

Directory Integrator
Version 7.0

Installation and Administrator Guide



Directory Integrator
Version 7.0

Installation and Administrator Guide



Note

Note: Before using this information and the product it supports, read the general information under Appendix C, "Notices," on page 275.

Product Version 7.0

This edition applies to version 7.0 of the IBM Tivoli Directory Integrator and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2003, 2009.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Preface

This document contains the information that you need to develop solutions using components that are part of the IBM® Tivoli® Directory Integrator.

Who should read this book

This book is intended for those responsible for the development, installation and administration of solutions with the IBM Tivoli Directory Integrator.

Tivoli Directory Integrator components are designed for network administrators who are responsible for maintaining user directories and other resources. This document assumes that you have practical experience installing and using both IBM Tivoli Directory Integrator.

The reader should be familiar with the concepts and the administration of the systems that the developed solutions connect to. Depending on the solution, these could include, but are not limited to, one or more of the following products, systems and concepts:

- IBM Directory Server
- IBM Tivoli Identity Manager
- IBM Java Runtime Environment (JRE) or Sun Java Runtime Environment
- Microsoft Active Directory
- Windows and UNIX operating systems
- Security management
- Internet protocols, including HTTP, HTTPS and TCP/IP
- Lightweight Directory Access Protocol (LDAP) and directory services
- Supported user registry
- Authentication and authorization
- SAP ABAP Application Server

Publications

Read the descriptions of the IBM Tivoli Directory Integrator library and the related publications to determine which publications you might find helpful. After you determine the publications you need, refer to the instructions for accessing publications online.

IBM Tivoli Directory Integrator library

The publications in the Tivoli Directory Integrator library are:

IBM Tivoli Directory Integrator V7.0 Getting Started

A brief tutorial and introduction to Tivoli Directory Integrator 7.0. Includes examples to create interaction and hands-on learning of IBM Tivoli Directory Integrator.

IBM Tivoli Directory Integrator V7.0 Installation and Administrator Guide

Includes complete information about installing, migrating from a previous version, configuring the logging functionality, and the security model underlying the Remote Server API of IBM Tivoli Directory Integrator. Contains information on how to deploy and manage solutions.

IBM Tivoli Directory Integrator V7.0 Users Guide

Contains information about using IBM Tivoli Directory Integrator 7.0. Contains instructions for

designing solutions using the Tivoli Directory Integrator designer tool (**ibmditk**) or running the ready-made solutions from the command line (**ibmdisrv**). Also provides information about interfaces, concepts and AssemblyLine creation.

IBM Tivoli Directory Integrator V7.0 Reference Guide

Contains detailed information about the individual components of IBM Tivoli Directory Integrator 7.0: Connectors, Function Components, Parsers and so forth – the building blocks of the AssemblyLine.

IBM Tivoli Directory Integrator V7.0 Problem Determination Guide

Provides information about IBM Tivoli Directory Integrator 7.0 tools, resources, and techniques that can aid in the identification and resolution of problems.

IBM Tivoli Directory Integrator V7.0 Messages Guide

Provides a list of all informational, warning and error messages associated with the IBM Tivoli Directory Integrator 7.0.

IBM Tivoli Directory Integrator V7.0 Password Synchronization Plug-ins Guide

Includes complete information for installing and configuring each of the five IBM Password Synchronization Plug-ins: Windows Password Synchronizer, Sun Directory Server Password Synchronizer, IBM Directory Server Password Synchronizer, Domino Password Synchronizer and Password Synchronizer for UNIX and Linux. Also provides configuration instructions for the LDAP Password Store and JMS Password Store.

IBM Tivoli Directory Integrator V7.0 Release Notes

Describes new features and late-breaking information about IBM Tivoli Directory Integrator 7.0 that did not get included in the documentation.

Related publications

Information related to the IBM Tivoli Directory Integrator is available in the following publications:

- IBM Tivoli Directory Integrator 7.0 uses the JNDI client from Sun Microsystems. For information about the JNDI client, refer to the *Java Naming and Directory Interface™ Specification* on the Sun Microsystems Web site at <http://java.sun.com/j2se/1.5.0/docs/guide/jndi/index.html>.
- The Tivoli Software Library provides a variety of Tivoli publications such as white papers, datasheets, demonstrations, redbooks, and announcement letters. The Tivoli Software Library is available on the Web at: <http://www.ibm.com/software/tivoli/library/>
- The *Tivoli Software Glossary* includes definitions for many of the technical terms related to Tivoli software. The *Tivoli Software Glossary* is available on the World-Wide Web, in English only, at <http://publib.boulder.ibm.com/tividd/glossary/tivoliglossarymst.htm>

Accessing publications online

The publications for this product are available online in Portable Document Format (PDF) or Hypertext Markup Language (HTML) format, or both in the Tivoli software library: <http://www.ibm.com/software/tivoli/library>.

To locate product publications in the library, click the **Product manuals** link on the left side of the Library page. Then, locate and click the name of the product on the Tivoli software information center page.

Information is organized by product and includes READMEs, installation guides, user's guides, administrator's guides, and developer's references as necessary.

Note: To ensure proper printing of PDF publications, select **Fit to page** in the Adobe Acrobat Print window (which is available when you click **File->Print**).

Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use information technology products successfully. With Tivoli Directory Integrator 7.0, you can use assistive technologies to hear and navigate the interface. After installation you also can use the keyboard instead of the mouse to operate all features of the graphical user interface.

Accessibility features

The following list includes the major accessibility features in Tivoli Directory Integrator 7.0:

- Supports keyboard-only operation.
- Supports interfaces commonly used by screen readers.
- Discerns keys as tactually separate, and does not activate keys just by touching them.
- Avoids the use of color as the only way to communicate status and information.
- Provides accessible documentation.

Keyboard navigation

This product uses standard Microsoft Windows navigation keys for common Windows actions such as access to the File menu, copy, paste, and delete. Actions that are unique to Tivoli Directory Integrator use Tivoli Directory Integrator keyboard shortcuts. Keyboard shortcuts have been provided wherever needed for all actions.

Interface Information

The accessibility features of the user interface and documentation include:

- Steps for changing fonts, colors, and contrast settings in the Configuration Editor:
 1. Type **Alt-W** to access the Configuration Editor **Window** menu. Using the downward arrow, select **Preferences...** and press **Enter**.
 2. Under the **Appearance** tab, select **Colors and Fonts** settings to change the fonts for any of the functional areas in the Configuration Editor.
 3. Under **View and Editor Folders**, select the colors for the Configuration Editor, and by selecting colors, you can also change the contrast.
- Steps for customizing keyboard shortcuts, specific to IBM Tivoli Directory Integrator:
 1. Type **Alt-W** to access the Configuration Editor **Window** menu. Using the downward arrow, select **Preferences...**
 2. Using the downward arrow, select the General category; right arrow to open this, and type downward arrow until you reach the entry **Keys**.
Underneath the **Scheme** selector, there is a field, the contents of which say "type filter text." Type `tivoli directory integrator` in the filter text field. All specific Tivoli Directory Integrator shortcuts are now shown.
 3. Assign a keybinding to any Tivoli Directory Integrator command of your choosing.
 4. Click **Apply** to make the change permanent.

The Configuration Editor is a specialized instance of an Eclipse workbench. More detailed information about accessibility features of applications built using Eclipse can be found at <http://help.eclipse.org/help33/topic/org.eclipse.platform.doc.user/concepts/accessibility/accessmain.htm>

- The information center and its related publications are accessibility-enabled for the JAWS screen reader and the IBM Home Page Reader. You can operate all documentation features using the keyboard instead of the mouse.

Vendor software

The IBM Tivoli Directory Integrator installer uses the InstallShield Multiplatform 11.5 wizard.

The IBM Tivoli Directory Integrator 7.0 installer has accessibility features that are independent from the product. The installer supports 3 UI modes:

GUI Keyboard-only operation is supported in GUI mode, and the use of a screen reader is possible. In order to get the most from a screen reader, you should use the Java Access Bridge and launch the installer with a Java access Bridge enabled JVM, for example:

```
install_tdiv70_win_x86.exe -is:javahome "C:\Program Files\IBM\Java50"
```

The JVM used should be a 1.5 JRE.

Console

In console mode, keyboard-only operation is supported and all displays and user options are displayed as text that can be easily read by screen readers. There is also an ISMP option, -accessibility, that can be used when running the installer which not only makes the installer run in console mode, but makes it makes the information appear on the console in a way that makes more sense when detected by a screen reader:

```
install_tdiv70_win_x86.exe -accessibility "C:\Program Files\IBM\Java50"
```

Console mode is the suggested install method for accessibility.

Silent In silent mode, user responses are given through a response file, and no user interaction is required.

Related accessibility information

Visit the *IBM Accessibility Center* at <http://www.ibm.com/able> for more information about IBM's commitment to accessibility.

Contacting IBM Software support

Before contacting IBM Tivoli Software support with a problem, refer to IBM System Management and Tivoli software Web site at:

<http://www.ibm.com/software/sysmgmt/products/support/IBMDirectoryIntegrator.html>

If you need additional help, contact software support by using the methods described in the *IBM Software Support Handbook* at the following Web site:

<http://techsupport.services.ibm.com/guides/handbook.html>

The guide provides the following information:

- Registration and eligibility requirements for receiving support.
- Telephone numbers and e-mail addresses, depending on the country in which you are located.
- A list of information you must gather before contacting customer support.

A list of most requested documents as well as those identified as valuable in helping answer your questions related to IBM Tivoli Directory Integrator can be found at <http://www-01.ibm.com/support/docview.wss?rs=697&uid=swg27009673>.

Contents

Preface iii

Who should read this book	iii
Publications	iii
IBM Tivoli Directory Integrator library	iii
Related publications	iv
Accessing publications online	iv
Accessibility	v
Accessibility features	v
Keyboard navigation	v
Interface Information	v
Vendor software	v
Related accessibility information	vi
Contacting IBM Software support	vi

Chapter 1. Introduction 1

IBM Tivoli Directory Integrator 7.0 Editions	1
--	---

Chapter 2. Installation instructions for IBM Tivoli Directory Integrator 3

Before you install	3
Disk space requirements	3
Memory requirements	3
Platform requirements	3
Components in IBM Tivoli Directory Integrator	3
Other requirements	6
Installing IBM Tivoli Directory Integrator	8
Launching the appropriate installer	9
Using the platform-specific TDI installer	12
Installing using the command line	23
Performing a silent install	24
Post-installation steps	24
Installing local Help files	26
Deploying AMC to a custom ISC SE/AE	27
Installing or Updating using the Eclipse Update Manager	28
Post-installation steps	29
Uninstalling	29
Launching the uninstaller	29
Performing a silent uninstallation	31
Default installation locations	31

Chapter 3. Update Installer 33

The .registry file	34
Installing fixes	36
Rollback	36
Troubleshooting	36

Chapter 4. Supported platforms 37

Virtualization support	39
----------------------------------	----

Chapter 5. Migrating 41

Migrate files to a different location	41
Which files do not need to be modified to be used in another location?	41

Which files need to be modified before they can be used in another location?	42
Which files should not be used in another location under normal circumstances?	42
Migrating files that contain encrypted data	43
Migrate files to a newer version	43
Installer-assisted migration	43
Tool-assisted migration	43
Manual migration	44
Backing up important data	51
Migrating AMC 7.0 configuration settings to another AMC deployment	52
Converting from EventHandlers to corresponding AssemblyLines	53
TCP Server Connector	54
Mailbox Connector	54
JMX Connector	55
SNMP Server Connector	55
IBM Directory Server Changelog Connector	55
HTTP Server Connector	56
LDAP Server Connector	56
Netscape/iPlanet/Sun Directory Changelog Connector	56
Active Directory Change Detection Connector	57
z/OS LDAP Changelog Connector	58
DSMLv2SOAPServerConnector	58
Migrating BTree tables and BTree Connector to System Store	59
Migrating Cloudscape database to Derby	60
Migrating global and solution properties files using the migration tool	60

Chapter 6. Security and TDI. 63

Introduction	63
Secure Sockets Layer (SSL) Support	63
Server SSL configuration of TDI components	64
Client SSL configuration of TDI components	65
SSL client authentication	65
IBM Tivoli Directory Integrator and Microsoft Active Directory SSL configuration	65
Summary of properties for enabling SSL and PKCS#11 support	67
Keystore and truststore management	68
SSL example	71
Remote Server API	72
Introduction	72
Configuring the Server API	73
Server API access options	76
Server API SSL remote access	76
Server API authentication	77
Server API Authorization	85
Server Audit Capabilities	90
TDI Server Instance Security	92
Stash File	93
Server Security Modes	93

Working with encrypted TDI configuration files	94
Standard TDI encryption of global.properties or solution.properties	96
Encryption of properties in external property files	96
The TDI Encryption utility	96
The TDI secret key tool	98
TDI System Store Security	101
Miscellaneous Config File features	102
The "password" configuration parameter type	102
Component Password Protection	102
Protecting attributes from being printed in clear text during tracing	103
Encryption of TDI Server Hooks	104
Remote Configuration Editor and SSL	104
Using the Remote Configuration Editor	104
Summary of configuration files and properties dealing with security	105
Web Admin Console Security	108
Miscellaneous security aspects	108
Ports and files used by TDI	108
HTTP Basic Authentication	108
Lotus Domino SSL specifics	108
Certificates for the TDI Web service Suite	109
MQe authentication with mini-certificates	109

Chapter 7. Reconnect Rule Engine . . . 111

Introduction	111
Reconnect Rules	111
User-defined rules configuration	112
Examples	113
Exception considerations	114
General reconnect configuration	114

Chapter 8. System Queue . . . 117

System Queue Configuration	117
WebSphere MQe parameters	118
WebSphere MQ parameters	118
Microbroker parameters	118
JMSScript Driver parameters	119
System Queue Configuration Example	120
Security and Authentication	121
MQe Configuration Utility	121
Authentication of MQe messages to provide MQe Queue Security	122
Support for DNS names in the configuration of the MQe Queue	123
Configuration of High Availability for MQe transport of password changes	123
Providing remote configuration capabilities in the MQe Configuration Utility	123

Chapter 9. Encryption and FIPS mode 125

Configuring Tivoli Directory Integrator to run FIPS mode	125
Symmetric cipher support	125
Configuring SSL and PKI certificates	133
Encrypting and decrypting using CryptoUtils	134
Working with certificates	134
Using cryptographic keys located on hardware devices	135

Using IBMPCSKS11 to access devices and to store SSL keys and certificates	136
Enabling or disabling padding	136
Maintaining encryption artifacts – keys, certificates, keystores, encrypted files	137
Changed encryption key	137
Changed password for encryption key or keystore	137
Expired encryption certificate	137

Chapter 10. Configuring the TDI Server API. . . 139

Server ID	139
Exception for password protected Configs	139
Server RMI	139
Config load time-out interval	140

Chapter 11. Properties . . . 141

Working with properties	141
Migrating using properties and the tdimigbl tool	142
Global properties	142
Solution properties	142
Java properties	142
System properties	144

Chapter 12. System Store . . . 145

Property stores	145
Password Store	145
User property stores	145
Third-party RDBMS as System Store	146
Oracle	146
MS SQL Server	147
IBM DB2 for z/OS	148
DB2 for other OS	148
Using Derby to hold your System Store	149
Configuring Derby Instances	150
Starting Derby in networked mode	150
Enabling user authentication in System Store	150
Create statements for System Store tables	150
Backing up Derby databases	152
Troubleshooting Derby issues	152
Pre-6.0 (properties file) configuration of Cloudscape	153
See also	155

Chapter 13. Command-line options 157

Configuration Editor	157
Server	157
Command Line Interface – tdisrvctl utility	160
Command Line Reference	160

Chapter 14. Logging and debugging 171

Script-based logging	172
Logging using the default Log4J class	172
Log Levels and Log Level control	176
Log4J default parameters	176
Creating your own log strategies	177

Chapter 15. Tracing and FFDC 179

Tracing Enhancements	179
Understanding Tracing	179
Configuring Tracing	180
Setting trace levels dynamically	180
Useful JLOG parameters.	181

Chapter 16. Administration and Monitoring 183

Installation and Configuration.	183
Deploying AMC into the Integrated Solutions Console (ISC)	183
Starting the Administration and Monitoring Console and Action Manager and logging in	184
Enabling AMC	184
AMC Logs	186
Backward Compatibility with previous versions of Tivoli Directory Integrator	187
AMC in the Integrated Solutions Console	188
Console user authority	188
Action Manager	189
Enabling Action Manager	193
Action Manager status in real time	194
AMC force trigger for a given rule	195
AMC and Action Manager security	195
Introduction	195
AMC and SSL	195
AMC and remote TDI server	196
AMC and role management	197
AMC and passwords	198
AMC and encrypted configs	198
Administration and Monitoring Console User Interface	198
Log in and logout of the console	198
AMC Console Layout	199
Logging off the console	200
Using AMC tables.	200
Servers	202
Console Properties	203
Solution Views	205
Monitor Status and Action Manager.	208
Property Stores.	218
Log Management	219
Preferred Solution Views	220
AMC and AM Command line utilities	220
Example walkthrough of creating a Solution View and Rules	225

Chapter 17. Tombstone Manager 231

Introduction	231
Configuring Tombstones.	231
Configuration Editor Configuration screen.	231
AssemblyLine Configuration screen	233
The Tombstone Manager	234

Chapter 18. Multiple TDI services . . . 237

IBM Tivoli Directory Integrator as Windows Service	237
Introduction	237
Installing and uninstalling the service	237
Starting and stopping the service.	238
Logging	238
Configuring the service	238
IBM Tivoli Directory Integrator as Linux/UNIX Service	240
Deployment methods.	240
Tailoring /etc/inittab.	240
IBM Tivoli Directory Integrator as z/OS Service	241
USS process	241
Normal z/OS started task	242
IBM Tivoli Directory Integrator as i5/OS Service	243

Chapter 19. z/OS environment Support 245

Post install configuration	245
Using MQe for system queue	245
Default encoding different than IBM-1047	245
JDK 5.0 not located at /usr/lpp/java/J5.0.	246
Running Tivoli Directory Integrator	246
Reading License Files.	246
Using the Remote Configuration Editor on z/OS	247
Handling configuration and properties files	247
Using ASCII mode	248
Configuring the TDI task to log to its SYSOUT	249

Appendix A. Dictionary of terms 251

IBM Tivoli Directory Integrator terms	251
---	-----

Appendix B. Example Property files 265

Log4j.properties	265
jlog.properties	266
derby.properties	268
global.properties	268

Appendix C. Notices 275

Third-Party Statements	276
ICU License - ICU 1.8.1 and later.	276
Trademarks	277

Chapter 1. Introduction

For an overview of the general concepts of the IBM Tivoli Directory Integrator 7.0, refer to "IBM Tivoli Directory Integrator concepts," in *IBM Tivoli Directory Integrator V7.0 Users Guide*.

