvalid	{ trapTime, trapServerHostName,	A connection pool has an invalid

SNMP ID	Event	Variables	Description
		trapServerPort, trapConnectionPoolId }	connection.
20517	csConnFail	{ trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId }	A connection pool has a failed connection.
20518	csConnPoolExhaust	{ trapTime, trapServerHostName, trapServerPort, trapConnectionPoolId }	A connection pool has exhausted its connections.
20519	csStatisticsProcessing StartProcess	{ trapTime, trapServerHostName, trapServerPort, trapDataSourcePath }	A data source started the statistics processing process.
20520	csStatisticsProcessing Complete	{ trapTime, trapServerHostName, trapServerPort, trapDataSourcePath }	A data source completed the statistics processing process.
20521	csStatisticsProcessing CompletePartial	This event/message is deprecated. Statistics processing with ID # is partially completed.	This event/message is deprecated.
20522	csStatisticsProcessing Failed	{ trapTime, trapServerHostName, trapServerPort, trapDataSourcePath, trapMessage }	A data source failed to complete the statistics processing process.

SNMP ID	Event	Variables	Description
20523	csStatisticsProcessing Update	This event/message is deprecated. Statistics processing with ID # updated the information.	This event/message is deprecated.

SNMP Details for Sessions

SNMP ID	Event	Variables	Description
20700	csSessionLoginFail	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName }	A session login has failed for a user.
20701	csSessionStart	{ trapTime, trapServerHostName, trapServerPort, trapSessionId, trapUserName, trapDomainName }	A session is started for a user.
20702	csSessionEnd	{ trapTime, trapServerHostName, trapServerPort, trapSessionId, trapUserName, trapDomainName }	A session is ended for a user.
20703	csSessionTerminate	{ trapTime, trapServerHostName, trapServerPort, trapSessionId,	A session is terminated for a user.

SNMP ID	Event	Variables	Description
		trapUserName, trapDomainName }	
20705	csSessionMax ConnectionsExhaust	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName, trapHostName, trapLocalHostName, trapLocalHostIP }	A session creation request is denied for a user.

SNMP Details for Resources

SNMP ID	Event	Variables	Description
20800	csResourceCreate	{ trapTime, trapServerHostName, trapServerPort, trapResourceName, trapDataSourcePath, trapResourceType }	A resource is created.
20801	csResourceDelete	{ trapTime, trapServerHostName, trapServerPort, trapResourceName, trapDataSourcePath, trapResourceType }	A resource is deleted.
20802	csStatisticsResource ProcessingStartProcess	{ trapTime, trapServerHostName, trapServerPort, trapResourcePath }	A resource starts the statistics gathering process.

SNMP ID	Event	Variables	Description
20803	csStatisticsResourceProcessingComplete	{ trapTime, trapServerHostName, trapServerPort, trapResourcePath }	A resource completes the statistics processing process.
20804	csStatisticsResourceProcessingFailed	{ trapTime, trapServerHostName, trapServerPort, trapResourcePath, trapMessage }	A resource fails to complete the statistics processing process.
20805	csResourceLock	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName, trapLockTime, trapResourcePath, trapResourceType, trapResourceSubType }	A resource is locked. Only the topmost parent node is reported as locked.
20806	csResourceUnlock	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName, trapUnlockTime, trapResourcePath, trapResourceType, trapResourceSubType, trapComment }	A resource is unlocked. Only the topmost parent node is reported as unlocked.

SNMP Details for Storage

SNMP ID	Event	Variables	Description
21000	csStorageLowWarning	{ trapTime, trapServerHostName, trapServerPort }	A storage low warning has occurred on a machine.
21001	csStorageLowCritical	{ trapTime, trapServerHostName, trapServerPort }	A storage low critical event has occurred on a machine.