For more detailed information about IBM Tivoli Directory Integrator 7.0 concepts, see *IBM Tivoli Directory Integrator V7.0 Reference Guide*.

IBM Tivoli Directory Integrator 7.0 Editions

The IBM Tivoli Directory Integrator 7.0 exists in two different editions (for which different licensing agreements apply):

Identity Edition

The Identity Edition contains the full set of Connectors, Parsers, Function Components, Password Interceptor Plug-ins and miscellaneous other components. Licensing is done on a per-user basis.

General Purpose Edition

The General Purpose Edition differs from the Identity Edition in that some components have been removed. Those removed are the specific Identity Management components as listed below:

- Windows Users and Groups Connector
- z/OS Changelog Connector
- IDS Changelog Connector
- Netscape/iPlanet/Sun Directory Changelog Connector
- Active Directory Change Detection Connector
- ITIM DSMLv2 Connector
- TAM Connector
- JMS Password Store Connector
- User Registry Connector for SAP R/3
- LDIF Parser
- SPMLv2 Parser
- Password Interceptor Plug-ins

Examples applying to these components have been removed as well.

Licensing for this edition is done on a per-processor basis.

Chapter 2. Installation instructions for IBM Tivoli Directory Integrator

Before you install

The Tivoli Directory Integrator 7.0 installer uses the InstallShield Multiplatform 11.5 wizard. Before you install, read the following sections and make sure your system meets the minimum requirements.

Disk space requirements

The IBM Tivoli Directory Integrator 7.0 installer requires 450 MB of temporary disk space during installation, and additionally the following amount of disk space for the TDI components that remain on the computer after installation:

- Windows: 551 MB
- Linux: 522 MB
- AIX: 499 MB
- Solaris: 642 MB
- HP-UX: 467 MB
- i5/OS: 201 MB

The precise amount of required disk space depends on the components you choose to install; the amounts listed above are applicable for a full Custom option installation. To calculate precisely the necessary disk space, add together the disk space requirements for each component you want to install. See “Components in IBM Tivoli Directory Integrator” for the required disk space for each TDI component.

Memory requirements

The IBM Tivoli Directory Integrator 7.0 installer requires 512 MB of memory. The precise amount of required memory after installation depends on the components you choose to install.

To calculate the necessary memory requirements, add together the memory requirements for each component you want to install. See “Components in IBM Tivoli Directory Integrator” for the memory requirements of each TDI component.

Memory requirements for a Typical installation: 484 MB

Memory requirements for a Custom option installation with all components: 868 MB

Platform requirements

See Chapter 4, “Supported platforms,” on page 37

Components in IBM Tivoli Directory Integrator

With some exceptions, the following components are available and selectable for installation as part of IBM Tivoli Directory Integrator 7.0:

Runtime Server

A rules engine used to deploy and run TDI integration solutions.

- Disk space requirements: 48 MB.
- Memory requirements: Each TDI server instance requires at least 256 MB. NOTE: More RAM may be required depending on the size and complexity of the solution being created.

Configuration Editor

A development environment for creating, debugging and enhancing TDI integration solutions.

Note: IBM Tivoli Directory Integrator does not support the Configuration Editor (CE) on the following operating systems:

- HP-UX Integrity*
- HP PA-RISC
- Solaris Opteron
- z/OS
- i5/OS
- Linux PPC
- Linux 390

* On HP-UX Integrity, you have the option of installing the Tivoli Directory Integrator Eclipse CE plug-ins into an existing Eclipse Workbench; see below for more details. If you choose not to do that, and for the other platforms, see “Using the Remote Configuration Editor” on page 104 and “Using the Remote Configuration Editor on z/OS” on page 247.

Disk space requirements for the Configuration Editor for each supported operating system:

- Microsoft Windows: 139 MB
- Linux: 139 MB
- AIX 139 MB
- Solaris 138 MB

Memory requirements for CE on each supported operating system: 128 MB.

Configuration Editor Update Site (Eclipse update site for CE)

Use the CE Update Site folder to install the Tivoli Directory Integrator Configuration Editor into an existing Eclipse installation. Use the Eclipse software update tool and use this folder as a local update site. The CE update site is only supported for deployment on Eclipse 3.3.2.

Note: IBM Tivoli Directory Integrator does not support the Configuration Editor Update Site on the following operating systems:

- HP PA-RISC
- Solaris Opteron
- z/OS
- i5/OS
- Linux PPC
- Linux 390

See “Using the Remote Configuration Editor” on page 104 and “Using the Remote Configuration Editor on z/OS” on page 247.

CE Update Site requirements are:

- Disk space requirements: 6 MB
- Memory requirements: Not applicable

Java API documentation

Full HTML documentation of TDI internals. Essential reference material for scripting in solutions, as well as for developing custom components.

- Disk space requirements: 48 MB
- Memory requirements: Not applicable

Examples

A series of short, illustrative example Configs that highlight specific Tivoli Directory Integrator features or components.

- Disk space requirements: 3 MB
- Memory requirements: Not applicable

IEHS v3.11 (Host TDI help locally. The default is online.)

You can install an IBM Eclipse Help System locally as an alternative to using the global online help service. This option requires manual download and deployment of the TDI help files after installation.

Disk space requirements by platform:

- Windows: 24 MB
- Linux: 18 MB
- AIX: 18 MB
- Solaris: 18 MB
- HP-UX: 18 MB
- i5/OS: 18 MB

Memory requirements: 128 MB (256 MB or more is recommended.)

Note: You must increase memory according to the size of the documentation plug-ins. For example, if the size of the documentation is 100 MB, add at least 80 MB of additional RAM.

If your platform meets these requirements, you can proceed with the download and installation instructions documented in “Installing local Help files” on page 26.

embedded Web platform (includes Integrated Solutions Console SE)

Tivoli Directory Integrator includes an embedded lightweight Web server platform. This server platform is based on the Eclipse and Open Services Gateway Initiative (OSGI) architecture and supports running web applications and Web services. The runtime provides a secure infrastructure with a small footprint and minimal configuration. The embedded Web platform includes Integrated Solution Console SE, which is used as the default alternative for deploying AMC on an existing ISC installation. The embedded Web platform provides an OSGI based lightweight infrastructure for hosting Web applications and Web services with the following characteristics:

- Minimal footprint
- Minimal configuration
- Compatibility with OSGI based ISC

AMC installation into the embedded Web platform requires at least 94 MB or more on each of the following supported operating systems:

- Windows
- Linux
- AIX
- Solaris
- HP-UX

Memory requirements: a minimum of 512 MB is recommended.

Note: For the i5/OS platform, the Integrated Web Application Server must already be installed on the target computer. See “Installing IBM Tivoli Directory Integrator on i5/OS” on page 12 for details.

AMC: Administration and Monitoring Console

A browser-based application for monitoring and managing running Tivoli Directory Integrator Servers. AMC runs in the Integrated Solutions Console (ISC). In previous releases, AMC was a servlet application that was deployed into an embedded or existing instance of WebSphere Application Server (WAS).

- Disk space requirements: 74 MB
- Memory requirements: 128 MB

Tivoli Directory Integrator 7.0 supports ISC SE 7.1 and ISC AE 7.5 only. There is no support for ifix2 for ISC AE.

Additional components automatically installed that are not selectable:

JRE (Java Runtime Environment) 5.0 SR9

A subset of the Java Development Kit (JDK) that contains the core executable files and other files that constitute the standard Java platform. The JRE includes the Java Virtual Machine (JVM), core classes, and supporting files.

Note: The JRE used for any of the installed Tivoli Directory Integrator packages is independent of any system-wide JRE or JDK you may have installed on your system. The JRE is installed no matter what features are selected. The uninstaller requires the JRE, so it is always installed.

Disk space requirements by platform:

- Windows: 100 MB
- Linux: 66 MB
- AIX: 69 MB
- Solaris: 124 MB
- HP: 220 MB

Memory requirements: Not applicable

Password Synchronization Plug-ins

All supported platforms: 8 MB

Miscellaneous

Contains the License Package, the Uninstaller, the Update Installer and TDI overhead.

The Tivoli Directory Integrator 7.0 License Package contains the license files for Tivoli Directory Integrator.

Disk space requirements by platform:

- Windows: 20 MB
- Linux: 20 MB
- AIX: 20 MB
- Solaris 19 MB
- HP-UX: 20 MB
- i5/OS: 20 MB

Memory requirements: Not applicable.

Other requirements

Root or Administrator Privileges

Note the following differences when installing Tivoli Directory Integrator with administrator as opposed to non-administrator privileges:

- Anyone installing Tivoli Directory Integrator must have write privileges when installing to the specified installation location.
- Non-administrator users have different Configuration Editor shortcuts from administrative users.
- Users who do not have administrator privileges when installing Tivoli Directory Integrator do not see the "Register AMC as a Service" window.
- Once Tivoli Directory Integrator is installed using one particular non-root user ID, that same user ID must be used to carry out any further maintenance on that installation, like un-installation or migration to newer versions.

Security Enhanced (SELinux)

RedHat Linux (RHEL) has a security feature known as Security Enhanced Linux or SELinux. SELinux provides security that protects the host from certain types of malicious attacks. A less secure version of SELinux was included in RHEL version 4.0 and was disabled by default, but RHEL version 5.0 defaults SELinux to enabled. The RHEL 5.0 SELinux default settings have been known to prevent Java 5 from running properly. If you try to run the RHEL 5.0 Tivoli Directory Integrator installer, an error resembling the following output may display:

```
# ./install_tdiv70_linux_x86.bin
```

```
Initializing Wizard.....
```

```
Verifying JVM...
```

```
No Java Runtime Environment (JRE) was found on this system.
```

The reason for this error is that the Java Runtime Environment (JRE) that InstallShield Multi-Platform (ISMP) extracts to the /tmp directory does not have the required permissions to run. To avoid this error:

1. Disable SELinux: `setenforce 0`.
2. Run the Tivoli Directory Integrator installer.
3. Enable SELinux again: `setenforce 1`.

You can also edit the /etc/selinux/config configuration file to enable and disable SELinux. The default settings for the /etc/selinux/config file resemble the following lines:

```
# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
# enforcing - SELinux security policy is enforced.
# permissive - SELinux prints warnings instead of enforcing.
# disabled - SELinux is fully disabled.
SELINUX=enforcing
# SELINUXTYPE= type of policy in use. Possible values are:
# targeted - Only targeted network daemons are protected.
# strict - Full SELinux protection.
SELINUXTYPE=targeted
```

Modifying SELINUX to either SELINUX=permissive or SELINUX=disabled allows the Tivoli Directory Integrator installer to run. However, both modifications of the SELINUX property, to either SELINUX=permissive or to SELINUX=disabled, affect the level of security for the host.

The Tivoli Directory Integrator installer uses a JRE located at *install_dir/jvm* that cannot run with the SELinux default settings. The installer makes a best effort to avoid the problems with the SELinux default settings by trying to change the Tivoli Directory Integrator JRE security permissions that are blocking the installer. The Tivoli Directory Integrator installer issues a command that changes the security permissions for the Tivoli Directory Integrator JRE so that it can run. The Tivoli Directory Integrator installer issues the following command:

```
chcon -R -t textrel_shlib_t install_dir/jvm/jre
```

Note: If the installer cannot issue the `chcon` command, or if there is an error when issuing the command, you must edit the permissions manually.

Errors that resemble the following output indicate that the `chcon` command did not work:

```
[root@dyn9-37-225-164 V7.0]# ./ibmdisrv
Failed to find VM - aborting
```

```
[root@dyn9-37-225-164 V7.0]# ./ibmditk
Failed to find VM - aborting
```

```
[root@dyn9-37-225-164 V7.0]# bin/amc/start_tdiamc.sh
Failed to find VM - aborting
```

Authentication of AMC on Unix/Linux

On some UNIX platforms (for example, SLES 10) the Administration and Monitoring Console (AMC) in ISE SE fails consistently to authenticate users, even when correct credentials are specified. Such behavior is observed when AMC is run as a non-root user and the operating system uses a password database (for example, a `/etc/shadow` file). For more information on this issue, and for a workaround see "Authentication failure on UNIX when LWI runs as non-root user" in *IBM Tivoli Directory Integrator V7.0 Problem Determination Guide*.

Installing IBM Tivoli Directory Integrator

The Tivoli Directory Integrator installer allows you to install Tivoli Directory Integrator 7.0 in its entirety, only those Tivoli Directory Integrator components that you need, upgrade a previous version of Tivoli Directory Integrator (versions 6.0, 6.1 or 6.1.1), or add features to an existing Tivoli Directory Integrator 7.0 installation.

Note: IBM Tivoli Directory Integrator does not support the Configuration Editor (CE) on the following operating systems:

- HP-UX Integrity*
- HP PA-RISC
- Solaris Opteron
- z/OS
- i5/OS
- Linux PPC
- Linux 390

* On HP-UX Integrity, you have the option of installing the CE Plugins into an existing Eclipse Workbench; see "Platform requirements" on page 3. If you choose not to do that, and on the other platforms, see "Using the Remote Configuration Editor" on page 104 and "Using the Remote Configuration Editor on z/OS" on page 247 for more information on using the product without a locally-installed Configuration Editor.

Installing Tivoli Directory Integrator V7.0 uninstalls a previous version; the uninstall does not remove any files that the user has created. User created files are still available after the new installation completes. Configuration files such as `global.properties` and `am_config.properties` are migrated to Tivoli Directory Integrator 7.0, keeping any custom configuration changes that have been made. The Tivoli Directory Integrator 7.0 installation continues to use the features available in previous versions of Tivoli Directory Integrator:

- Administration and Monitoring Console (AMC)
- Configuration Editor (CE)
- Examples
- IBM Eclipse Help System (IEHS)
- Java API Documentation

- Runtime Server

Note: For the remainder of this *IBM Tivoli Directory Integrator V7.0 Installation and Administrator Guide*, the variable *install_dir* refers to the installation directory location chosen by the user on the **Destination Panel** during installation. See “Default installation locations” on page 31 for information on where Tivoli Directory Integrator is usually installed.

Launching the appropriate installer

You can launch the IBM Tivoli Directory Integrator 7.0 Installer by using one of the following methods:

Launch the installer from the Launchpad

The Tivoli Directory Integrator Launchpad provides essential getting started installation information and links to more detailed information on various installation, migration, and post installation topics. In addition, Launchpad allows you to launch the Tivoli Directory Integrator installer.

Notes:

1. The Launchpad is not available on z/OS and i5/OS.
2. Using the Launchpad requires that you have a supported Web browser installed and configured; if this is not the case, you cannot use the Launchpad. However, you can still use the platform-specific installer directly; see “Using the platform-specific TDI installer” on page 12 for instructions on how to use the Tivoli Directory Integrator Installer.

Note:

1. Open the Tivoli Directory Integrator Launchpad by typing the following command at the command prompt:
 - For Windows platforms, type:
`Launchpad.bat`
 - For all other platforms, type:
`Launchpad.sh`

The menu on the left of the Launchpad allows you to navigate the Launchpad windows. Click a menu item to view information about it. The following menu items are available:

Welcome

The installation Welcome window contains links to:

- IBM Tivoli Directory Integrator Web site
- 7.0 Documentation
- Support Web site
- Tivoli Directory Integrator news group



The following options on the left are TDI Launchpad windows:

Release Information

This window contains a list of some of the new and improved features available this release, as well as links to documentation about the release.

Prerequisite Information

This window contains links to information about platform support and hardware requirements.

Installation scenarios

This window contains a description of the TDI components available for installation. You can install some or all of these components during installation. This window also contains a description of the Password Synchronization Plug-ins components available for installation.

Migration Information

This window contains a link to information about migrating from Tivoli Directory Integrator 6.0 or 6.1.X to 7.0. It also contains information about migrating the Derby System Store.

Install IBM Tivoli Directory Integrator

This window contains links to the IBM Tivoli Directory Integrator Installer, as well as links to installation, migration and supported platforms documentation. See "Using the platform-specific TDI installer" on page 12 for instructions on how to use the IBM Tivoli Directory Integrator Installer.

Install IBM Tivoli Directory Integrator Password Synchronization Plug-ins

This window contains links to the IBM Tivoli Directory Integrator Password Synchronizer Plug-ins Installer, as well as links to installation and supported platforms documentation.

Note: This window is not available on Linux PPC and Linux 390 platforms.

Exit Exits the Launchpad, without installing anything.

2. On the installation window, click IBM Tivoli Directory Integrator **Installer**. This launches the installer. See “Using the platform-specific TDI installer” on page 12 for instructions on how to use the installer.

Launch the installer directly

You can launch the installer directly using the installation executable file:

1. Locate the installation executable file for your platform in the `tdi_installer` directory on the product CD (on i5/OS this directory is called `TDI_INST`).

Windows Intel

`install_tdiv70_win_x86.exe`

Windows AMD64/EM64T

`install_tdiv70_win_x86_64.exe`

AIX `install_tdiv70_aix.bin`

Linux `install_tdiv70_linux_x86.bin`

Linux AMD 64

`install_tdiv70_linux_x86_64.bin`

Power PC Linux

`install_tdiv70_ppclinux.bin`

z/OS Linux

`install_tdiv70_zlinux.bin`

Solaris Sparc

`install_tdiv70_solaris_sparc.bin`

Solaris

`install_tdiv70_solaris_x86_64.bin`

HP-UX PA-RISC

`install_tdiv70_hpux_parisc.bin`

HP-UX Integrity

`install_tdiv70_hpux_ia64.bin`

i5/OS INST_TDI.SH

2. Double-click the executable file, or type the executable file name at the command prompt. This launches the installer. See “Using the platform-specific TDI installer” on page 12 for information on how to use the installer.

Once you have launched the installer (using the Launchpad or by starting the platform-dependent installer directly), you are ready to begin the process of “Using the platform-specific TDI installer” on page 12.

Notes:

1. By default, the Windows installers will return immediately when started from a command prompt. This makes it impossible to write a script that will run the installer on Windows and wait for it to complete. Additional executables, `install_tdiv70_win_x86_console.exe` and `install_tdiv70_win_x86_64_console.exe` are provided; these do not return until completed, so these executables should be used if such behavior is desired. There is no other difference in their behavior.
2. Non-administrators can install Tivoli Directory Integrator, with the following caveats: users installing Tivoli Directory Integrator must have write privileges to the installation location; non-administrators do not see the "Register AMC as a service" window, and non-administrator Configuration Editor shortcuts differ from administrator Configuration Editor shortcuts.

Using the platform-specific TDI installer

The platform-specific Tivoli Directory Integrator installer is launched either from the Launchpad or from the command line. The Tivoli Directory Integrator installer can be used to install a new copy of Tivoli Directory Integrator, add a feature to an existing instance of Tivoli Directory Integrator, or upgrade a previous version of Tivoli Directory Integrator. The default install location on your computer for Tivoli Directory Integrator varies with the platform.