SNMP Details for Server Events

SNMP ID	Event	Variables	Description
21500	csClusterServerDisconnected	{ trapTime, trapServerHostName, trapServerPort, trapClusterServerName }	This trap is generated when a server has been disconnected from the cluster.
21501	csClusterServerConnected	{ trapTime, trapServerHostName, trapServerPort, trapClusterServerName }	This trap is generated when a server has been connected to the cluster.
21502	csClusterServerShunned	{ trapTime, trapServerHostName, trapServerPort }	This trap is generated when a server has been

SNMP ID	Event	Variables	Description
			shunned from the cluster.
21503	csClusterServerJoined	{ trapTime, trapServerHostName, trapServerPort }	This trap is generated when a server has joined the cluster.
22000	csSecurityRBSCreate	{ trapTime, trapServerHostName, trapServerPort, trapPolicyName, trapResourceName }	This trap is generated when a Row Based Security policy has been created.
22001	csSecurityRBSUpdate	{ trapTime, trapServerHostName, trapServerPort, trapPolicyName, trapOriginalAssignment, trapNewAssignment }	This trap is generated when a Row Based Security policy has been updated.
22002	csSecurityRBSDelete	{ trapTime, trapServerHostName, trapServerPort, trapPolicyName, trapResourceName }	This trap is generated when a Row Based Security policy has been deleted.
22003	csSecurityRBSEnable	{ trapTime, trapServerHostName, trapServerPort }	This trap is generated when a Row Based Security policy has been enabled.

SNMP ID	Event	Variables	Description
22004	csSecurityRBSDisable	{ trapTime, trapServerHostName, trapServerPort }	This trap is generated when a Row Based Security policy has been disabled.
22005	csSecurityRBSAssign	{ trapTime, trapServerHostName, trapServerPort, trapPolicyName, trapResourceName }	This trap is generated when a Row Based Security policy has been assigned.
22006	csSecurityRBSRemove	{ trapTime, trapServerHostName, trapServerPort, trapPolicyName, trapResourceName }	This trap is generated when a Row Based Security policy has been removed.
22007	csSecurityUserLockStatus	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapUserLocked, trapUserLockedBy }	This trap is generated when the user lock status has been changed.
22008	csSecurityUserImplicitlyLocked	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName }	This trap is generated when a user has been implicitly locked.

SNMP Details for Security

SNMP ID	Event	Variables	Description
22010	csSecurityCBSCreate	{ trapTime, trapServerHostName, trapServerPort, trapPolicyName }	This trap is generated when a cbs policy is created.
22011	csSecurityCBSUpdate	{ trapTime, trapServerHostName, trapServerPort, trapPolicyName }	This trap is generated when a cbs policy is updated.
22012	csSecurityCBSDelete	{ trapTime, trapServerHostName, trapServerPort, trapPolicyName }	This trap is generated when a cbs policy is deleted.
22013	csSecurityCBSEnable	{ trapTime, trapServerHostName, trapServerPort }	This trap is generated when a cbs feature is enabled.
22014	csSecurityCBSDisable	{ trapTime, trapServerHostName, trapServerPort }	This trap is generated when a cbs feature is disabled.
22015	csSecurityCBSAssign	{ trapTime, trapServerHostName, trapServerPort, trapResourcePath,	This trap is generated when a cbs policy is

SNMP ID	Event	Variables	Description
		trapColumnName, trapPolicyName }	assigned to a resource column.
22016	csSecurityCBSUpdateAssign	{ trapTime, trapServerHostName, trapServerPort, trapResourcePath, trapColumnName, trapOriginalAssignment, trapNewAssignment }	This trap is generated when a cbs policy assignment is updated.
22017	csSecurityCBSDeAssign	{ trapTime, trapServerHostName, trapServerPort, trapResourcePath, trapColumnName, trapPolicyName }	This trap is generated when a cbs policy is removed from a resource column.

SNMP Details for Workload

SNMP ID	Event	Variables	Description
	csWorkloadDelete	{ trapTime, trapServerHostName, trapServerPort, trapRuleName }	This trap is generated when a workload rule is deleted.
	csWorkloadEvent	{ trapTime, trapServerHostName, trapServerPort, trapRuleName }	This trap is generated when a workload rule is violated.

SNMP ID	Event	Variables	Description
	csWorkloadCreate	{ trapTime, trapServerHostName, trapServerPort, trapRuleName }	This trap is generated when a workload rule is created.
	csWorkloadEnable	{ trapTime, trapServerHostName, trapServerPort }	This trap is generated when Workload management feature is enabled.
	csWorkloadDisable	{ trapTime, trapServerHostName, trapServerPort }	This trap is generated when Workload management feature is disabled.
	csWorkloadUpdate	{ trapTime, trapServerHostName, trapServerPort, trapRuleName }	This trap is generated when a workload rule is updated.
	csWorkloadRename	{ trapTime, trapServerHostName, trapServerPort, trapRuleName, trapNewRuleName }	This trap is generated when a workload rule is renamed.

SNMP Details for KPI

SNMP ID	Event	Variables	Description
	csKPIPersistentStart	{ trapTime, trapServerHostName, trapServerPort, trapMessage }	This trap is generated when KPI persistent is started.
	csKPIPersistentEnd	{ trapTime, trapServerHostName, trapServerPort, trapMessage }	This trap is generated when KPI persistent ends.
	csKPIPersistentFailed	{ trapTime, trapServerHostName, trapServerPort, trapMessage, trapException }	This trap is generated when KPI persistent fails.
	csKPITruncationStart	{ trapTime, trapServerHostName, trapServerPort, trapMessage }	This trap is generated when KPI truncation starts.
	csKPITruncationEnd	{ trapTime, trapServerHostName, trapServerPort, trapMessage }	This trap is generated when KPI truncation ends.
	csKPITruncationFailed	{ trapTime, trapServerHostName, trapServerPort, trapMessage, trapException }	This trap is generated when KPI truncation fails
23006	csKPIBackupStart	{ trapTime, trapServerHostName, trapServerPort, trapMessage }	This trap is generated when KPI backup starts.

SNMP ID	Event	Variables	Description
23007	csKPIBackupEnd	{ trapTime, trapServerHostName, trapServerPort, trapMessage }	This trap is generated when KPI backup ends.
23008	csKPIBackupFailed	{ trapTime, trapServerHostName, trapServerPort, trapMessage, trapException }	This trap is generated when a KPI backup fails.
23009	csKPIRestoreStart	{ trapTime, trapServerHostName, trapServerPort, trapMessage }	This trap is generated when KPI restore starts.
23010	csKPIRestoreEnd	{ trapTime, trapServerHostName, trapServerPort, trapMessage }	This trap is generated when KPI restore ends.
23011	csKPIRestoreFailed	{ trapTime, trapServerHostName, trapServerPort, trapMessage, trapException }	This trap is generated when KPI restore fails.

SNMP Details for Audit Startup and Shutdown

SNMP ID	Event	Variables	Description
24000	csAuditConfigChange	{trapTime, trapServerHostName, trapServerPort,	This trap is generated when Audit

SNMP ID	Event	Variables	Description
		trapPropertyName, trapOldValue, trapNewValue}	config change happens.
24001	csAuditSessionAuthFail	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName }	This trap is generated when audit session auth fails.
24002	csAuditFetchRequestInvalid	{ trapTime, trapServerHostName, trapServerPort, trapUserName, trapDomainName }	This trap is generated when audit fetch request is invalid.
24003	csAuditSessionIdValidationFailed	{ trapTime, trapServerHostName, trapServerPort, trapSessionId }	This trap is generated when audit session Id validation fails.

TDV Event Log Message Reference

TDV generates over 20 different log files that contain hundreds of different messages. This reference attempts to capture some of the event log messages. It is not a comprehensive list of all messages generated by TDV.

- [TDV Trap Messages and Variables](#)
- [TDV Monitor Events](#)
- [TDV Server Events](#)
- [Active Cluster Events](#)

TDV Trap Messages and Variables

- Trap messages and variables are described in [SNMP Trap Message Reference](#) .

TDV Monitor Events

The config.key entry is only needed for monitor event logging. The server event logging uses the ServerEvents class to perform event filtering. These messages can be found in <TDV_install_dir>/logs/cs_server_events.log.