Before you install

Note: In addition to being unavailable on the i5/OS operating system, Tivoli Directory Integrator supports neither the Configuration Editor (CE) nor the Configuration Editor Update Site on the following operating systems:

- HP-UX Integrity
- HP PA-RISC
- Solaris Opteron
- z/OS
- Linux PPC
- Linux 390

See “Using the Remote Configuration Editor” on page 104 for information on how to develop solutions without a local Configuration Editor.

Installing IBM Tivoli Directory Integrator on i5/OS

InstallShield Multi Platform (ISMP) 11.5 Tivoli Directory Integrator supports installation on i5/OS. The i5/OS platform does not support a GUI interface (Launchpad). i5/OS supports the command line installation -console option. See “Installing using the command line” on page 23.

The following Tivoli Directory Integrator features are not available on i5/OS, and are not listed as installable features during -console installation:

- Configuration Editor (CE/Integrated Development Environment (IDE)) – This component is the Tivoli Directory Integrator IDE. Also see “Using the Remote Configuration Editor on z/OS” on page 247.
- Configuration Editor Update Site (Eclipse Update site for CE) – An Eclipse update site used for CE maintenance and for allowing the customer to install the CE to an existing Eclipse workbench if they do not want to use the stand alone Rich Client Platform (RCP) application.

Prior to installing Tivoli Directory Integrator, certain software must already be installed on the i5/OS operating system running Tivoli Directory Integrator.

Java virtual machine version: Tivoli Directory Integrator requires the IBM J9 32-bit JVM on all versions of i5/OS. If the J9 32-bit JVM is not found, an error message displays:

The install was unable to detect the IBM J9 VM (32-bit). The IBM J9 VM is required for this product. Please install this JVM then try again.

If you see this message, you must cancel the installation, install the IBM J9 VM, and begin the installation again.

If you choose the embedded Web platform feature, the install will check to make sure that LWI 7.0.1(v5r4) or 7.1.1(v6r1) is resident on the target machine.

PTF versions: For i5/OS V5R4 the installer will check that the following items are installed:

1. Product 5722JV1 option 8 (J2SE 5.0 32 bit)
2. PTF group SF99291 level 17 or higher (Java 5.0 SR9)
3. PTF group SF99114 level 13 or higher (embedded Web platform 7.0.1)
4. PTF SI34226 for product 5722SS1 (SSL support for embedded Web platform)

5. PTF SI34483 for product 5722SS1 (embedded Web platform fixes)

For i5/OS V6R1 the installer will check that the following items are installed:

1. Product 5761JV1 option 8 (J2SE 5.0 32 bit)
2. PTF group SF99562 level 5 or higher (Java 5.0 SR9)
3. Product 5761DG1, *BASE (IBM HTTP Server for i5/OS – Contains embedded Web platform 7.1.1)
4. PTF SI33788 for product 5761SS1 (Administration and Monitoring Console Roles fix)
5. PTF SI34021 for product 5761SS1 (SSL fixes)

If these PTFs or products are not found, an error message appears:

The installer was unable to detect the i5/OS product/PTFs required by the embedded Web platform feature. You may choose to continue the install without the embedded Web platform feature or you may exit now and refer to the install log for a list of the missing requirements.

Installation:

Note: The installer and uninstaller on i5/OS are called INST_TDI.SH and uninstaller.sh, respectively.

To begin installing on i5/OS:

1. Locate the installation executable file for i5/OS in the TDI_INST directory on the product CD. Launchpad is not available on i5/OS. The default location to install i5/OS on your computer is: /QIBM/ProdData/TDI/V7.0
2. On i5/OS, in order to extract the Tivoli Directory Integrator 7.0 installer from a TAR image, you must set environment variable "QIBM_CCSD" to 819; that is, run the command
`export QIBM_CCSD=819`

before un-tarring the Tivoli Directory Integrator installer TAR image.

Another i5/OS difference is on the Tivoli Directory Integrator Solutions directory panel. On i5/OS, there is a specific place for user data. As a result, instead of giving you the option to make the installation directory the same as the solutions directory, the option reads: Use the TDI User Product Directory.

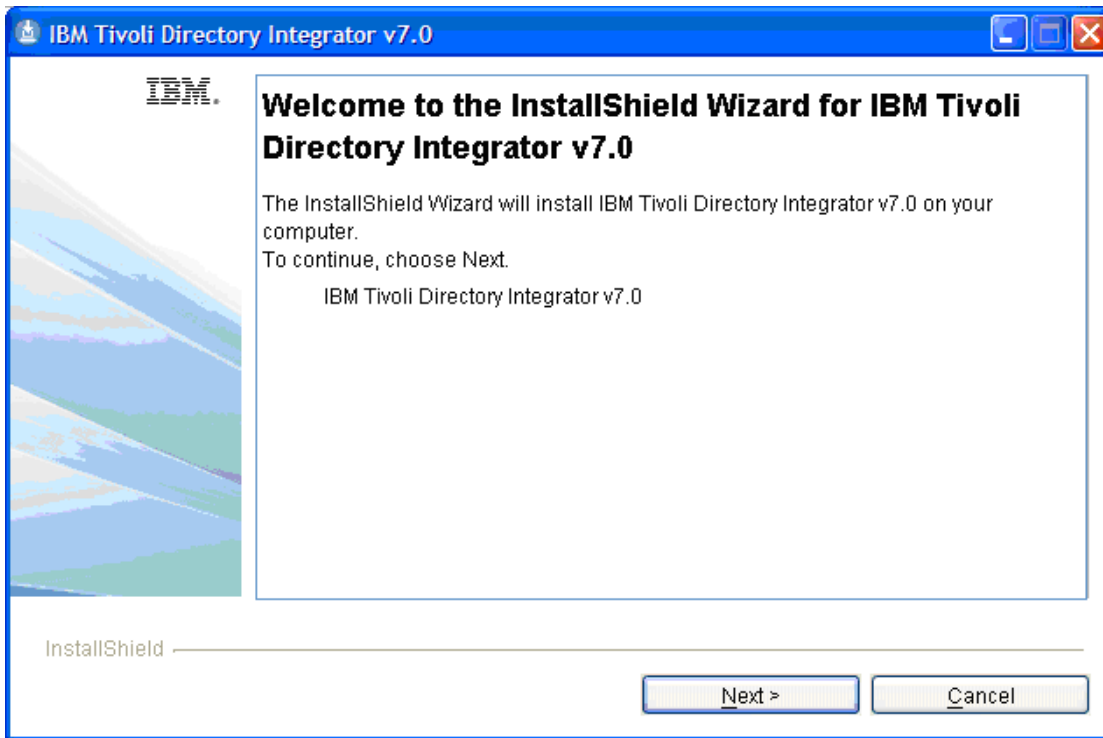
Installing a new version of Tivoli Directory Integrator 7.0

Tivoli Directory Integrator is either not present on the target system, or you chose to install a new version when prompted.

When you run the installer:

A window appears, stating "The InstallShield Wizard is preparing the installation, please wait."

After the installation program is loaded, the Welcome window appears, stating "Welcome to the InstallShield Wizard for IBM Tivoli Directory Integrator v7.0."



1. Click **Next** to continue.

A window appears, stating "Click next to search for previously installed versions of Tivoli Directory Integrator. This may take a while..."

2. Click **Next** to search for previously installed versions of Tivoli Directory Integrator.
3. If there is a version of Tivoli Directory Integrator already installed on your computer, a new window lists several options:
 - Install Tivoli Directory Integrator version 7.0 to a new location.
 - Add features to a current version of Tivoli Directory Integrator.
 - Upgrade one of the following Tivoli Directory Integrator installed versions to version 7.0.
4. Select **Install TDI 7.0 to a new location**. Click **Next**.
5. After reading the software license agreement, select one of the following options:
 - **I accept both the IBM and the non-IBM terms.**
 - **I do not accept the terms in the license agreement.**

If you do not accept the terms of the software license agreement, you cannot continue with the installation. Click **Next** to continue.

A new window appears, showing the default directory for the installation of Tivoli Directory Integrator 7.0.

6. To accept the default directory, click **Next**. To have InstallShield install Tivoli Directory Integrator to a different directory, click **Browse**, navigate to the directory you want, and click **Open** in Windows Explorer. Click **Next**.

Notes:

- a. Non-ASCII characters and the following list of characters are not supported in the install path:
";|*?!#&\$',=^@%
- b. You cannot choose an installation location of an existing installation (an "instance") of Tivoli Directory Integrator. If you try to select the directory of an existing Tivoli Directory Integrator version, an error appears, stating:

The directory you have chosen already contains an installation of IBM Tivoli Directory Integrator. Choose another directory.

Note:

7. Select the installation type you prefer:

- Select **Typical** to install the Java Runtime Environment, the Runtime Server, Configuration Editor, Java API documentation, Examples, embedded Web platform (includes Integrated Solutions Console SE), and Administration and Monitoring Console (AMC). This option does not give you the opportunity to install individual features.
- Select **Custom** to choose individual or additional features to install.

Click **Next**.

8. If you are performing a custom installation, select the Tivoli Directory Integrator components you want to install. See "Components in IBM Tivoli Directory Integrator" on page 3 for a list and descriptions of components available for installation. When you have finished making your selections, click **Next**.
9. If you are performing a typical installation, or if you selected to install the Runtime Server, you must select a solution directory. The solution directory is where the Tivoli Directory Integrator Server and Configuration Editor locate your solutions, like Config files and properties files. To set up your solution directory, select one of the following options:
 - a. **Use a subdirectory named TDI under my home directory.** This creates a subdirectory under your home directory, for example for Windows: C:\Documents and Settings\username\My Documents\TDI.
 - b. **Use Install Directory.** This is the directory you specified in "Installing a new version of Tivoli Directory Integrator 7.0" on page 13, step6 on page 14.
 - c. **Select a directory to use.** Navigate to or type the name of a directory. Any directory reachable on the local computer can be established as the default solution directory.
 - d. **Do not specify.** Use the current working directory at startup time.

Click **Next**.

Note: On i5/OS, the "Use Install Directory" is replaced with "Use the TDI User Product Directory". The i5/OS platform has a specific location (/QIBM/UserData) for user data. The installation directory is not appropriate.

10. If you selected to deploy Administration and Monitoring Console (AMC) to ISC, a window appears, providing you with options for deploying AMC into the Integrated Solutions Console (ISC):
 - **Embedded instance of Integrated Solutions Console SE.**
 - **Existing instance of Integrated Solutions Console.** When selecting this option in Windows, you can use **Browse** to locate the instance of ISC.
 - **Do not specify. I will manually deploy AMC at a later time.**

Note: If you did not select the embedded Web platform component, the option **Embedded instance of Integrated Solutions Console SE** is not available.

11. If you chose to deploy AMC to ISC, enter port-related values to be used by ISC and AMC for the following fields:
 - **HTTP port.** The default HTTP port is 13100.
 - **HTTPS port.** The default HTTPS port is 13101.
 - **Action Manager API port.** The default Action Manager API port is 13104.

Enter ISC port information and click **Next**. If you have selected to deploy AMC to the embedded Web platform, a window appears giving you the option to register AMC as a system service.

12. If you want to register AMC as a service for your operating system, select **Register AMC as a system service**. The default is cleared, or do not register AMC as a service with the OS.

The installer tries to provide a valid default value for **Service Name**. If the installer is unable to determine a valid **Service Name**, the installer leaves the field blank. You cannot continue with the installation until you enter a valid service name.

Note: The register AMC as a service window appears only if these conditions are met:

- The embedded Web platform and AMC features are selected.
- The userid has Administrator privileges.

A window appears, stating "Please read the summary information below."

13. Review the installation information you have selected. The summary shows what is going to be installed. Click **Back** to make any changes. When you are ready to begin installation, click **Install**.

A window showing installation progress appears. The progress bar applies to all features selected for installation. Each file or group of files being installed, along with its path, displays rapidly to let you know what is being installed.

14. Click **Cancel** only if you want to cancel the installation.

A summary information window appears when the installation process is complete.

15. If you chose to install the Configuration Editor and if you want the Configuration Editor to start immediately, select **Start the Configuration Editor**. Click **Finish** to complete the installation process.

Adding features to an existing (installed) version of Tivoli Directory Integrator

Tivoli Directory Integrator 7.0 is present on the target system, and some features supported for your platform have not yet been installed.

Notes:

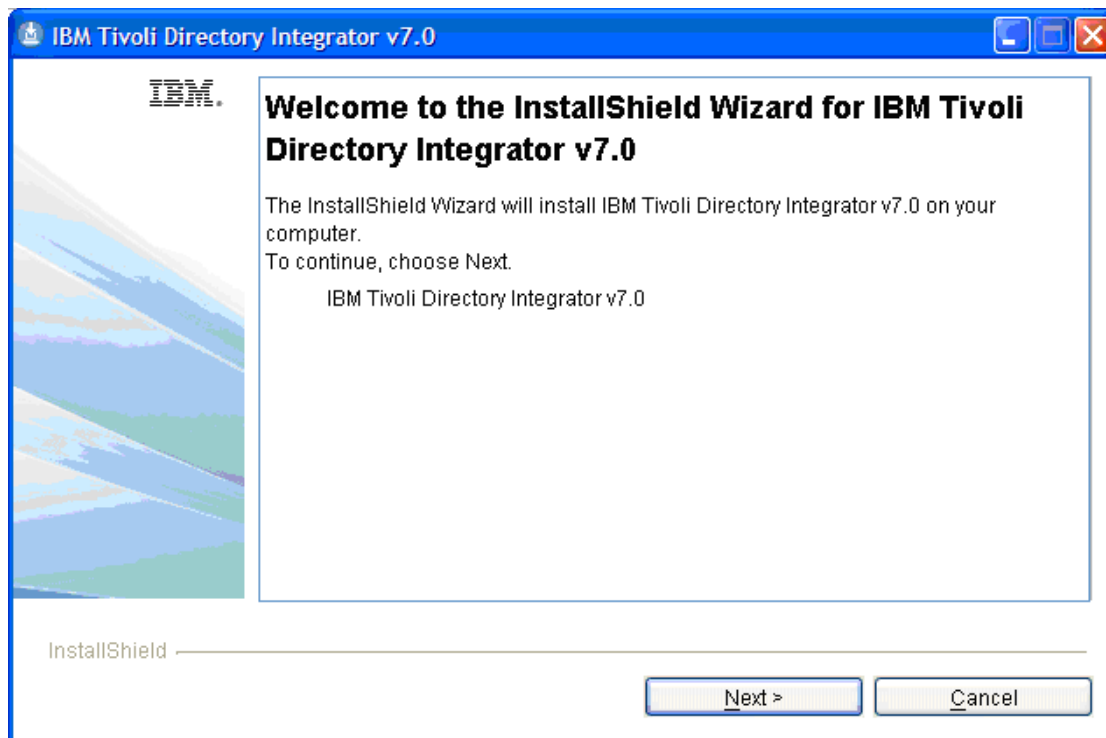
1. You cannot choose the **Add Features** option if there is not an installed Tivoli Directory Integrator 7.0 instance available. The Tivoli Directory Integrator 7.0 selections are enabled if the **Add Features** option is chosen.
2. You cannot remove features using this procedure. To remove features, you must use the Tivoli Directory Integrator uninstaller. See "Uninstalling" on page 29.

When you run the installer:

A window appears, stating "The InstallShield Wizard is preparing the installation, please wait."

After the installation program is loaded, the Welcome window appears, stating "Welcome to the InstallShield Wizard for IBM Tivoli Directory Integrator v7.0."

1. Click **Next** to continue.



A window appears, stating "Click next to search for previously installed versions of Tivoli Directory Integrator. This may take a while..."

2. Click **Next** to search for previously installed versions of Tivoli Directory Integrator.
3. If there is a version of Tivoli Directory Integrator already installed on your computer, a new window lists several options:
 - Install Tivoli Directory Integrator version 7.0 to a new location.
 - Add features to a current version of Tivoli Directory Integrator.
 - Upgrade one of the following Tivoli Directory Integrator installed versions to version 7.0.
4. Select **Add features to a current version of Tivoli Directory Integrator**.
5. Immediately below the **Add features** option, click the down arrow and select the version of Tivoli Directory Integrator to which you want to add features. Click **Next**.
 You are not prompted for an installation location. The feature selection window appears. The features that are installed are shown already selected. You cannot clear these features.
6. Choose the features you would like to add to the installed version of TDI that you have selected. Click **Next**.
7. If you are adding the Runtime Server, you must select a solutions directory. The solutions directory is where the Tivoli Directory Integrator server and Configuration Editor locate your solutions, like Config files and properties files. To set up your solution directory, select one of the following options:
 - a. **Use a subdirectory named TDI under my home directory.** This creates a subdirectory under your home directory, for example for Windows: C:\Documents and Settings\username\My Documents\TDI.
 - b. **Use the install directory.** This is the directory you specified in "Installing a new version of Tivoli Directory Integrator 7.0" on page 13, step 6 on page 14.
 - c. **Select a directory to use.** Navigate to or type the name of a directory. Any directory reachable on the local computer can be established as the default solution directory.
 - d. **Do not specify.** Use the current working directory at startup time.

Click **Next**.

Note: On i5/OS, the "Use Install Directory" is replaced with "Use the TDI User Product Directory." The i5/OS platform has a specific location (/QIBM/UserData) for user data. The installation directory is not appropriate.

8. If you chose to install the Administration and Monitoring Console (AMC) feature and deploy to ISC, a window appears, providing you with options for deploying AMC into the Integrated Solutions Console (ISC):
 - **Embedded instance of Integrated Solutions Console SE.**
 - **Existing instance of Integrated Solutions Console.** When selecting this option in Windows, you can use **Browse** to locate the instance of ISC.
 - **Do not specify. I will manually deploy AMC at a later time.**

Note: If you did not select the embedded Web platform component, the option **Integrated Solutions Console SE** is not available.

9. If you chose to deploy AMC to ISC, enter port-related values to be used by ISC and AMC for the following fields:
 - HTTP port. The default HTTP port is 13100.
 - HTTPS port. The default HTTPS port is 13101.
 - Action Manager API port. The default Action Manager API port is 13104.

Enter ISC port information and click **Next**. If you have selected to deploy AMC to the embedded Web platform, a window appears giving you the option to register AMC as a system service.

10. On platforms other than i5/OS, if you want to register AMC as a service for your operating system, select **Register AMC as a system service**. The default is cleared, or do not register AMC as a service with the OS.

The installer tries to provide a valid default value for **Service Name**. If the installer is unable to determine a valid **Service Name**, the installer leaves the field blank. You cannot continue with the installation until you enter a valid service name.

Note: The register AMC as a service window appears only if these conditions are met:

- The bundled ISC platform and AMC features are selected.
- The userid has Administrator privileges.

11. Review the installation information you have selected. The summary shows what is going to be installed. Click **Back** to make any changes. When you are ready to begin the installation, click **Install**.

A window showing installation progress appears. The progress bar applies to all features selected for installation. Each file or group of files being installed, along with its path, displays rapidly to let you know what is being installed.

12. Click **Cancel** only if you want to cancel the installation.

A summary information window appears when the installation process is complete.

13. If you chose to install the Configuration Editor and if you want the Configuration Editor to start immediately, select **Start the Configuration Editor**. Click **Finish** to complete the installation process.

Upgrading an installed version of Tivoli Directory Integrator

Note: You cannot select the **Upgrade** option if there are no previous versions of Tivoli Directory Integrator available.

Upgrading a version of Tivoli Directory Integrator from Tivoli Directory Integrator 6.0 to Tivoli Directory Integrator 7.0

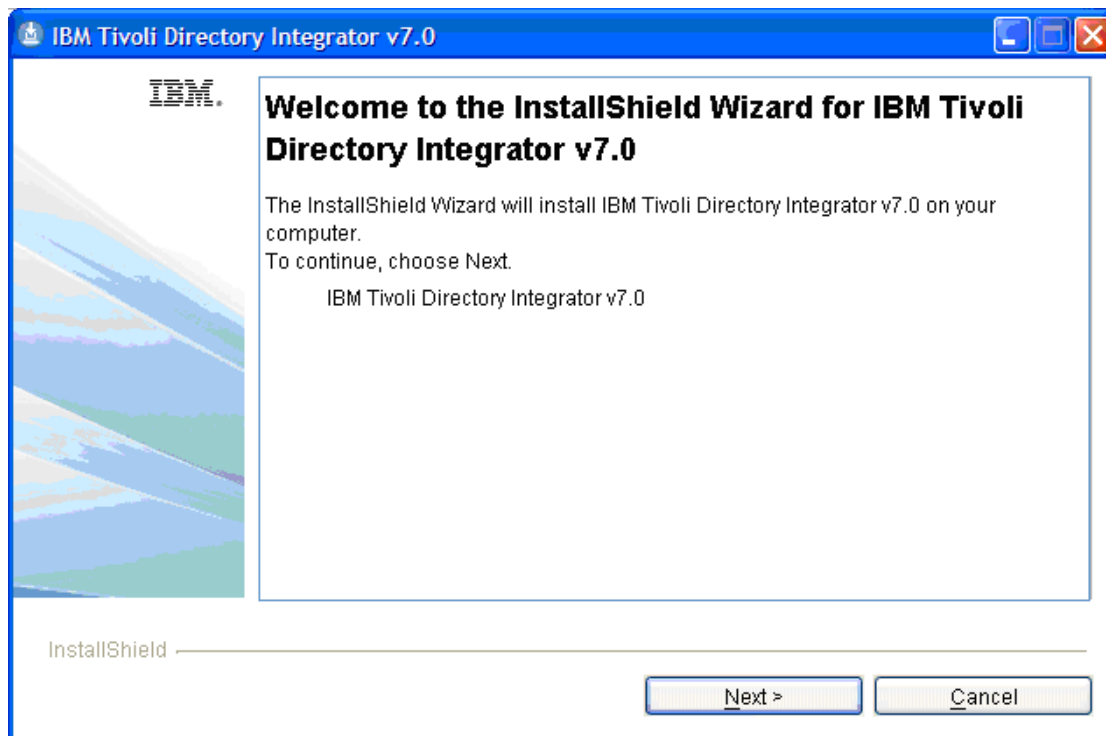
Notes:

1. If you have installed IBM Tivoli Directory Integrator 6.0 as a Windows Service, see section "Uninstalling the service" on page 238 first.
2. Not all files are upgraded when performing the installer upgrade. See section "Migrate files to a newer version" on page 43 for information what files are automatically upgraded and which need to be updated manually.

After some time, a window appears, stating "The InstallShield Wizard is preparing the installation, please wait."

After the installation program is loaded, the Welcome window appears, stating "Welcome to the InstallShield Wizard for IBM Tivoli Directory Integrator v7.0."

1. Click **Next** to continue.



A window appears, stating "Click next to search for previously installed versions of Tivoli Directory Integrator. This may take a while..."

2. Click **Next** to search for previously installed versions of Tivoli Directory Integrator.
If there is a version of TDI already installed on your computer, a new window lists several options:
 - Install Tivoli Directory Integrator version 7.0 to a new location.
 - Add features to a current version of Tivoli Directory Integrator.
 - Upgrade one of the following Tivoli Directory Integrator installed versions to version 7.0.
3. Select **Upgrade one of the following Tivoli Directory Integrator installed versions to version 7.0.**
4. Immediately below the **Upgrade** option, click the down arrow and select a TDI 6.0 version. Click **Next**.

Note: There is no version detection and automated upgrade for versions prior to Tivoli Directory Integrator 6.0.

5. Select **Typical** or **Custom** and click **Next**. If you select custom, the feature selection window appears. The Default Install Location window is skipped, and the value for the installation location is taken from the TDI 6.0 installation.
6. If you are performing a custom installation, select the Tivoli Directory Integrator components you want to install. When you have finished making your selections, click **Next**.

Note: See “Components in IBM Tivoli Directory Integrator” on page 3 for a list and descriptions of components available for installation.

7. If you are upgrading the Runtime Server feature, the Tivoli Directory Integrator Solutions Directory window is skipped, and the value for the Solutions Directory location is taken from the TDI 6.0 installation.
8. If you chose to install the AMC feature and deploy to ISC, the ISC window appears, providing you with options for deploying AMC into the Integrated Solutions Console (ISC):
 - **Embedded instance of Integrated Solutions Console SE.**
 - **Existing instance of Integrated Solutions Console.** When selecting this option in Windows, you can use **Browse** to locate the instance of ISC.
 - **Do not specify. I will manually deploy AMC at a later time.**

Note: If you did not select the embedded Web platform component, the option **Embedded instance of Integrated Solutions Console SE** is not available.

9. If you chose to deploy AMC to ISC, enter port-related values to be used by ISC and AMC for the following fields:
 - HTTP port. The default HTTP port is 13100.
 - HTTPS port. The default HTTPS port is 13101.
 - Action Manager API port. The default Action Manager API port is 13104.

Enter ISC port information and click **Next**. If you have selected to deploy AMC to the embedded Web platform, a window appears giving you the option to register AMC as a system service.

10. On platforms other than i5/OS, if you want to register AMC as a service for your operating system, select **Register AMC as a system service**. The default is cleared, or do not register AMC as a service with the OS.

The installer tries to provide a valid default value for **Service Name**. If the installer is unable to determine a valid **Service Name**, the installer leaves the field blank. You cannot continue with the installation until you enter a valid service name.

Note: The register AMC as a service window appears only if these conditions are met:

- The bundled ISC platform and AMC features are selected.
- The userid has Administrator privileges.

A window appears, stating "Please read the summary information below."

11. Review the installation information you have selected. The summary shows what is going to be installed. Click **Back** to make any changes. When you are ready to begin installation, click **Install**.

A window showing installation progress appears. The progress bar applies to all features selected for installation. Each file or group of files being installed, along with its path, displays rapidly to let you know what is being installed.
12. Click **Cancel** only if you want to cancel the installation.

A summary information window appears when the installation process is complete.
13. If you chose to install the Configuration Editor and if you want the Configuration Editor to start immediately, select **Start the Configuration Editor**. Click **Finish** to complete the installation process.

Updating a version of Tivoli Directory Integrator 6.1 or 6.1.1 to Tivoli Directory Integrator 7.0

A version of Tivoli Directory Integrator 6.1 or 6.1.1 is present on the system, and the Tivoli Directory Integrator 7.0 installer is invoked to upgrade the existing installation.

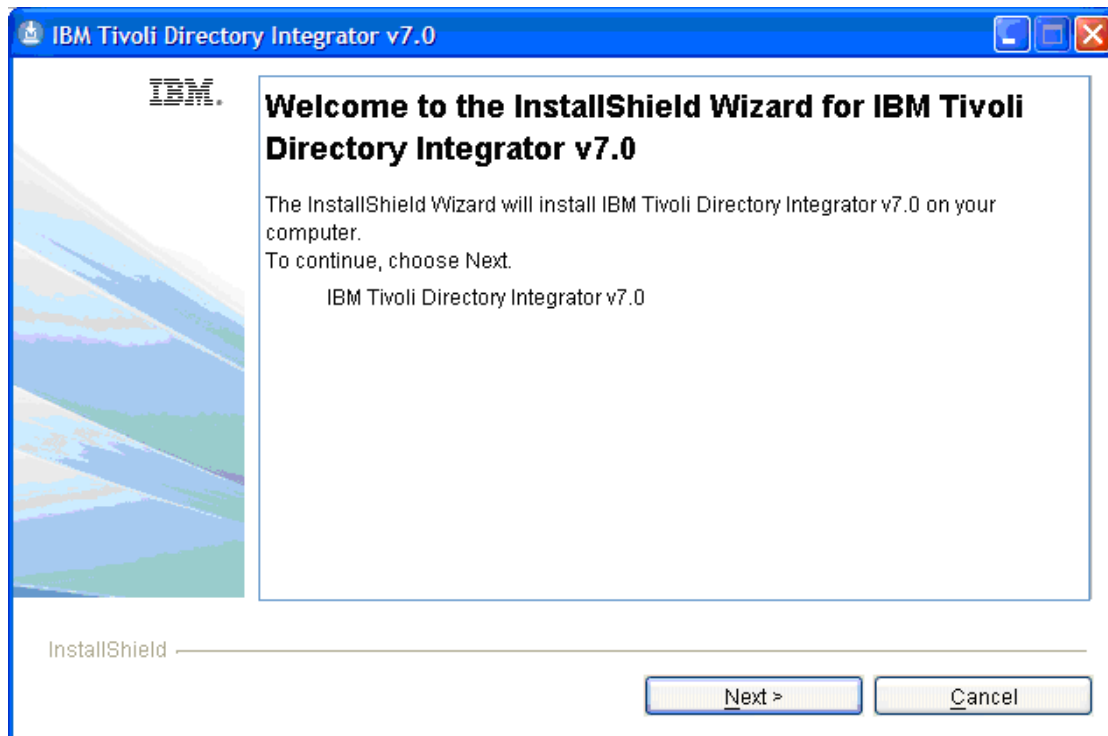
Notes:

1. If you have installed IBM Tivoli Directory Integrator 6.1 or 6.1.1 as a Windows Service, see section "Uninstalling the service" on page 238 first.
2. Not all files are upgraded when performing the installer upgrade. See section "Migrate files to a newer version" on page 43 for information what files are automatically upgraded and which need to be updated manually.

After some time, a window appears, stating "The InstallShield Wizard is preparing the installation, please wait."

After the installation program is loaded, the Welcome window appears, stating "Welcome to the InstallShield Wizard for IBM Tivoli Directory Integrator v7.0."

1. Click **Next** to continue.



A window appears, stating "Click next to search for previously installed versions of Tivoli Directory Integrator. This may take a while..."

2. Click **Next** to search for previously installed versions of Tivoli Directory Integrator.
If there is a version of Tivoli Directory Integrator already installed on your computer, a new window lists several options:
 - Install Tivoli Directory Integrator 7.0 to a new location.
 - Add features to a current version of **Tivoli Directory Integrator**.
 - Upgrade one of the following Tivoli Directory Integrator installed versions to version 7.0.
3. Select **Upgrade one of the following Tivoli Directory Integrator installed versions to version 7.0**.

4. Immediately below the Upgrade option, click the down arrow and select a Tivoli Directory Integrator 6.1 or 6.1.1 version. Click **Next**.
5. You are not prompted for an installation location. The feature selection window appears. The features already installed are shown already selected. You cannot clear these features. Choose the features you would like to add to the installed version of Tivoli Directory Integrator that you have selected. Click **Next**.
6. If you are adding the Runtime Server, you must select a solutions directory. The solutions directory is where the Tivoli Directory Integrator server and CE locate your solutions, like Config files and properties files. To set up your solution directory, select one of the following options:
 - **Use a subdirectory named TDI under my home directory.** This creates a subdirectory under your home directory, for example for Windows: C:\Documents and Settings\username\My Documents\TDI.
 - **Use the installation directory.** This is the directory you specified in 6 on page 14.
 - **Select a directory to use.** Navigate to or type the name of a directory. Any directory reachable on the local computer can be established as the default solution directory.
 - **Do not specify.** Use the current working directory at startup time.

Click **Next**.

Note: On i5/OS, the phrase "Use installation directory"; is replaced with "Use the TDI User Product Directory."; The i5/OS platform has a specific location (/QIBM/UserData) for user data. The installation directory is not appropriate.

7. If you chose to install the AMC feature and to deploy to ISC, the ISC window appears, providing you with options for deploying AMC into the Integrated Solutions Console (ISC):
 - **Embedded instance of Integrated Solutions Console SE.**
 - **Existing instance of Integrated Solutions Console.** When selecting this option in Windows, you can use **Browse** to locate the instance of ISC.
 - **Do not specify. I will manually deploy AMC at a later time.**

Note: If you did not select the embedded Web platform component, the option **Embedded instance of Integrated Solutions Console SE** is not available.

8. If you chose to deploy AMC to ISC, enter port-related values to be used by ISC and AMC for the following fields:
 - HTTP port. The default HTTP port is 13100.
 - HTTPS port. The default HTTPS port is 13101.
 - Action Manager API port. The default Action Manager API port is 13104.

Enter ISC port information and click **Next**. If you have selected to deploy AMC to the embedded Web platform, a window appears giving you the option to register AMC as a system service.

9. On platforms other than i5/OS, if you want to register AMC as a service for your operating system, select **Register AMC as a system service**. The default is cleared, or do not register AMC as a service with the OS.

The installer tries to provide a valid default value for **Service Name**. If the installer is unable to determine a valid **Service Name**, the installer leaves the field blank. You cannot continue with the installation until you enter a valid service name.

Note: The register AMC as a service window appears only if these conditions are met:

- The bundled ISC platform and AMC features are selected.
- The userid has Administrator privileges.

A window appears, stating "Please read the summary information below."

10. Review the installation information you have selected. The summary shows what is going to be installed. Click **Back** to make any changes. When you are ready to begin installation, click **Install**.

A window showing installation progress appears. The progress bar applies to all features selected for installation. Each file or group of files being installed, along with its path, displays rapidly to let you know what is being installed.

11. Click **Cancel** only if you want to cancel the installation.

A summary information window appears when the installation process is complete.

12. If you chose to install the Configuration Editor and if you want the Configuration Editor to start immediately, select **Start the Configuration Editor**. Click **Finish** to complete the installation process.

Installing using the command line

The following command line options are supported by the Tivoli Directory Integrator 7.0 installer:

-console

Specifies to use the console interface mode, where messages during installation are displayed on the Java console and the wizard is run in console mode.

```
install_tdiv70_win_x86.exe -console
```

-options-record

Specifies that the TDI installation Wizard should automatically generate a response file for the project after the completion of the installation or uninstallation.

```
install_tdiv70_win_x86.exe -options-record Response File Name
```

-options

Specifies that a response file be used to run the installation or uninstallation of TDI. A response file is usually used when a silent installation is run (see the next option).

```
install_tdiv70_win_x86.exe -options Response File Name
```

Note: The directory *TDI_install_dir/examples/install* contains a number of example response files for various installation and de-installation scenarios.

- silent** Specifies to install or uninstall the product in silent mode, where the installation or uninstallation is performed with no user interaction. The **-options** command line option is used here to specify what response file to use.

```
install_tdiv70_win_x86.exe -silent -options Response File Name
```

```
TDI_install_dir/_uninst/uninstaller.exe -silent -options Uninstall Response File Name
```

-is:javahome *java home directory*

Specifies the home directory location of the JVM to the launcher. The JVM that is specified must be at the 1.4.2 level or later. Currently, version 1.5 SR9 is included.

```
install_tdiv70_win_x86.exe -is:javahome c:\Java142
```

Note: This option should only be used if the installer will not run with the bundled JRE.

-is:log *filename*

This option is useful for setting up launchers that hide the Java console because it logs detailed information about the launcher, including the actual Java commands that were used to start the Java program, to the specified filename. This includes all of the "std out" and "std err" messages from the Java process.

```
install_tdiv70_win_x86.exe -is:log c:\temp\Log.txt
```

-is:tempdir *directory*

Sets the path to the temporary directory to which the launcher should write its temporary files. If the specified directory does not exist or is not a directory, the launcher uses the system temp directory instead, and no error message is provided.

```
install_tdiv70_win_x86.exe -is:tempdir c:\privateTemp
```

The following command line option is unique to the Tivoli Directory Integrator installation Wizard:

-V TDI_BACKUP

This parameter should only be passed in on an uninstallation. This parameter is provided for future migration considerations.

```
TDI_install_dir\_uninst/uninstaller.exe -V TDI_BACKUP="true"
```

-V TDI_NOSHORTCUTS

This parameter should only be passed in on an installation. This parameter is provided so that installers may avoid creating shortcuts if they are not needed.

```
install_tdiv70_win_x86.exe -V TDI_NOSHORTCUTS="true"
```

Performing a silent install

To perform a silent installation you must first generate a response file. To generate this file, perform a non-silent installation with the `-options-record` option specified, for example:

```
install_tdiv70_win_x86.exe -options-record Response File Name
```

The response file is created in the directory that you specify during installation.

Note: The directory `TDI_install_dir/examples/install` contains a number of example response files for various installation and uninstallation scenarios.

Once the response file is created, you can install silently using the following command:

```
install_tdiv70_win_x86.exe -silent -options Response File Name
```

Note: The examples in this document use the Windows platform installation executable file. See “Launching the appropriate installer” on page 9 for a list of executable file names for each supported platform.

Post-installation steps

CE Update Site

If CE Update site was installed, you now have to manually deploy into Eclipse. See the section entitled “Installing or Updating using the Eclipse Update Manager” on page 28 for more information.

Plugins

If any of the password synchronization plugins were installed, see the *IBM Tivoli Directory Integrator V7.0 Password Synchronization Plug-ins Guide* for information on how to deploy the plugin code.

Administration and Monitoring Console (AMC)

General information:

- For more information on AMC, see Chapter 16, “Administration and Monitoring,” on page 183.
- When you are ready to log into the console, browse to `http://hostname:port/ibm/console`. For more information, see section “Log in and logout of the console” on page 198.
- For more information on adding users and user roles, see section “Console user authority” on page 188.