Number	Event	Description	Severity	Status
10000	monitorStart	TDV Server Monitor start	INFO	START
10001	monitorStop	TDV Server Monitor stop	INFO	STOP
10002	monitorFail	TDV Server	ERROR	FAIL

Number	Event	Description	Severity	Status
		Monitor fail		
10003	serverStopUnplanned	TDV Server unplanned stop	ERROR	FAIL
10004	serverStopPlanned	TDV Server planned stop	INFO	STOP
10005	serverRestart	TDV Server restart started	INFO	RESTART
10006	serverRestartFail	TDV Server restart failed	ERROR	FAIL
10007	serverRepositoryUp	TDV Server Repository up	INFO	UP
10008	serverRepositoryDown	TDV Server Repository down	ERROR	FAIL

TDV Server Events

These messages can be found in <TDV_install_dir>/logs/cs_server_events.log.

Number	Event	Description	Severity	Status
20000	csServerStart	Composite Server started	INFO	
20001	csServerStop	Composite Server stopped	INFO	
20002	csUserCreate	user={0}	INFO	

Number	Event	Description	Severity	Status
		created in domain={1}		
20003	csGroupCreate	group={0} created in domain={1}	INFO	
20004	csUserDelete	user={0} deleted in domain={1}	INFO	
20005	csGroupDelete	group={0} deleted in domain={1}	INFO	
20006	csUserAddToGroup	user={0} in domain={1} added to group={2} in domain={3}	INFO	
20007	csUserRemoveFromGroup	user={0} in domain={1} removed from group= {2} in domain={3}	INFO	
20008	csDomainCreate	created domain={0}	INFO	
20009	csDomainDelete	deleted domain={0}	INFO	
20010	csUserPasswordModify	password changed for user={0} in	INFO	

Number	Event	Description	Severity	Status
		domain={1}		
20100	csRequestStart	request id={0} started	INFO	
20101	csRequestWait	request id={0} waiting to run		
20102	csRequestEnd	request id={0} completed		
20103	csRequestFail	request id={0} failed. exception={1}		
20104	csRequestCancel	request id={0} canceled		
20105	csRequestWaitQueueThreshold Pass	wait queue threshold passed with request id={0}		
20106	csRequestWaitQueueThreshold Reset	wait queue threshold reset with request id={0}		
20107	csPreparedStatementSuccess	prepared statement with transaction		

Number	Event	Description	Severity	Status
		id={0} successful. sql={1}		
20108	csPreparedStatementFail	prepared statement with transaction id={0} failed. sql={1} exception= {2}		
20109	csRequestInactive	request is inactive for {3} minutes: session id= {2},transacti on id= {1},request id={0},sql= {4}	WARNIN G	WARNING
20110	csRequestRunForTooLong	request has run for {3} minutes: session id= {2},transacti on id= {1},request id={0},sql= {4}	WARNIN G	WARNING
20200	csTransactionStart	transaction id={0} started		

Number	Event	Description	Severity	Status
20201	csTransactionCommit	transaction id={0} committed		
20202	csTransactionFail	transaction id={0} failed. message={0} stack trace= {1}		
20203	csTransactionRollBack	transaction id={0} rolled back		ROLLBACK
20204	csTransactionCompensate	transaction id={0} compensate d. message= {1}		COMPENSATE
20300	csCacheEnable	cache={0} enabled		
20301	csCacheDisable	cache={0} disabled		
20302	csCacheClear	cache={0} cleared, params={1}		
20303	csCacheRefreshStart	cache refresh={0} started, params={1}		
20304	csCacheRefreshEnd	cache refresh={0}		

Number	Event	Description	Severity	Status
		end, params={1}		
20305	csCacheRefreshFail	cache refresh={0} failed, params={1}, message={2}		
20400	csTriggerStart	trigger name={0} type={1} action={2}		
20401	csTriggerEnd	trigger name={0} type={1} action={2}		
20402	csTriggerFail	trigger name={0} type={1} action={2}, reason failed: {3}		
20500	csDataSourceOn	data source= {0} with type={1} on		
20501	csDataSourceOff	data source= {0} with type={1} off		
20502	csDataSourceUp	data source= {0} with type={1} up		

Number	Event	Description	Severity	Status
20503	csDataSourceDown	data source={0} with type={1} down		
20504	csDataSourceModify	data source={0} with type={1} modified		
20505	csIntroSpectStart	introspection of data source={0} with type={1} started		
20506	csIntroSpectEnd	introspection of data source={0} with type={1} completed. Report={2}.		
20507	csIntroSpectCancel	introspection of data source={0} with type={1} cancelled		
20508	csIntroSpectFail	introspection of data source={0} with type={1} failed. Exception={2}.		

Number	Event	Description	Severity	Status
20509	csTestStart	test data source={0} with type={1} started		
20510	csTestSuccess	test data source={0} with type={1} successful		
20511	csTestFail	test data source={0} with type={1} failed		
20512	csConnPoolSizeIncrease	connection pool={0} size increased		
20513	csConnPoolSizeDecrease	connection pool={0} size decreased		
20514	csConnCheckOut	connection pool={0} has connection checked out		CHECK_OUT
20515	csConnCheckIn	connection pool={0} has connection checked in		
20516	csConnInvalid	connection pool={0} has invalid connection		

Number	Event	Description	Severity	Status
20517	csConnFail	connection pool={0} has failed connection		
20518	csConnPoolExhaust	connection pool={0} exhausted		
20519	csStatisticsDSProcessingStartProcess	statistics processing={0} started processing		
20520	csStatisticsDSProcessingComplete	statistics processing={0} completed processing successfully		
20522	csStatisticsDSProcessingFailed	statistics processing={0} failed to process		
20701	csSessionStart	Session={0} started for user={1} in domain={2}		
20702	csSessionEnd	Session={0} completed for user={1} in domain={2}		

Number	Event	Description	Severity	Status
20703	csSessionTerminate	Session={0} terminated for user={1} in domain= {2}		
20704	csSessionNonLocalhostRequestFail	Session creation denied for user={0} in domain={1} at hostname= {2}. Only session creation requests from localhost, 127.0.0.1, {3} and {4} are allowed.		
20705	csSessionMaxConnectionsExhaust	Session creation denied for user={0} in domain={1} at hostname= {2}. Maximum session connection limit of {3} has been exhausted.		

Number	Event	Description	Severity	Status
20706	csSessionRunForTooLong	Session run for {3} minutes for user={0} in domain={1} at hostname={2}.	WARNING	
20707	csSessionLoginUserFail	Session login failed in domain={0}		
20708	csSessionLoginPwdFail	session login failed for user={0} in domain={1}		
20800	csResourceCreate	resource={0} with path={1} and type={2} created		
20801	csResourceDelete	resource={0} with path={1} and type={2} deleted		
20802	csStatisticsResourceProcessingStartProcess	statistics processing={0} started processing		
20803	csStatisticsResourceProcessingComplete	statistics processing={0} completed		

Number	Event	Description	Severity	Status
		processing successfully		
20804	csStatisticsResourceProcessing Failed	statistics processing={0} failed to process		
20805	csResourceLock	locked resource at path={3}		
20806	csResourceUnlock	unlocked resource at path={3}		
21000	csStorageLowWarning	storage low warning	WARNING	
21001	csStorageLowCritical	storage low critical		
21500	csClusterServerDisconnected	server {0} has been disconnected	WARNING	DISCONNECTED
21501	csClusterServerConnected	server {0} has been activated		
21502	csClusterServerShunned	server has been shunned		
21503	csClusterServerJoined	server has joined to the		

Number	Event	Description	Severity	Status
		cluster		
22000	csSecurityRBSCreate	RBS {0} has been created		
22001	csSecurityRBSUpdate	RBS {0} has been updated		
22002	csSecurityRBSDelete	RBS {0} has been deleted		
22003	csSecurityRBSEnable	RBS Feature has been enabled		
22004	csSecurityRBSDisable	RBS Feature has been disabled		
22005	csSecurityRBSAssign	RBS {0} has been assigned		
22006	csSecurityRBSRemove	RBS {0} has been removed		
22007	csSecurityUserLockStatus	User {2} has changed the account lock status of user {0} to {1}		
22008	csSecurityUserImplicitlyLocked	User {0}@{1} has become		