Bundled embedded web platform deployment:

- If you installed AMC with the bundled embedded web platform and are ready to use AMC, you will need to run the commands to start AMC and Action Manager (AM) before you can log into the ISC console. For more information, see section “Starting the Administration and Monitoring Console and Action Manager and logging in” on page 184.

Note: On Windows, a shortcut to the `launchAMC.html` file is created in the start menu under Program Files.

- By default, the user who installs Tivoli Directory Integrator is the only one with access to log in to the console.

Customer or deferred deployment:

- If you chose a custom ISC AE to deploy AMC to, and are now ready to deploy, see “Deploying AMC to a custom ISC SE/AE” on page 27. When deploying AMC in ISC AE the installer does not automatically assign the current user the TDI AMC Admin role. This right needs to be manually authorized by an administrator of the ISC console. This is typically done using the **Users and Groups** -> **Administrative User Roles** panel of the ISC AE console. Alternatively this role could be assigned using the `setAMCRoles` command.
- If you chose to defer deployment of AMC into an ISC and are ready to do so, see section “Deploying AMC to a custom ISC SE/AE” on page 27.

Note: If you have done a custom ISC SE/AE deployment then at a minimum you will need to ensure that AM is started after you start the ISC SE/AE that AMC was installed into.

Documentation

The documentation system used by Tivoli Directory Integrator is the IBM Eclipse Help System. After you have done a default installation, this means that IBM Tivoli Directory Integrator 7.0 documentation is made available to you online, on the Web in the form of an Infocenter hosted by IBM. You may, however, choose to deploy the documentation locally. For more information, see “Installing local Help files” on page 26.

If you are new to Tivoli Directory Integrator, we recommend that you read and step through the *IBM Tivoli Directory Integrator V7.0 Getting Started* in order to get used to the concepts used.

If you have used earlier versions of Tivoli Directory Integrator, then chapter 3 of the *IBM Tivoli Directory Integrator V7.0 Users Guide* will be very beneficial to you in order to understand the new IDE framework and layout. It will also explain how you can import and open your existing configurations; and how the Server still uses the Config model at runtime.

Migration

If you have used an earlier version of Tivoli Directory Integrator, then you will most likely need to migrate certain aspects of your previous deployment. More information on what to do in this case can be found under Chapter 5, “Migrating,” on page 41.

Attention: Support for running the Configuration Editor (the GUI for developing solutions in IBM Tivoli Directory Integrator 7.0) has changed. The Configuration Editor (CE) is not supported on the following platforms:

- z/OS
- Linux 390
- Linux PPC
- HP PA-RISC
- HP-UX Integrity

If you wish to develop solutions for these platforms, you should use the Remote Configuration Editor functionality, meaning that you run the CE on a supported platform, while in contact with a Tivoli Directory Integrator Server on one of the aforementioned platforms. See “Using the Remote Configuration Editor” on page 104 for details.

Installing local Help files

The IBM Tivoli Directory Integrator installer does not contain any user documentation, other than the Java API documentation, which can be displayed by selecting the **Help -> Welcome** screen, **JavaDocs** link in the Configuration Editor. IBM provides the user documentation in online form in an information center, at http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.0/welcome.htm.

IBM Tivoli Directory Integrator is equipped with code¹ to provide you with context-dependent online help that you can launch from the Configuration Editor (CE). By default, this code resolves the documentation from the online information center as referenced above. You can, however, install the documentation locally, such that you are not dependent upon the Internet to be able to read it.

These are the steps you must take to install documentation locally:

- The code to handle the documentation files, the IBM Eclipse Help System, is not installed by default. In order to install the Help system, you will need to do a custom install, and install the IEHS v3.11 feature into your existing Tivoli Directory Integrator installation.
- All the manuals are stored together in one compressed directory, which when uncompressed contains an Eclipse *Document plug-in*.
- All the manuals can be downloaded, in their compressed form, from the Tivoli Directory Integrator documentation site, at http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.0/welcome.htm
- The entire documentation package, *di_plug-ins-7.0.zip*, should be uncompressed into the right place: *TDI_install_dir/ibm_help/eclipse/plug-ins* folder (or uncompress somewhere else, and move into the right place). The package contains the actual Tivoli Directory Integrator documentation in *com.ibm.IBMDI.doc_7.0*, alongside a number of other directories whose names end in *.doc*; all of those directories should be at the same aforementioned *plug-ins* level.
- The location of the documentation that the CE tries to access is set in the *global.properties* file, which resides at the root level of the installation directory of IBM Tivoli Directory Integrator, or *solutions.properties* in the Solutions Directory. By default, this points to the online information center, but if you comment the line:

```
## Name of help server, comment out if you want local help system
com.ibm.di.helpHost=publib.boulder.ibm.com/infocenter/tiv2help/index.jsp?topic=
```

such that it reads:

```
## Name of help server, comment out if you want local help system
#com.ibm.di.helpHost=publib.boulder.ibm.com/infocenter/tiv2help/index.jsp?topic=
```

then next time you run the CE and launch Help, it starts a local task to serve the documentation, from the content in the *plug-ins* directory.

- The location of the documentation server that AMC tries to access is set in its *web.xml* file. Open the *web.xml* file which is located in the WEB-INF folder of *tdiamc* webapp and list the IP address (or hostname) and port of the help server, for both occurrences of the following attributes: *InfocenterHostName* and *InfocenterPort*.

After you install the documentation in the *plug-ins* directory as outlined above, you can also decide to host the documentation on that computer for other installations of IBM Tivoli Directory Integrator in your environment. In the *TDI_install_dir/ibm_help* directory there are a number of *.bat* files (Windows) or *.sh* files (Unix/Linux) that enable you to do this.

IC_start.bat or IC_start.sh

If you run this script, the script starts an information center on http://your_IP_address:8888

1. The help system is powered by Eclipse™ technology. (<http://www.eclipse.org>)

By editing this file, you can change the port number from the default, 8888; if you want to change this, to for example 80, change "-port 8888" to "-port 80". On those clients that are trying to access this information center, the port must match another property in the `global.properties` or `solution.properties` file, `com.ibm.di.helpPort` – its default is set to 80. Also, the `com.ibm.di.helpHost` property should read something like `infocenter_IP_address/help`, where `infocenter_IP_address` is the address of your local information center. In addition, in order for AMC to find this information center, you must update the parameters `InfoCenterHostname` and `InfoCenterPort` attributes in its configuration file, `web.xml`, to match the values above.

IC_stop.bat or IC_stop.sh

Stops the help system, a Java program, that serves the local information center.

help_start.bat or help_start.sh

Similar to `IC_start`, except the port used is a random one, and it also launches a local browser showing the start page. As the port is random, unsuitable for use other than on the local computer.

help_stop.bat or help_stop.sh

Stop the local Java task that was started by `WebSphere_help_start`.

Deploying AMC to a custom ISC SE/AE

If you chose to defer deployment of AMC to ISC, and are now ready to deploy, follow these steps:

- Execute the following scripts:

```
TDI_install_dir/bin/setISCHome.bat(sh) ISC location
TDI_install_dir/bin/amc/install.bat(sh)
TDI_install_dir/bin/amc/setAMCRoles.bat(sh) username
```

Notes:

1. Calling the `setAMCRoles` script is optional for both SE and AE. If executed, `username` should be an already existing one on the ISC/WAS environment. As an alternative, you can use the ISC console ("Console User Authority" panel specifically) to manually assign one of the roles that came with AMC - "TDI AMC Admin" and "TDI AMC User" to a user.
 2. See the section "AMC in the Integrated Solutions Console" on page 188 for more information on AMC roles.
- Alter the `amc.properties` file so the lines specifying `amc.api.port` and `amc.help.port` have appropriate port values. For ISC SE this file is located in `ISC location/runtime/isc/eclipse/plugins/AMC_7.0.0/` and for ISC AE this file is located in `ISC location/systemApps/isclite.ear/tdiamc.war`

If you chose a custom ISC AE to deploy AMC to, and are now ready to deploy, follow this step:

- Execute the following script:

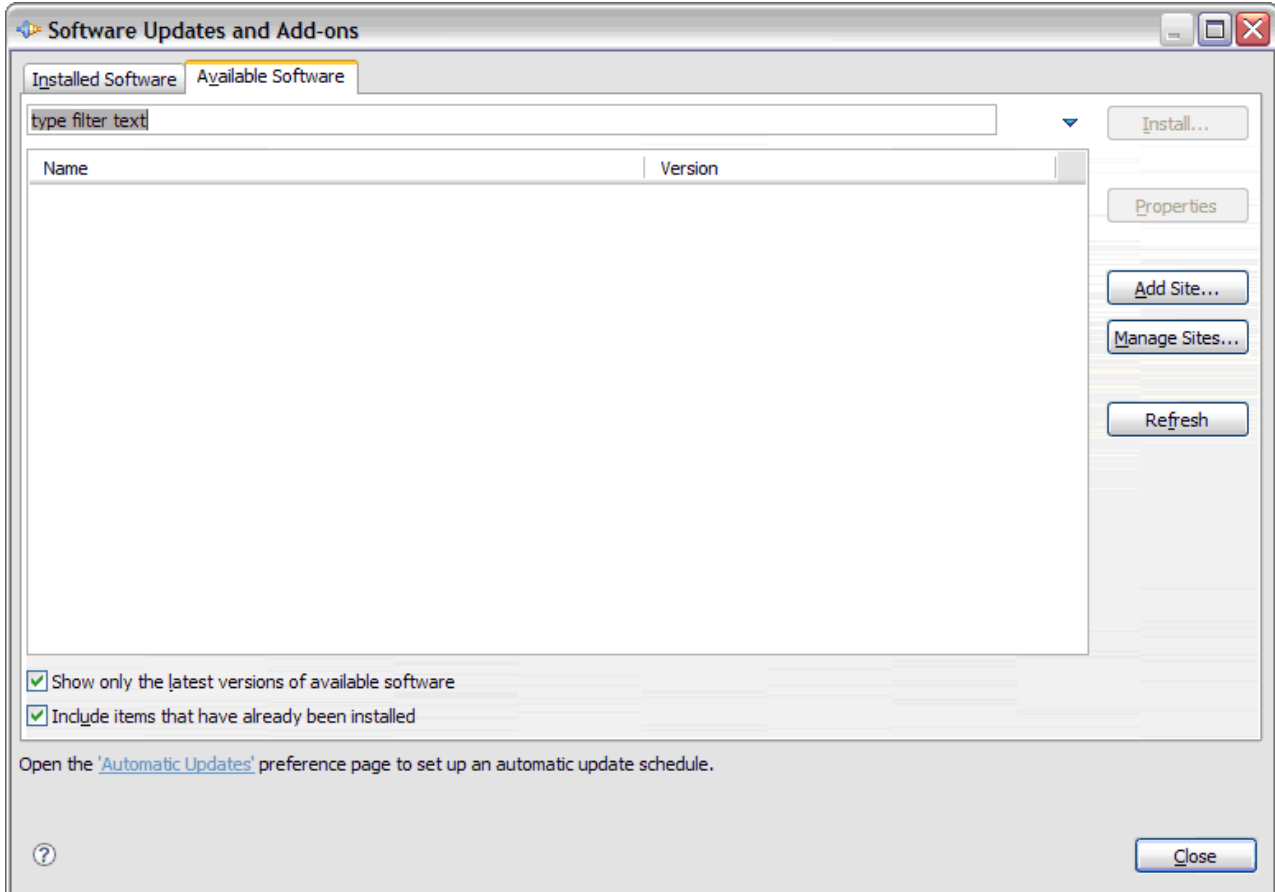
```
TDI_install_dir/bin/amc/setAMCRoles.bat(sh) username
```

Notes:


1. Calling the `setAMCRoles` script is optional for both SE and AE. If executed, `username` should be an already existing one on the ISC/WAS environment. As an alternative, you can use the ISC console ("Console User Authority" panel specifically) to manually assign one of the roles that came with AMC - "TDI AMC Admin" and "TDI AMC User" to a user.
2. See the section "AMC in the Integrated Solutions Console" on page 188 for more information on AMC roles.

Installing or Updating using the Eclipse Update Manager

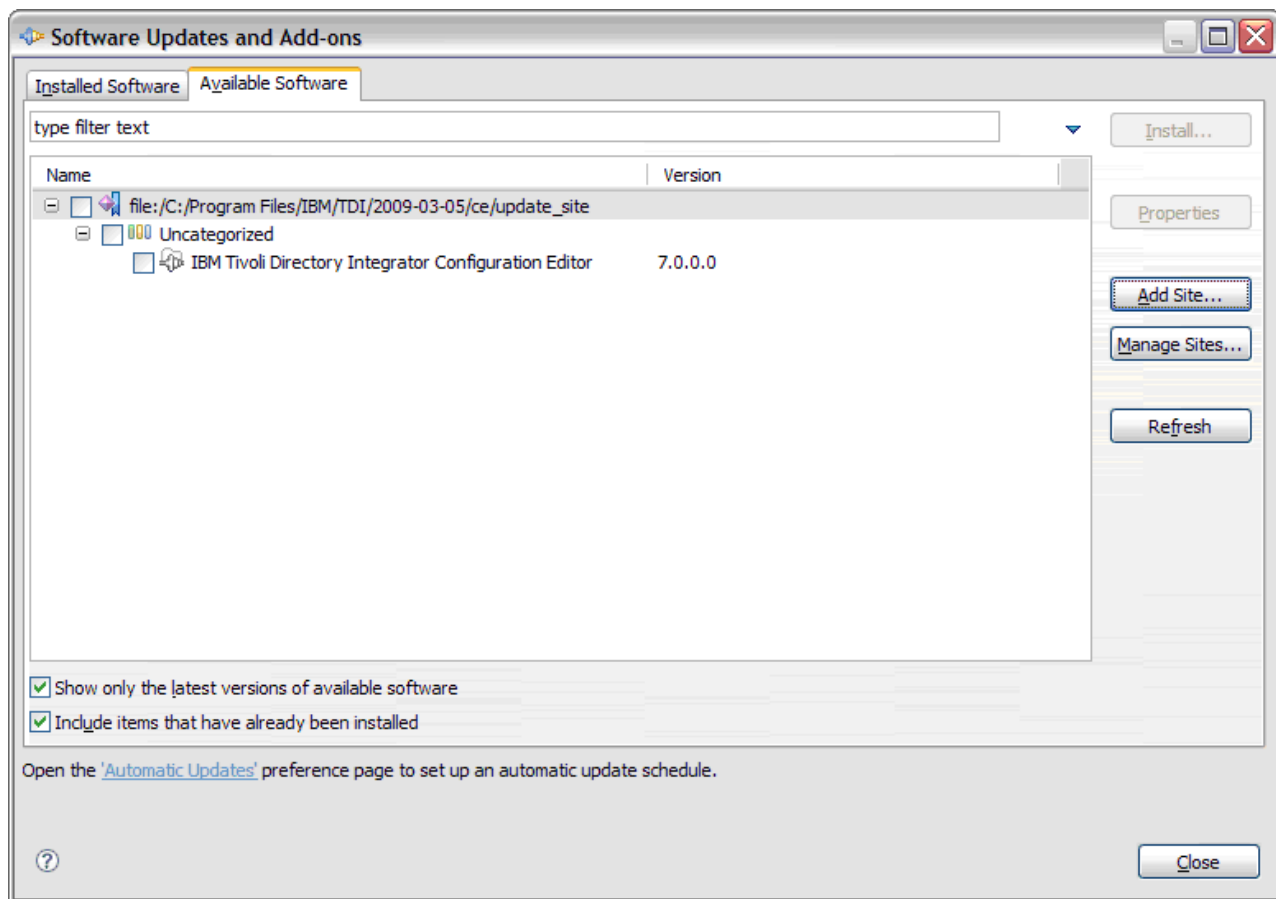
The Tivoli Directory Integrator Rich Client Platform contains a complete runtime environment to run the Tivoli Directory Integrator CE. However, it is possible to install the Tivoli Directory Integrator Eclipse plug-in into an existing Eclipse installation. This is done using the Eclipse Update Manager. Open the Eclipse Update Manager through the **Help ->Software Updates** menu option.



Before you have the Tivoli Directory Integrator plug-in installed you will want to add a new update site. Choose the **Add Site...** button and specify the location of the update site.



Depending on the location of the update site choose the appropriate action. In this example we choose a directory on the local file system. Using the **Local** button you are prompted to choose a directory which is then filled into the location input field. When you press **OK** the new update site and updates should be available:



Check the plugins you want to install and press **Install**. As the software update manager updates your installation you may be prompted to confirm the installation and you are also usually encouraged to restart the workbench after installation. After installation is complete you should see Tivoli Directory Integrator in the **Installed Software** tab.

Post-installation steps

The Tivoli Directory Integrator CE relies on a system property or environment variable named *TDI_install_dir* that points to a Tivoli Directory Integrator installation. This installation is used to create the local development server that the CE uses.

When the CE is installed as a plug-in in another Eclipse installation like in the procedure described above, two specific properties must be set to include the *TDI loader*. The TDI loader is an `org.eclipse.osgi` fragment that provides class loading for the CE.

```
# TDI class loader
org.eclipse.osgi.framework.extensions=com.ibm.tdi.loader
org.eclipse.osgi.hook.configurators.include=com.ibm.tdi.loader.TDIClassLoaderHook
```

There are several ways to set these properties. One is to update the configuration/config.ini file of the Eclipse installation.

Uninstalling

You can uninstall IBM Tivoli Directory Integrator in its entirety, or uninstall only certain components.

Launching the uninstaller

To uninstall Tivoli Directory Integrator, you must first launch the uninstaller:

Note: Before uninstalling, stop any component that you intend to remove, for example an instance of the Tivoli Directory Integrator Runtime, an AMC service that is running, or a Password Synchronization plug-in. Not stopping running components may cause some files to not be removed (to remain after the uninstallation). On Windows, a restart may be required and the Tivoli Directory Integrator Web Admin (AMC) service may remain in the Services list, requiring manual deletion.

1. Navigate to the Tivoli Directory Integrator `_uninst` directory, for example:

`install_path/_uninst`

2. Launch the uninstaller by executing the uninstall executable file.

For Windows platforms, the uninstall executable file is called `uninstaller.exe`. For all other platforms, the uninstall executable file is called `uninstaller.bin`, except for i5/OS where it is called `uninstaller.sh`.

Note: On Windows platforms, you can also uninstall using the **Add/Remove** programs in the Control Panel window.

The uninstall Welcome window appears, stating "The InstallShield Wizard will uninstall IBM Tivoli Directory Integrator 7.0 from your computer." To continue, click **Next**.