Number	Event	Description	Severity	Status
		implicitly locked		
23000	csKPIPersistentStart	{0}		
23001	csKPIPersistentEnd	{0}		
23002	csKPIPersistentFailed	{0} exception= {1}		
23003	csKPITruncationStart	{0}		
23004	csKPITruncationEnd	{0}		
23005	csKPITruncationFailed	{0} exception= {1}		
23006	csKPIBackupStart	{0}		
23007	csKPIBackupEnd	{0}		
23008	csKPIBackupFailed	{0} exception= {1}		
23009	csKPIRestoreStart	{0}		
23010	csKPIRestoreEnd	{0}		
23011	csKPIRestoreFailed	{0} exception= {1}		
22200	csAuditConfigChange	property={0} has been		

Number	Event	Description	Severity	Status
		changed from {1} to {2}		
22201	csAuditSessionAuthFail	Session authentication for user={0} in domain={1} has been failed		
22202	csAuditFetchRequestInvalid	Invalid dataset fetch request from user={0} in domain={1}		

Active Cluster Events

These messages can be found in <TDV_install_dir>/logs/cluster/cs_cluster.log

Event Message Code	Description
cluster.CODE_REGROUPING_ALREADY_IN_PROGRESS	Regrouping initiated by node {0} is already in progress on node {1}.
cluster.CODE_REGROUPING_FAILED_NOT_ANESHEZIZED	Regrouping failed on node {0}.
cluster.CODE_REGROUPING_FAILED_UNEXPECTED_ARBITER	Regrouping failed on node {0}. New arbiter {1} not expected. Expected: {2}
cluster.CODE_REGROUPING_	Regrouping failed on node {0}. Node cluster

Event Message Code	Description
FAILED_COULD_NOT_RESET_NODE_CLUSTER_CONNECTION	connection could not be reset. The node must be restarted. Cause: {1}
cluster.CODE_REGROUPING_FAILED	Regrouping failed. Cause: {0}
cluster.CODE_REGROUPING_FAILED_TIMEOUT	Regrouping timed out after {0} seconds. {1}
cluster.CODE_REGROUP_OUTCOME	Regroup outcome - Attempted to regroup {0} node(s): {1}. \nActivated {2} node(s): {3}.
cluster.CODE_REGROUP_NODE	Attempted to regroup node: {0}
cluster.CODE_REGROUP_OUTCOME_ACTIVATED_NODE	Activated node: {0}.
cluster.CODE_REGROUP_OUTCOME_FAILED_TO_ACTIVATE	Failed to activate {0} node(s): {1} cluster.CODE_REGROUP_OUTCOME_FAILED_TO_ACTIVATE_NODE = Failed to activate node: {0} cluster.CODE_REGROUP_OUTCOME_SPLIT = Detected metadata SPLIT and evicted {0} node(s): {1} cluster.CODE_REGROUP_OUTCOME_SPLIT_NODE = Detected metadata SPLIT and evicted node: {0}

TIBCO Documentation and Support Services

For information about this product, you can read the documentation, contact TIBCO Support, and join TIBCO Community.

How to Access TIBCO Documentation

Documentation for TIBCO products is available on the [Product Documentation website](#), mainly in HTML and PDF formats.

The [Product Documentation website](#) is updated frequently and is more current than any other documentation included with the product.

Product-Specific Documentation

The following documentation for this product is available on the [TIBCO® Data Virtualization](#) page.

Users

- TDV Getting Started Guide
- TDV User Guide
- TDV Web UI User Guide
- TDV Client Interfaces Guide
- TDV Tutorial Guide
- TDV Northbay Example

Administration

- TDV Installation and Upgrade Guide
- TDV Administration Guide
- TDV Active Cluster Guide
- TDV Security Features Guide

Data Sources

- TDV Adapter Guides

TDV Data Source Toolkit Guide (Formerly Extensibility Guide)

References

TDV Reference Guide

TDV Application Programming Interface Guide

Other

TDV Business Directory Guide

TDV Discovery Guide

TDV and Business Directory Release Notes - Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

Release Version Support

TDV 8.5 and 8.8 are designated as Long Term Support (LTS) versions. Some release versions of TIBCO® Data Virtualization products are selected to be long-term support (LTS) versions. Defect corrections will typically be delivered in a new release version and as hotfixes or service packs to one or more LTS versions. See also [Long Term Support](#).