3. Select the features for "IBM Tivoli Directory Integrator v7.0 you would like to uninstall:" The features listed for uninstallation vary according to what was selected for installation. All installed features are selected for uninstallation as a default. To exclude a listed feature from uninstallation, click the check box for that item. When the item is cleared, it is not to be uninstalled. Once you have selected the features you want to uninstall, click **Next**.
4. Summary information appears for the features you have selected for uninstallation. To make changes in the uninstallation, click **Back**. To confirm the uninstallation shown in the summary information, click **Uninstall**; alternatively, click **Cancel** to stop the uninstallation.
A progress window appears, stating "Uninstalling IBM Tivoli Directory Integrator v7.0..."
5. When the uninstallation is complete, a window appears, stating "Please read the summary information below. The InstallShield Wizard has successfully uninstalled IBM Tivoli Directory Integrator v7.0. Choose Finish to exit the wizard."
6. Select from the following two restart options:

Note: This panel displays only if the installer detects that a restart is required. Generally, a restart is required if a file is locked or a if a process is running during an uninstall.

- Yes, restart my computer.
- No, I will restart my computer at a later time.

Once you have selected the restart option you want, click **Finish**.

Attention: During an uninstall, a number of directories on the computer are emptied and removed. These are:

- *TDI_install_dir/lwi* - There is the possibility that some files are left over here, or files can get created by LWI that the installer doesn't lay down. This directory is deleted on uninstall.
- *TDI_install_dir/ce/eclipsece/features/com.ibm.tdi.*.jar*
- *TDI_install_dir/ce/eclipsece/plugins/com.ibm.tdi.*.jar*
- *TDI_install_dir/ce/eclipsece/configuration*
- *TDI_install_dir/ce/update_site/features/com.ibm.tdi.*.jar* - If any features have been added that match this wildcard, they will be deleted.
- *TDI_install_dir/ce/update_site/plugins/com.ibm.tdi.*.jar* - same as previous.
- *TDI_install_dir/maintenance/BACKUP* - This directory may be created by the update installer.
- *TDI_install_dir/_uninst/** - This is needed because of the way the uninstall on iOS happens. It will be deleted regardless.

Anything you may have put into these directories yourself that matches any of these criteria, will be removed as well during an uninstall.

Performing a silent uninstallation

To perform a silent uninstallation of IBM Tivoli Directory Integrator you must first generate a response file. To generate this file, you must perform a full GUI or console uninstallation with the `-options-record` option specified; for example:

```
TDI_install_dir/_uninst/uninstaller.exe -options-record UninstallResponseFileName
```

The response file is created in the directory that you specify during uninstallation.

Note: The directory *TDI_install_dir/examples/install* contains a number of example response files for various installation or uninstallation scenarios.

Once the response file is created, you can uninstall silently using the following command:

```
TDI_install_dir/_uninst/uninstaller.exe -silent -options UninstallResponseFileName
```

Note: The examples in this document use the Windows platform uninstallation executable file.

Default installation locations

IBM Tivoli Directory Integrator installs to the following default locations:

Windows platforms

c:\Program Files\IBM\TDI\V7.0

Linux and UNIX platforms (AIX, HP-UX, Solaris)

/opt/IBM/TDI/V7.0

i5/OS /QIBM/ProdData/TDI/V7.0

Chapter 3. Update Installer

The Tivoli Directory Integrator Update Installer is an installer included with updates to existing Tivoli Directory Integrator installations; it is used whenever you need to perform maintenance in order to install fixes.

The regular installer lays down a file named `.registry` in the install directory that represents the current level of installed components. A script named `tdiSetBackupDir.bat` or `tdiSetBackupDir.sh` is created in the bin directory of the installation that sets the location of the backup directory; this will be a directory named `BACKUP` in the maintenance directory by default. You can change the backup location by running the `tdiSetBackupDir` script. So for example, if a fix is named "ifix1", backup files and directories would be under `install dir/maintenance/BACKUP/ifix1` in this scenario. The update installer will harvest the name of the backup directory when performing maintenance. The user performing maintenance to Tivoli Directory Integrator should have write permission for the install and backup directories. You should also be aware that during a complete uninstall, the uninstaller will attempt to delete the default backup directory.

The regular installer also handles maintenance of the `.registry` file during uninstalling and adding features. When performing a full uninstall, the `.registry` file will be deleted along with the other files. When performing a partial uninstall, only the components being uninstalled will be removed from the registry file, and when adding features, the `.registry` file will be updated to contain the newly installed features. After adding a feature, you should immediately install all of the fixes that are currently applied. Installing a fix that has previously been applied will only update the newly added features.

The update installer will be comprised of several Java files, but to avoid you having to specify the java executable, a wrapper script is created in the bin directory called `applyUpdates.bat(sh)`. This script will use existing scripts to find the right Java JRE to use and call the underlying code. The script's usage is as follows:

```
applyUpdates -update fix_file.zip [-clean [-silent]]
applyUpdates -rollback
applyUpdates -queryreg
applyUpdates -queryfix fix_file.zip
applyUpdates -?
```

The options are as follows:

-update

This is used to apply a fix. `fix_file` is the name of the zip file containing the fix; it can be a relative or absolute path. The **-clean** option, which is only available for a fixpack, will erase all of the files that have been backed up prior to applying the current fixpack. You will be prompted to ascertain you want to delete the old data. The **-silent** option suppresses the confirmation prompt. In the event that a fixpack is being reapplied, for example if new features have been added that need the fixpack, the **-clean** option will be ignored. Another thing to note about the **-clean** option is that cleaning the backup directories will limit the ability to rollback to a single level.

-rollback

This is used to rollback to the state Tivoli Directory Integrator was in before the most recent fix was applied.

-queryreg

This will show the features that are in the current installation as well as all of the fixes applied.

Example output:

```
Information from .registry file in: C:\Program Files\IBM\TDI\V7.0
Edition: General Purpose
Level: Full
License: 7.0.0.0
```

```
Fixes Applied
=====
None
```

```
Components Installed
=====
BASE
SERVER
CE
JAVADOCS
EXAMPLES
EMBEDDED WEB PLATFORM
AMC
Deferred: false
```

-queryfix

This will show information about the fix contained in fix_file.zip.

Example output:

```
Information from fix file: C:\fixes\TDI-7.0-TESTFP0001.zip
Name: fixpack1
Minimum level required to apply fix: 7.0.0.0
Maximum level allowed to apply fix: 7.0.0.0
```

```
Prereq
=====
None
```

```
Components Affected
=====
BASE
CE
EXAMPLES
```

The zip file representing a fix (fix_file.zip in the examples above) will contain a manifest file named .manifest which contains information about applying the fix.

-? This is for usage information.

The .registry file

Inside the install directory will be a file named .registry. This file represents the level of all Tivoli Directory Integrator components currently installed on the system in this particular install directory. This file is initially created by the installer based on the options chosen at install time.

When a fix is installed, the backed up files will be stored in a directory with the name of the fix inside the backup directory, and additional entries will be made to the .registry file that represent the changes made to components by a fix. There will be a FIXES section of the .registry file that represents the fixes that have been applied, and each component will have entries representing which fixes have been applied that have altered them.

The Update Installer recognizes the following components:

- BASE
- SERVER
- Configuration Editor (CE)
- CE_UPDATE

- JAVADOCS
- EXAMPLES
- IEHS
- embedded Web platform
- Administration and Monitoring Console (AMC)
- PLUGINS

The plugins component may require some steps that cannot be performed by the update installer and must be performed manually. If there is something to be changed in any of the `pwsync.props` files (not very likely), you must perform this manually following steps in the `manual_readme.txt` file of the fix pack. This must happen after installing the fix pack, but before any of the post-install steps listed below. The readme will warn you that the Update Installer will only update files, which the Installer has put down. Files that are copied by you, need to be updated manually as described in the post-install steps.

Here is a list of steps that you need to perform pre/post installing the fix pack (as noted, these steps should be executed only if you have registered the corresponding Password Synchronizers into target systems):

Windows Password Synchronizer

Pre-install steps: None

Post-install steps (only if the fix pack contains an update of the DLL of the Password Synchronizer):

1. Remove the name of the DLL of the Password Synchronizer from the following registry key: `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA\Notification Packages`. The DLL file is named "tdipwflt" on 32-bit Windows and "tdipwflt_64" on 64-bit Windows.
2. Reboot Windows, so that the Local Security Authority (LSA) process unloads the DLL of the Password Synchronizer.
3. Replace the DLL inside the `system32` Windows folder with the one from the installation of the Password Synchronizer. The DLL path after installation is either `install_dir/pwd_plugins/windows/tdipwflt.dll` or `install_dir/pwd_plugins/windows/tdipwflt_64.dll` depending on the version of Windows.
4. Add the name of the DLL (without the ".dll" extension) inside the registry key: `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\LSA\Notification Packages`
5. Reboot Windows again, so that the LSA loads the new DLL.

Afterwards the Password Synchronizer should run normally, using the updated files.

IBM Directory Server Password Synchronizer

Pre-install steps:

1. Stop the Directory Server.
2. Stop the Proxy process of the Password Synchronizer using the `stopProxy` command-line utility. This is necessary because the IBM Directory Server Password Synchronizer does not automatically stop its Proxy when being terminated.

Post-install steps: None.

Sun Directory Server Password Synchronizer

Pre-install steps: Stop the Directory Server.

Post-install steps: None.

PAM Password Synchronizer

Pre-install steps: If possible, avoid any password changes while the update takes place. Otherwise unregister the Password Synchronizer from the PAM configuration file.

Post-install steps: If you have unregistered the Password Synchronizer before the update, register it again; see *IBM Tivoli Directory Integrator V7.0 Password Synchronization Plug-ins Guide* for more information.

Domino® Password Synchronizer

Pre-install steps: None

Post-install steps: follow the post-install instructions from the chapter "Domino HTTP Password Synchronizer" in *IBM Tivoli Directory Integrator V7.0 Password Synchronization Plug-ins Guide*, followed by a new setup of the Domino plug-in as described in the section "Deployment on a single Domino Server".

Installing fixes

If a fix file contains a fix for a component and that component is installed on the system, a number of programmed actions will be performed for the individual components.

If there are any manual steps that must be performed outside of the update installer, instructions will be included in the readme file for the fix; in addition see here for additional remarks about Password Plugins fixes.

Rollback

During a Rollback, the Update Installer uses information previously laid out during a fix, and files backed up, to restore a previous state.

Troubleshooting

The Update Installer creates a log file named `updateinstaller.log` in the `install_dir/logs` directory. The default level of messages is INFO, but this can be changed by altering the `install_dir/etc/updateinstaller-log4j.properties` file so that DEBUG messages are also logged.

Chapter 4. Supported platforms

Some of the software components in IBM Tivoli Directory Integrator, like the Launchpad and the Administration and Monitoring Console, require a configured Web browser. Supported Web browsers are:

- Microsoft Internet Explorer
- Mozilla Firefox

These are the platforms supported for IBM Tivoli Directory Integrator 7.0. For the latest list of platforms supported see: <https://www-304.ibm.com/support/docview.wss?uid=swg21376850>.

This URL can also be used to see the platform supported by the Tivoli Directory Integrator plug-ins and the bit support per platform.

Note: On AIX®, Linux on Power, Sun Solaris SPARC, native z/OS and i5/OS, the Tivoli Directory Integrator Server, Configuration Editor (CE) and other components run in 32-bit tolerance mode because it ships and uses a 32-bit JRE. This would be 31-bit tolerance on native z/OS.

- Microsoft Windows Intel x86
 - Windows 2003 Standard, Enterprise and Datacenter Edition (32-bit supported)
 - Windows 2008 Standard, Enterprise and Datacenter Edition (32-bit supported)
- Microsoft Windows Intel x86-64
 - Windows 2003 Standard, Enterprise and Datacenter Edition (64-bit supported)
 - Windows 2008 Standard, Enterprise and Datacenter Edition (64-bit supported)
- AIX
 - AIX 5L 5.3 (32/64-bit supported) (Recommended Maintenance package 5300-03 is required)
 - AIX 6.1 (11/2007) (32/64-bit supported) (WPARs are supported)
- Sun Solaris SPARC
 - Solaris 9 (32/64-bit supported)
 - Solaris 10 (32/64-bit supported) (Zones are supported)
- Sun Solaris Opteron
 - Solaris 10 (64-bit supported)
- HP-UX PA-RISC
 - HP-UX11iv2 (11.23) (32/64-bit supported)
 - HP-UX11iv3 (32/64-bit supported)
- HP-UX Integrity
 - HP-UX11iv2 (11.23) (64-bit supported)
 - HP-UX11iv3 (64-bit supported)
- Linux Intel x86
 - RedHat Enterprise Linux ES/AS 4.0 (32-bit supported)
 - RedHat Enterprise Linux ES/AS 5.0 (32-bit supported)
 - SLES 9 (32-bit supported)
 - SLES 10 (32-bit supported)
 - Red Flag Data Center 5.0 SP1 / Asianix 2.0 SP1 (32-bit supported)

Note: A prerequisite on RedHat Enterprise Linux ES/AS 4.0 and Red Flag Data Center 5.0 SP1 is that library package compat-libstdc++-296-2.96-132.7.2 or above is installed for the Installer to work.

For RedHat Enterprise Linux ES/AS 5.0, compat-libstdc++-33-3.2.3-61.i386.rpm and libXp-1.0.0-8.i386.rpm provided on 2nd installation CD for RHEL5 are required.

- Linux Intel x86-64
 - RedHat Enterprise Linux ES/AS 4.0 (64-bit supported)
 - RedHat Enterprise Linux ES/AS 5.0 (64-bit supported)
 - SLES 9 (64-bit supported)
 - SLES 10 (64-bit supported)

Note: A prerequisite on RedHat Enterprise Linux ES/AS 4.0 is that library package compat-libstdc++-296-2.96-132.7.2 or above is installed for the Installer to work.

For RedHat Enterprise Linux ES/AS 5.0, compat-libstdc++-33-3.2.3-61.x86_64.rpm and libXp-1.0.0-8.x86_64.rpm provided on 3rd installation CD for RHEL5 are required.

- Linux on Power (pSeries, iSeries, OpenPower and JS20 Blades)
 - RedHat Enterprise Linux ES/AS 4.0 (32/64-bit supported)
 - RedHat Enterprise Linux ES/AS 5.0 (32/64-bit supported)
 - SLES 9 (32/64-bit supported)
 - SLES 10 (32/64-bit supported)

Note: A prerequisite on RedHat Enterprise Linux ES/AS for this architecture is that library package compat-libstdc++-296-2.96-132.7.2 or above is installed for the Installer to work.

For RedHat Enterprise Linux ES/AS 5.0, compat-libstdc++-33-3.2.3-61.ppc.rpm and libXp-1.0.0-8.ppc.rpm are required.

- Linux s/390 and zSeries®
 - RedHat Enterprise Linux ES/AS 4.0 (64-bit supported)
 - RedHat Enterprise Linux ES/AS 5.0 (64-bit supported)
 - SLES 9 (64-bit supported)
 - SLES 10 (64-bit supported)

Note: A prerequisite on RedHat Enterprise Linux ES/AS for this architecture is that library package compat-libstdc++-296-2.96-132.7.2 or above is installed for the Installer to work.

For RedHat Enterprise Linux ES/AS 5.0, compat-libstdc++-33-3.2.3-61.s390.rpm and libXp-1.0.0-8.s390.rpm are required.

- Native s/390 and zSeries
 - z/OS 1.8, 1.9 and 1.10 (31/64-bit supported) – with limitations; refer to Chapter 19, “z/OS environment Support,” on page 245
- iSeries
 - i5/OS V5R4 and V6R1 (32/64-bit supported); the CE is not supported on this platform. In addition, the JVM or embedded Web platform to be used is the system-installed one(s); no JVM or embedded Web platform is installed as part of the Tivoli Directory Integrator product install.
The Tivoli Directory Integrator server on i5/OS requires that IBM Technology for Java Virtual Machine (also known as “J2SE 5.0 32-bit JVM” or “IBM J9 VM”) be installed on the system
The installation can only be performed locally, as an Administrator, and using a console installation. No GUI installation facility exists for i5/OS.

Notes:

1. “x86” denotes an Intel-architecture 32-bit CPU; “x86-64” denotes an Intel-architecture 64-bit CPU. 64-bit CPUs generally support 32-bit and 64-bit versions of assorted operating systems; 32-bit CPUs only support 32-bit versions.

2. The Tivoli Directory Integrator Config Editor is not supported locally on HP-UX Integrity, HP-UX PA-RISC, Solaris Opteron, Linux on Power, Linux s/390, Native s/390 and zSeries, or i5/OS. The Config Editor will have to be run remotely from a supported operating system.
3. The Administration and Monitoring Console is not supported locally on Native s/390 and zSeries. The Administration and Monitoring Console will have to be run remotely from a supported operating system.
4. The Windows x86 and Linux x86 installers can be used on Windows x86-64 and Linux x86-64 platforms if a 32-bit JRE is needed on these 64-bit platforms. However, the 32-bit installer is not able to install the Password interceptor Plug-ins on a 64-bit platform, as the libraries to support such operations do not exist. If you need both a 32-bit Tivoli Directory Integrator installation and Password Plug-ins, you will need to install a second, 64-bit version of Tivoli Directory Integrator in a different location first, and then use a Custom install to install the Plug-ins using that 64-bit installation.

Virtualization support

Virtualization technology allows a single physical machine to be partitioned into multiple physical or logical partitions with each partition providing the look and feel of an independent operating system. Each partition is called a Virtual Environment. Each Virtual Environment represents a complete system, with processors, memory, networking and other system resources. Examples of some virtualization technologies include Sun Solaris Zones, AIX WPARs, and VMware.

IBM Tivoli Directory Integrator 7.0 currently provides support for these virtualization technologies:

- VMware; see the following for specifics: <http://www-01.ibm.com/support/docview.wss?rs=697&context=SSCQGF&dc=DB520&q1=VMWare&uid=wws1e333ce0912f7b152852571f60074d175&loc>
- Sun Solaris Zones
- AIX WPARs

Chapter 5. Migrating

In the context of Tivoli Directory Integrator, "migrating" can mean a number of things.

- Prepare relevant files (and their contents) to be used in a new location, on the same machine or a different one; or
- prepare relevant files to be used with a new version of the product.

The following table summarizes migration scenarios:

Table 1. Migration scenarios

Source & Destination versions equal?	Source & Destination install paths equal?	Scenario description
no	no	Migrate files to a new version, which is installed in a different location
no	yes	Migrate files to a new version, which will be installed in the same location
yes	no	Migrate files to a different installation of the same version
yes	yes	Restore backed up files to their original location

If you have to both migrate to a new version and a new location, you should do the version upgrade first, because here we will cover location migration only for the current release (7.0).

The IBM Tivoli Directory Integrator Installer can assist in migrating from Tivoli Directory Integrator 6.0 and Tivoli Directory Integrator 6.1.X to Tivoli Directory Integrator 7.0.

Migrate files to a different location

In this section we cover only 7.0.

Which files do not need to be modified to be used in another location?