How to Contact Support for TIBCO Products

You can contact the Support team in the following ways:

- To access the Support Knowledge Base and getting personalized content about products you are interested in, visit our [product Support website](#).
- To create a Support case, you must have a valid maintenance or support contract with a Cloud Software Group entity. You also need a username and password to log in to the [product Support website](#). If you do not have a username, you can request one by clicking **Register** on the website.

How to Join TIBCO Community

TIBCO Community is the official channel for TIBCO customers, partners, and employee subject matter experts to share and access their collective experience. TIBCO Community offers access to Q&A forums, product wikis, and best practices. It also offers access to extensions, adapters, solution accelerators, and tools that extend and enable customers to gain full value from TIBCO products. In addition, users can submit and vote

on feature requests from within the [TIBCO Ideas Portal](#). For a free registration, go to [TIBCO Community](#).

Legal and Third-Party Notices

SOME CLOUD SOFTWARE GROUP, INC. ("CLOUD SG") SOFTWARE AND CLOUD SERVICES EMBED, BUNDLE, OR OTHERWISE INCLUDE OTHER SOFTWARE, INCLUDING OTHER CLOUD SG SOFTWARE (COLLECTIVELY, "INCLUDED SOFTWARE"). USE OF INCLUDED SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED CLOUD SG SOFTWARE AND/OR CLOUD SERVICES. THE INCLUDED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER CLOUD SG SOFTWARE AND/OR CLOUD SERVICES OR FOR ANY OTHER PURPOSE.

USE OF CLOUD SG SOFTWARE AND CLOUD SERVICES IS SUBJECT TO THE TERMS AND CONDITIONS OF AN AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER AGREEMENT WHICH IS DISPLAYED WHEN ACCESSING, DOWNLOADING, OR INSTALLING THE SOFTWARE OR CLOUD SERVICES (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH LICENSE AGREEMENT OR CLICKWRAP END USER AGREEMENT, THE LICENSE(S) LOCATED IN THE "LICENSE" FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENT IS SUBJECT TO THOSE SAME TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

This document is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of Cloud Software Group, Inc.

TIBCO, TIBCO logo, TIBCO O logo, ActiveSpaces, Enterprise Messaging Service, Spotfire, TERR, S-PLUS, and S+ are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

All other product and company names and marks mentioned in this document are the property of their respective owners and are mentioned for identification purposes only. You acknowledge that all rights to these third party marks are the exclusive property of their respective owners. Please refer to Cloud SG's Third Party Trademark Notices (<https://www.cloud.com/legal>) for more information.

This document includes fonts that are licensed under the SIL Open Font License, Version 1.1, which is available at: <https://scripts.sil.org/OFL>

Copyright (c) Paul D. Hunt, with Reserved Font Name Source Sans Pro and Source Code Pro.

Cloud SG software may be available on multiple operating systems. However, not all operating system platforms for a specific software version are released at the same time. See the “readme” file for the availability of a specific version of Cloud SG software on a specific operating system platform.

THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENT COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THIS DOCUMENT. CLOUD SG MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S), THE PROGRAM(S), AND/OR THE SERVICES DESCRIBED IN THIS DOCUMENT AT ANY TIME WITHOUT NOTICE.

THE CONTENTS OF THIS DOCUMENT MAY BE MODIFIED AND/OR QUALIFIED, DIRECTLY OR INDIRECTLY, BY OTHER DOCUMENTATION WHICH ACCOMPANIES THIS SOFTWARE, INCLUDING BUT NOT LIMITED TO ANY RELEASE NOTES AND “README” FILES.

This and other products of Cloud SG may be covered by registered patents. For details, please refer to the Virtual Patent Marking document located at <https://www.cloud.com/legal>.

Copyright © 2002-2025. Cloud Software Group, Inc. All Rights Reserved.