- User configurations, data files (.xml, .xsd, .xsl, .txt ...), key store files (.jks, ...), certificate files (.der, ...), and so forth.

Also consider the implications of section "Maintaining encryption artifacts – keys, certificates, keystores, encrypted files" on page 137.

- Scripts (see this section for exceptions)
- .bat, .sh and .vbs files
- JAR files
- Native binaries - .exe, .dll, .so, and so forth.
- Server API registry file
- Server Stash File
- Derby databases:

For example the default System Store database "TDISysStore" and the default AMC database "tdiamcdb".

Note that you must move the database as a whole (the whole folder). You should not merge the files of two databases.

For more complicated scenarios, transfer data between databases using the JDBC Connector.

- Action Manager property files (located in the *TDI_install_dir/bin/amc/ActionManager* folder)

- Configuration files from the "etc" folder except these:
 - build.properties
 - global.properties
 - updateinstaller-log4j.properties
 - tdisrvctl-log4j.properties
- AMC configuration files:
 - amc.properties
 - amcdbhandler.properties
 - amcdbschema.xml
 - idiamc.sth

Which files need to be modified before they can be used in another location?

In general, location-sensitive files will contain the absolute path of the installation folder in one or more places. These occurrences need to be replaced with the new location path, so that the file becomes relevant to the new location.

Below is a list of the files that need migration and hints about which fields to update. These hints are based on the default content of these files. If you have modified the files, there may be other fields that are also location specific and need to be updated too.

- bin/amc/amcwinbservice.ini:
This is the configuration file for AMC when registered as a Windows service.
Update the "WorkingDirectory", "StartCommand" and "StopCommand" properties.
- global.properties/solution.properties:
Update the "com.ibm.di.store.database" property.
Also consider these properties: "api.config.folder", "systemqueue.jmsdriver.param.mqe.file.ini" and "com.ibm.di.loader.userjars".
If you migrate the file to an installation that uses a different encryption key, see section "Maintaining encryption artifacts – keys, certificates, keystores, encrypted files" on page 137.
- etc/updateinstaller-log4j.properties:
Update the "log4j.appender.Default.file" property.
- etc/tdisrvctl-log4j.properties:
Update the "log4j.appender.Default.file" property.
- ibmdiservice.props:
This is the configuration file for the Server when registered as a Windows service.
Update the "path", "ibmdiroot" and "jvmRoot" properties.
- pwsync.props:
These are the configuration files of the Password Synchronizers.
Update the "proxyStartExe", "logFile", "javaLogFile" and "mqe.file.ini" properties.

Which files should not be used in another location under normal circumstances?

- Certain scripts
The sole purpose of the existence of these files is to convey location specific data.
There is virtually nothing else in them, so they would be of little value in another location.
Consider these scripts from the *TDI_install_dir*/bin folder: "javaHome", "defaultSolDir", "backupDir", "tdiISCHome".

- .reg files
These are used by the Windows Password Synchronizer.
- MQe queue manager files
Although you cannot easily migrate MQe files to another location, you can transfer data from one MQe queue to another using the JMS Connector with an MQe JMS driver.
- CE workspace
To reuse Directory Integrator projects from the Config Editor workspace, export them as Directory Integrator configurations and import them into the new workspace.
- etc/build.properties
This file contains time and version information about the release of the product.

Migrating files that contain encrypted data

See “Maintaining encryption artifacts – keys, certificates, keystores, encrypted files” on page 137.

Migrate files to a newer version

Installer-assisted migration

The installer migrates certain files automatically during upgrade to a newer version. Note that the installer considers only the installation folder of the Directory Integrator.

All solution folders that are different than the installation folder must be migrated manually (or using some of the tools described in section “Tool-assisted migration”).

Which files does the installer migrate automatically?:

- 6.0 to 7.0
 - global.properties
 - Cloudscape database (if used for System Store) is upgraded to Derby 10.3. See “Migrating Cloudscape database to Derby” on page 60.
- 6.1.x to 7.0
 - global.properties
 - the AMC database
 - amc.properties
 - am_config.properties

Which files need to be migrated manually?:

Everything mentioned in section “Manual migration” on page 44, except those mentioned in its first subsection, Property Files.

Tool-assisted migration

These tool are used by the installer for the installer-assisted migration. You can use them for manual migration.

Property files migration:

- global.properties:
Use the "tdimigbl" tool from *TDI_install_dir/bin*; see section “Migrating global and solution properties files using the migration tool” on page 60.
- amc.properties:

Use the "tdimigamc" tool from *TDI_install_dir/bin/amc*; see "AMC and AM Command line utilities" on page 220.

- am_config.properties:

Use the "tdimigam" tool from *TDI_install_dir/bin/amc*; see "AMC and AM Command line utilities" on page 220.

AMC database migration:

Use the "backupamcdb"/"restoreamcdb" tools from *TDI_install_dir/bin/amc*; see "AMC and AM Command line utilities" on page 220.

Cloudscape System Store migration (only for 6.0):

See the more detailed instructions in section "Migrating Cloudscape database to Derby" on page 60.

Manual migration

Copy your Config files and any other custom files, including Derby databases from your old installation directory to the new installation directory. Tivoli Directory Integrator 7.0 supports a Solution Directory, and we recommend you copy the Config files, property files, Derby databases, and so on, to such a solution directory instead of to the installation directory of Tivoli Directory Integrator version.

Once you have copied the objects referenced above to a new location, you can set out to manually migrate their contents to adapt them for use with IBM Tivoli Directory Integrator 7.0 as described in the sections below:

1. Property Files
2. Configurations
3. Customized scripts
4. Added or replaced JAR files in the installation
5. Password Synchronizer configurations

Note: Sandbox data is version-specific; data recorded under any previous version does not play in version 7.0.

Property files:

- global.properties:

The table below lists which properties have been deleted or changed in Tivoli Directory Integrator 7.0:

Table 2. Deleted and changed properties

Old property (pre-v7.0)	New property	Remarks
## Active Correlation Technology engine settings # act.engine.rule.set.file=myrules.acts	*DELETED*	Remove ACT Engine and ACT Connector
# Location of directory where the JRE TDI will use is installed com.ibm.di.jvmDir=\$jvmRoot\$	*DELETED*	No longer possible to specify.
com.ibm.di.scriptengine.precompile=true	*DELETED*	No longer possible to specify; the current script engine does not have this functionality.
com.ibm.di.scriptengine.regex=java	*DELETED*	No longer possible to specify - Java syntax is always followed.
ibmjs.options=com.ibm.di.script.ScriptEngineOptions	*DELETED*	Related to previous property; this is no longer a valid option.

Table 2. Deleted and changed properties (continued)

Old property (pre-v7.0)	New property	Remarks
com.ibm.di.store.create.checkpoint.store=<multiple statements>	*DELETED*	Checkpoint/Restart functionality is removed; any System Store create table statements related to this should be removed too.
com.ibm.di.admin.library.dir=	*DELETED*	The current Config Editor does not use this, so no longer possible to specify.
api.remote.on=false	api.remote.on=true	RMI enabled by default in TDI Server – Setting to true since it is enabled by default.
javax.net.ssl.trustStore= (protect)-javax.net.ssl.trustStorePassword= javax.net.ssl.trustStoreType=	javax.net.ssl.trustStore=serverapi\testadmin.jks (protect)-javax.net.ssl.trustStorePassword=administrator javax.net.ssl.trustStoreType=jks	RMI enabled by default in TDI Server – empty values replaced by the default truststore.
javax.net.ssl.keyStore= (protect)-javax.net.ssl.keyStorePassword= javax.net.ssl.keyStoreType=	javax.net.ssl.keyStore=serverapi\testadmin.jks (protect)-javax.net.ssl.keyStorePassword=administrator javax.net.ssl.keyStoreType=jks	RMI enabled by default in TDI Server – empty values replaced by the default keystore.
com.metamerge.securityTransformation=DES/ECB/NoPadding	com.ibm.di.securityTransformation=DES/ECB/NoPadding	FIPS 140-2 Certification – property name changed.
com.ibm.di.server.keystore=myKeyStore.jks com.ibm.di.server.key.alias=myKeyAlias	api.keystore=myKeyStore.jks api.key.alias=myKeyAlias	Server API keystore properties renamed.
com.ibm.di.store.database=TDISysStore com.ibm.di.store.jdbc.driver=org.apache.derby.jdbc.EmbeddedDriver com.ibm.di.store.jdbc.url=prefix=jdbc:derby: com.ibm.di.store.jdbc.user=APP	#com.ibm.di.store.database=TDISysStore #com.ibm.di.store.jdbc.driver=org.apache.derby.jdbc.EmbeddedDriver #com.ibm.di.store.jdbc.url=prefix=jdbc:derby: #com.ibm.di.store.jdbc.user=APP	The EMBEDDED MODE properties for the System Store have been commented out, since the System Store now runs in Network mode by default. The Installer never makes this change; if you have previously used Cloudscape/Derby in embedded mode you will need to make this change manually.
#com.ibm.di.store.database=jdbc:derby://localhost:1567/TD1SysStore;create=true #com.ibm.di.store.jdbc.driver=org.apache.derby.jdbc.ClientDriver #com.ibm.di.store.jdbc.url=prefix=jdbc:derby: #com.ibm.di.store.jdbc.user=APP #com.ibm.di.store.jdbc.password=APP #com.ibm.di.store.jdbc.start.mode=automatic #com.ibm.di.store.jdbc.host=localhost #com.ibm.di.store.jdbc.port=1567 #com.ibm.di.store.jdbc.sysIBM=true	com.ibm.di.store.database=jdbc:derby://localhost:1567/TD1SysStore;create=true com.ibm.di.store.jdbc.driver=org.apache.derby.jdbc.ClientDriver com.ibm.di.store.jdbc.url=prefix=jdbc:derby://localhost:1567/ com.ibm.di.store.jdbc.user=APP com.ibm.di.store.jdbc.password=APP com.ibm.di.store.jdbc.start.mode=automatic com.ibm.di.store.jdbc.host=localhost com.ibm.di.store.jdbc.port=1567 com.ibm.di.store.jdbc.sysIBM=true	These are the new, default properties for the System Store in Tivoli Directory Integrator 7.0. If you have migrated your installation, you will need to make these changes to your global.properties file as well, if you wish to run the System Store in Networked mode. The new architecture of the Configuration Editor in conjunction with other changes to the development process make that running System Store in embedded mode is very cumbersome. Therefore, we highly recommend that you run in Networked mode.
api.config.folder=\$change\$/configs	api.config.folder=configs	The configs folder is now always local to the Solution Directory.

Table 2. Deleted and changed properties (continued)

Old property (pre-v7.0)	New property	Remarks
<pre>##----- ## System Queue settings ##----- ## If set to "true" the System Queue is initialized on startup and can be used; ## otherwise the System Queue is not initialized and cannot be used. systemqueue.on=false ### MQe JMS driver initialization properties ## Specifies the location of the MQe initialization file. ## This file is used to initialize MQe on TDI server startup. systemqueue.jmsdriver.param.mqe.file.ini=\${changed}/MQePKStore/pwstore_server.ini</pre>	<pre>##----- ## System Queue settings ##----- ## If set to "true" the System Queue is initialized on startup and can be used; ## otherwise the System Queue is not initialized and cannot be used. systemqueue.on=true ### MQe JMS driver initialization properties ## Specifies the location of the MQe initialization file. ## This file is used to initialize MQe on TDI server startup. systemqueue.jmsdriver.param.mqe.file.ini=MQePKStore/pwstore_server.ini</pre>	The System Queue is now enabled by default in Tivoli Directory Integrator 7.0 (except on z/OS). Also, the MQe initialization file is now located in a directory subordinate to the Solution Directory.

The table below lists which properties have been added in Tivoli Directory Integrator 7.0:

Table 3. New properties

Property	Remarks
<code>com.ibm.di.server.fipsmode.on=false</code>	New property for enabling/disabling FIPS mode was added.
<pre>## To enable the built-in JAAS Authentication mechanism, ## set this property to "[jaas]". api.custom.authentication ## JAAS Authentication properties ## ----- ## java.security.auth.login.config=</pre>	Provide support for JAAS as a Server API Authentication provider; empty property is provided in which you can specify the JAAS Configuration file.
<pre>## Encryption certificate properties com.ibm.di.server.encryption.keystore = <<value of com.ibm.di.server.keystore from 6.1.1 global.properties>> com.ibm.di.server.encryption.key.alias = <<value of com.ibm.di.server.key.alias from 6.1.1 global.properties >> ## Server API keystore passwords {protect}-api.keystore.password= << keystore password from idisrv.sth>> {protect}-api.key.password= << key password from idisrv.sth if present>></pre>	Provide separate configuration options for certificate to be used for PKI Encryption and SSL.
<pre>## TDI Logging com.ibm.di.logging.enabled=true</pre>	Provide mechanisms to completely disable logging - set to "false" if you want to disable all logging.
<pre>derby.connection.requireAuthentication=true derby.authentication.provider=BUILTIN derby.database.defaultAccessMode=fullAccess</pre>	Additional parameters for the System Store (in Derby) in Networked mode.
<pre>##PKCS11 options ##Set the value of following properties to use PKCS11 enabled ## devices to store TDI servers private key / certificate. com.ibm.di.pkcs11cfg=etc\pkcs11.cfg com.ibm.di.server.pkcs11=false com.ibm.di.server.pkcs11.library= com.ibm.di.server.pkcs11.slot= {protect}-com.ibm.di.server.pkcs11.password=PASSWORD</pre>	Support TDI Server's private key/certificate on PKCS 11 compliant crypto devices.
<pre>## Specify the unique ID for the TDI Server ## ----- ## This property helps a client connecting to the TDI server to identify different servers ## running on the same IP and the same port in different time. (Default is blank) com.ibm.di.server.id=</pre>	TDI Server must provide a unique server ID available to remote server clients to detect the server being talked to.
<pre>## Timeout in minutes for loading configuration. api.config.load.timeout=2</pre>	Config initialization and Server API initialization need to be synchronized.
<pre>com.ibm.di.server.encryption.keystoretype = jks com.ibm.di.server.encryption.transformation = RSA</pre>	Symmetric Cipher Support (FIPS 140-2 compliance).
<pre>## Specifies a list of Server notification types, which will be suppressed. ## Notifications of suppressed types will not be propagated by the notifications framework. ## The notification types in the list are separated by spaces. Wildcards may be included. ## Example: ## api.notification.suppress=di.al.* di.ci.start ## The above example will suppress all AssemblyLine related notifications as well as ## notifications for starting a configuration instance. ## If the property is missing or is empty, no notifications will be suppressed. api.notification.suppress=di.server.api.authenticate di.server.api.authorize.*</pre>	Provide TDI Audit Capabilities - Server notification suppression.
<code>api.audit.on=false</code>	Provide TDI Audit Capabilities.

Table 3. New properties (continued)

Property	Remarks
<pre>## This property specifies whether LDAP Group authentication is turned on. ## If it is set to 'true', the group membership of the authenticating user ## will be resolved and will be taken into account during authorization. ## If it is missing, the default value 'false' is used. api.custom.authentication.ldap.groupsupport=false ## Specifies the name of the attribute of a user in LDAP that contains a list of the groups of which the user is a member. ## It is taken into account only if 'api.custom.authentication.ldap.groupsupport' is set to true. api.custom.authentication.ldap.usermembershipattribute= ## Specifies how groups are named in the membership attribute of a user. ## For example, if the user's membership attribute contains values, ## which correspond to the 'objectSID' attributes of groups, set this property to 'objectSID'. ## If the user's membership attribute contains distinguished names of groups, then set this property to 'dn'. ## The property is required in case 'api.custom.authentication.ldap.groupsupport' is set to true. api.custom.authentication.ldap.usermembershipattributecontent= ## Specifies the name of a group's attribute in LDAP which corresponds to the way the group is named in the TDI User Registry. ## For example, if LDAP groups are addressed in the TDI registry by their common name, then set this property to 'cn'. ## If the User Registry contains the distinguished names of the groups, then set this property to 'dn'. api.custom.authentication.ldap.groupnameattribute= ## Represents the LDAP directory context, where groups will be searched. ## It is required only when LDAP group support is enabled api.custom.authentication.ldap.groupsearchbase= ## Optional property, which represents a list of space-separated attribute names. ## Specifies attributes which have non-string syntax. ## api.custom.authentication.ldap.binaryattributes=</pre>	Enhance Authorization to support LDAP groups.

- amc.properties:

The table below lists which properties have been deleted or changed in Tivoli Directory Integrator 7.0:

Table 4. Deleted and changed properties in AMC

Old property (pre-v7.0)	New property	Remarks
AMC.auth	*DELETED*	
monitor.refresh.rate	*DELETED*	The refresh rate for the Monitor Status panel. The rate was specified in minutes.
monitor.startup	*DELETED*	To set the Monitor Status panel as the first panel to be seen by the user when (s)he logs in.
LDAPHostName LDAPPort LDAPAdminUIId LDAPAdminPwd LDAPServerType LDAPBindID LDAPBindPassword LDAPSuffix LdapUserPrefix LDAPUserSuffix LdapGroupPrefix LDAPGroupSuffix LDAPUserObjectClass LDAPGroupObjectClass LDAPGroupMember LDAPUserFilter LDAPGroupFilter LDAPsearchTimeout LDAPsslEnabled LDAPIgnoreCase	*DELETED*	LDAP Details.
com.ibm.di.amc.jdbc.start.mode	New default value: Automatic	
com.ibm.di.amc.jdbc.host	New default value: Localhost	
com.ibm.di.amc.jdbc.port	New default value: 1528	
com.ibm.di.amc.jdbc.sysibm	New default value: True	

The table below lists which properties have been added in Tivoli Directory Integrator 7.0:

Table 5. New properties in AMC

New property (default)	Remarks
al.workEntries.cacheSize (100)	This property is used by AMC when the AssemblyLine is started in Synchronous mode. The cache size specified here is used for determining the size of the work entries cache.
amc.db.type (derby)	Specifies the database being used by the AMC.
am.api.host (localhost)	Action Manager RMI Details.
am.api.port (13104)	Action Manager RMI Details.

- am_config.properties:

The table below lists which properties have been deleted or changed in Tivoli Directory Integrator 7.0:

Table 6. Deleted and changed properties in AM

Old property (pre-v7.0)	New property	Remarks
com.ibm.di.amc.am.serverapi.fail.interval.time=120 com.ibm.di.amc.am.queryProperty.interval.time=600 com.ibm.di.amc.am.healthAL.interval.time=5	*DELETED*	These properties should be commented out as the values for these properties will be configured by the user from AMC while creating each Server API Failure Trigger, On Property trigger and Configuring a Health AL respectively. Hence the properties mentioned in the am-config.properties file will not be used.
com.ibm.di.amc.am.queryAL.interval.time	*DELETED*	
javax.net.ssl.trustStore=\$change\$/bin/amc/ActionManager/testadmin.jks javax.net.ssl.keyStore=\$change\$/bin/amc/ActionManager/testadmin.jks	javax.net.ssl.trustStore=bin/amc/ActionManager/testadmin.jks javax.net.ssl.keyStore=bin/amc/ActionManager/testadmin.jks	Truststore files are now local to Solution Directory.

The table below lists which properties have been added in Tivoli Directory Integrator 7.0:

Table 7. New properties in AM

New property (default)	Remarks
smtp.host= smtp.port= smtp.user= smtp.password=	SMTP server details, added in TDI 7.0.

Configurations:

Certain Directory Integrator components/features have been modified or removed. Configurations that reference these need to be migrated manually. Here is a list of affected components/features:

- Checkpoint/restart functionality:

This functionality is removed in 7.0. This leaves Connectors that support Iterator mode with only the default ability to do a simple reconnect and automatically skip forward as many times as the number of successful reads. The assumption is that skipping forward this number of entries would get you back to where you last left off. Most Tivoli Directory Integrator Connectors will not automatically attempt to do this, because the behavior can be indeterminate or not appropriate. However, the default behavior is specific per Connector. The ability to automatically skip forward as many times as the number of successful reads is a new reconnect option available to each Connector and is configured in the Connection Errors panel, see **The Configuration Editor -> The Connector Editor -> Connection**

Errors in the *IBM Tivoli Directory Integrator V7.0 Users Guide*. If you require more than the ability to automatically skip entries processed, you need to use one of the following options in your solutions:

- Configure Delta for an Iterator mode for dynamically changing result sets.
- Override the `on_connection_failure` hook and do custom reconnect logic.
- Derby/Cloudscape in embedded mode as System Store used by multiple JVMs:
Default and recommended behavior in Tivoli Directory Integrator 7.0 is running Derby in networked mode. If you continue to use Derby in embedded mode, considerations regarding multiple JVMs attempting to use the same database simultaneously still apply; see “Using Derby to hold your System Store” on page 149. For migrating databases, see “Migrating Cloudscape database to Derby” on page 60.
- Exchange Changelog Connector:
This Connector is removed in v7.0. You may consider using the unsupported Exchange Changelog Connector connector that is now provided as an "example" in `TDI_install_dir/examples/ExchangeChangeLogConnector`.
- Btree Connector:
This Connector is removed from the default installation in v7.0. Use the System Store Connector instead as described in section “Migrating BTree tables and BTree Connector to System Store” on page 59; alternatively, use the (unsupported) Btree connector that is now provided as an "example" in `TDI_install_dir/examples/BTreeDBConnector`.
- Domino Change Detection Connector:
This applies to 6.0 and 6.1 only.
The **Delivery Mode** parameter is removed and **State Key Persistence** will be used instead. The behavior of old configurations which use this parameter will be as follows:
 - If the **Delivery Mode** parameter is set to "Assured once and only once delivery" mode then the **State Key Persistence** parameter will be set to "After read" which is the same behavior – the synchronization state is saved right after the notes document is read.
 - If the **Delivery Mode** parameter is set to "Normal assured delivery" mode then a check for a valid **State Key Persistence** parameter is made. If such is not found then the value of the **State Key Persistence** parameter is set to "After read". If the parameter is found in the configuration then the original value of it is used.
- IDS Changelog Connector:
The CRAM-MD5 option is no longer available in 7.0; you must manually choose another authentication mechanism.
In version 6.2 of Tivoli Directory Server the BEREncoder and BERDecoder classes have been moved from the `com.ibm.asn1` package to the `com.ibm.ldap.bp.asn1` package. Starting from Tivoli Directory Integrator v7.0 custom user solutions that directly use the old classes (`com.ibm.asn1.BEREncoder` and `com.ibm.asn1.BERDecoder`) need be updated to reflect this change.
- EventHandlers:
EventHandlers are removed in 7.0 (they have been deprecated already in Tivoli Directory Integrator v6.0). For more information, look here.
- EMF XMLToSDO and EMF SDOToXML Function Components:
These are deprecated in 7.0. Consider other functionality in the future.
- DSMLv2 Parser:
This applies only to 6.0.
The "dsml.request" and "dsml.response" attributes have been removed. These attributes used to provide the raw request and response objects from the ITIM DSMLv2 library. If you have old configurations using any of these Attributes, you to edit your old configurations so that these Attributes are no longer used. All the data available through the raw request and response objects is also available through the other Attributes delivered by the DSMLv2 Parser.
- ITIM Agent Connector:

If you have used the ITIM Agent Connector in a previous version of Tivoli Directory Integrator, you may have to change the way you configure SSL connections. The ITIM Agent Connector in IBM Tivoli Directory Integrator 7.0 uses JSSE (Java based keystore or truststore) for SSL authentication, and this requires that you configure the SSL related certificate details in the `global.properties` or `solution.properties` file; instead of mentioning the certificate name in the old ITIM Agent Connector's "CA Certificate File" Parameter. These are the steps involved:

1. Import the ITIM Agent's certificate that was previously mentioned in the "CA Certificate File" parameter into the Tivoli Directory Integrator truststore with for example *keytool* (Ikeyman can be used too):

```
keytool -import -file servercertificate.der -keystore tim.jks
```

In this example the truststore is stored in the file `tim.jks`.

2. Configure this truststore in the "server authentication" section of the `global.properties` or `solution.properties` file:

```
## server authentication
```

```
javax.net.ssl.trustStore=serverapi\tim.jks  
{protect}-javax.net.ssl.trustStorePassword=administrator  
javax.net.ssl.trustStoreType=jks
```

Now, the ITIM Agent Connector uses the same JSSE-based secure communications architecture as the rest of Tivoli Directory Integrator.

If you already have a truststore file configured in `global.properties` / `solution.properties`, then import the certificate into that store instead of creating a new one.

- XML Parser:

The pre-v7.0 XML Parser has been renamed and is now called the Simple XML Parser; the v7.0 XML Parser is a new parser with more functionality, especially regarding hierarchical objects. Config files created under earlier versions of Tivoli Directory Integrator referring to the XML Parser will, when imported into v7.0, refer to the Simple XML Parser (as the class name has not changed). If you want to use the new XML Parser instead, you will need to change that in your AssemblyLines and/or Connectors. In order to have the new XML Parser behave like the old one did you must set the both Entry Tag and Value Tag parameters to the values used in the Simple XML Parser.

Note: The Simple XML Parser exports a script variable named "xmldom", which is not exported by the new XML Parser. The new XML Parser represents the deeper hierarchy with the Entry itself. Any logic that relies on the "xmldom" variable and cannot be reworked to make use of the hierarchical structure provided by the Entry class, must not migrate to the new XML Parser.

Customized scripts:

If you have customized any of the Directory Integrator scripts (for example, adding items to the `PATH` or the `LD_LIBRARY_PATH` environment variables in the startup scripts - `ibmdisrv`, `ibmditk`), you should apply these customizations to the corresponding scripts of the new version.

Previous versions of Tivoli Directory Integrator used the `(MY)CLASSPATH` variable in these scripts; the current version has the required path information built in and does not require this variable anymore. If you had tailored the aforementioned scripts before to include some libraries of your own, you do not have to do anything with the `CLASSPATH` variable; just make sure your library is in the correct place (typically in the `jars/` directory) so it is found by Tivoli Directory Integrator. Alternatively, use the `com.ibm.di.loader.userjars` property in `global.properties` to point to your own directory to be included in the loader path. In Tivoli Directory Integrator 7.0, the property may specify several directories or jar files, separated by the Java Property "path.separator", which is ":" on Linux and ";" on Windows. The `TDILoader` for jar files searches directories recursively for files that contain classes and resources. Only files with a ".zip" or ".jar" extension are searched.

Added or replaced JAR files in the installation:

If you have added JAR files to the installation, you should copy them to the new version too.

IBM Tivoli Directory Integrator now requires and includes a J2SE version 1.5 compliant JVM (J2SE version 1.5 SR9). If you have developed your own code in Java, linked this code against the JVM libraries and integrated this with your IBM Tivoli Directory Integrator solution, you might have to recompile and re-link your code.

If you have overwritten any of the original JAR files of the installation (for example, putting any required MQ jars in *TDI_install_dir/jars/3rdparty/IBM*), you should do the same with the new version.

A 64-bit Java Runtime Environment (JRE) is used now on Windows x86-64, Linux x86-64 and Linux s390. Compared to a 32-bit JRE, some performance degradation has been observed in some scenarios; you can still use the Windows x86-32 or Linux x86-32 installer for non-password plugin activities if you believe you will have potential issues with performance degradation.

If you do use the 64-bit JRE, you need to be aware that 64-bit shared libraries will be needed for any custom component (connector, parser, FC) that depends on JNI.

Password Synchronizer configurations:

- Windows Password Synchronizer

Follow the steps described in "Migration from previous installations" in the Windows Password Synchronizer section of the *IBM Tivoli Directory Integrator V7.0 Password Synchronization Plug-ins Guide*.

- Other Password Synchronizers

There are no specific migration steps. Uninstall the old version, install v7.0 and configure it to suit your needs.

Backing up important data

It makes sense to backup each of the files of the original installation that you have modified in some way: property files, keystores, scripts, jars, and so forth.

Generally it is up to you to decide which items are important and have to be backed up. Here is an overview:

Files backed up by the Installer

Upgrade from version 6.0 to 7.0: If the Server feature is being upgraded the following files will be backed up:

```
TDI_install_dir\global.properties to TDI_install_dir\etc\global.properties.v60
TDI_install_dir\serverapi\testadmin.jks to TDI_install_dir\serverapi\testadmin.jks.v60
TDI_install_dir\serverapi\testadmin.der to TDI_install_dir\serverapi\testadmin.der.v60
TDI_install_dir\serverapi\registry.enc to TDI_install_dir\serverapi\registry.enc.v60
TDI_install_dir\serverapi\registry.txt to TDI_install_dir\serverapi\registry.txt.v60
TDI_install_dir\idisrv.sth to TDI_install_dir\idisrv.sth.v60
TDI_install_dir\testserver.jks to TDI_install_dir\testserver.jks.v60
TDI_install_dir\testserver.der to TDI_install_dir\testserver.der.v60
```

Upgrade from version 6.1.x to 7.0: If the Server feature is being migrated, the following files will be backed up: (The new suffix will be .v61 or .v611 depending on the previous version.)

```
TDI_install_dir\etc\global.properties to TDI_install_dir\etc\global.properties.v61x
TDI_install_dir\serverapi\testadmin.jks to TDI_install_dir\serverapi\testadmin.jks.v61x
TDI_install_dir\serverapi\testadmin.der to TDI_install_dir\serverapi\testadmin.der.v61x
TDI_install_dir\serverapi\registry.enc to TDI_install_dir\serverapi\registry.enc.v61x
TDI_install_dir\serverapi\registry.txt to TDI_install_dir\serverapi\registry.txt.v61x
```

TDI_install_dir\idisrv.sth to *TDI_install_dir\idisrv.sth.v61x*
TDI_install_dir\testserver.jks to *TDI_install_dir\testserver.jks.v61x*
TDI_install_dir\testserver.der to *TDI_install_dir\testserver.der.v61x*
TDI_install_dir\etc\reconnect.rules to *TDI_install_dir\etc\reconnect.rules.v61x*
TDI_install_dir\etc\derby.properties to *TDI_install_dir\etc\derby.properties.v61x*
TDI_install_dir\etc\jlog.properties to *TDI_install_dir\etc\jlog.properties.v61x*
TDI_install_dir\etc\log4j.properties to *TDI_install_dir\etc\log4j.properties.v61x*
TDI_install_dir\etc\tdisrvctl-log4j.properties to *TDI_install_dir\etc\tdisrvctl-log4j.properties.v61x*
TDI_install_dir\etc\act-jlog.properties to *TDI_install_dir\etc\act-jlog.properties.v611* (TDI 6.1.1 only)

Backup tools

These tools are used for backup/restore by the installer. They can also be used for manual migration.

backupamc/restoreamc

Use to backup/restore AMC configuration files.

backupamcdb/restoreamcdb

Use to backup/restore the AMC database.

backupam/restream

Use to backup/restore the Action Manager (AM) database.

See “AMC and AM Command line utilities” on page 220 for more details.

Manual backup

Manual backup means copy the file to some dedicated backup folder. Conversely, restore means copy the file from the dedicated backup folder to its original location.

Note that in some cases you have to consider dependencies between files. You need to backup a group of interdependent files as a whole. Such groups of files are:

Derby database files

To backup a database, backup the whole folder that contains the database files. For example copy the *TDI_install_dir/TDISysStore* folder to backup the default System Store database or copy the *TDI_install_dir/bin/amc/tdiamcdb* folder to backup the default AMC database.

MQe queue manager files

Backup the whole folder of the MQe queue manager. For example copy the *TDI_install_dir/MQePWStore* folder to backup the default System Queue.

CE workspace files

Backup the whole workspace folder.

Migrating AMC 7.0 configuration settings to another AMC deployment

The section explains the steps that are to be followed for migrating AMC configuration data to a new AMC deployment. These instructions are useful when migrating AMC 7.0 from ISC SE to ISC AE, ISC AE to ISC SE, or to create a mirror instance of AMC on another machine.

1. Backup all of the AMC configuration files and data:
 - a. Stop the AMC that you are migrating configuration data from using the *stop_tdiamc.bat*(sh) script. This script stops the server on which AMC is deployed.
 - b. Execute the *backupamc.bat* (.sh) script specifying the directory in which the AMC configuration files need to be backed up.
2. Migrate all of the AMC configuration data to the new AMC instance:
 - a. If the ISC you are migrating the AMC configuration to, is residing on another machine you will need to get the AMC backup directory copied to the new machine.
 - b. Ensure the AMC being migrated to is stopped. See step 1a for information on stopping AMC.

- c. Execute the `restoreamc.bat (.sh)` script specifying the directory you have the AMC configuration information backed up to. This command will place the AMC configuration files at the correct location in the AMC you are migrating to. This completes the migration process.

Notes:

1. The instructions are for migrating AMC 7.0 to any other ISC deployment.
2. The instructions assume that you already have AMC deployed already in both the system being migrated from and the system being migrated to using the TDI installer. The systems can be either running AMC in ISC SE or ISC AE.
3. If you are trying to migrate the AMC configuration to another AMC deployment on the same machine and you want all of the AMC commands shipped with your TDI deployment to use that AMC from that point on, you will need to update the configured ISC location in the AMC command line utility. This can be done using the following command: `setISCHome.bat(sh)`. The command takes as a parameter the location of the ISC installation directory, which is the installation location of Websphere for ISC AE and the location of the embedded Web platform for ISC SE. The command needs to be executed between steps 1 and 2 mentioned above.

Converting from EventHandlers to corresponding AssemblyLines

EventHandlers are removed in Tivoli Directory Integrator 7.0. Therefore, in order to replace the deleted functionality in old solutions, you need to migrate your EventHandler configurations to Server/Iterator or Changelog Connector configurations.

For each EventHandler a corresponding AssemblyLine must be created. Then a Server/Iterator Connector corresponding to the EventHandler must be inserted into the AssemblyLine "Feeds" section. Then the Connector parameters must be set – this is specific for each EventHandler/Connector pair, but generally the Connector parameters must be set to the same values as the corresponding EventHandler parameters (which usually have the same names).

Any processing configured in the EventHandler must be re-implemented in the AssemblyLine "Flow" section.

The functionality of the "enabled" EventHandler parameter (otherwise known as "Auto-start service") is also available for AssemblyLines. If you want your AssemblyLine to be started right after the Tivoli Directory Integrator Server is started, go to the **Log & Settings** section in your workspace in the Config Editor and add your AssemblyLine.

In general an EventHandler executes some piece of logic when a certain event occurs. "Event" has a different meaning for each EventHandler. For the HTTP EventHandler an "event" is an HTTP request. For the IBM Directory Server EventHandler an "event" is a change notification that comes from an IBM LDAP Directory.

Below are some general guidelines on migrating certain parts of a typical EventHandler. They are divided based on the titles of the UI tabs for an EventHandler in the pre-7.0 Config Editor:

Hooks

The "Prolog" hook of an EventHandler corresponds to the "Prolog – After Init" hook of an AssemblyLine. This hook is invoked for each incoming "event".

The "Epilog" hook of an EventHandler corresponds to the "Epilog – After Close" AssemblyLine hook. This hook is invoked once after each incoming "event" is processed.

In both the "Prolog" and "Epilog" EventHandler hooks the "event" Entry is accessible under the names "conn" and "event". However in the AssemblyLine hooks you should modify your script to use "work" instead of "conn" or "event".

The "Shutdown Request" hook of an EventHandler corresponds the "Shutdown Request" AssemblyLine hook.

Action Map

The Action Map of an EventHandler defines what actions should be taken when an "event" arrives. You should build the same actions into the logic of the AssemblyLine that you are preparing as a replacement of the EventHandler.

For example if the Action Map prescribed that a custom script should be executed if Attribute "x" of the event equals "3", then you could add an "IF" component to the AssemblyLine that checks for Attribute "x" being equal to 3 and executes a Script Component.

Logging

If you have configured custom log appenders for the EventHandler, you should configure the same appenders in the logging settings of the AssemblyLine(s) that you are preparing as a replacement for the EventHandler.

Config

These configuration parameters are specific to each EventHandler. See the subsections below for instructions on how to migrate them. The subsections are named after the corresponding Connectors.

TCP Server Connector

You must do the following to reproduce an old EventHandler's configuration into a new Connector's configuration:

1. Create a new AssemblyLine and insert the TCP Server Connector in it.
2. Set the **tcp.port** and **debug** Connector parameters to the values of the corresponding EventHandler parameters.
3. Set the **useSSL** and **requireClientAuth** Connector parameters to false (unchecked in the Config Editor).

Mailbox Connector

There is no need to migrate existing configurations that use the Tivoli Directory Integrator v6.0 Mailbox Connector, because the Tivoli Directory Integrator v7.0 Mailbox Connector is compatible with the Tivoli Directory Integrator v6.0 Mailbox Connector.

Configurations that use the Mailbox EventHandler, however, need to be migrated by following these steps:

1. Create a new AssemblyLine and insert the Mailbox Connector in it.
2. Copy the contents of the **mailServer** EventHandler parameter to the Connector parameter with the same name.
3. Set the **mailProtocol** Connector parameter to the value of the EventHandler parameter with the same name.
4. Copy the contents of the **mailUser** and **mailPassword** EventHandler parameters to the Mailbox Connector parameters with the same names.
5. Copy the contents of the **mailFolder** EventHandler parameter to the Connector parameter with the same name.
6. Copy the contents of the **pollInterval** EventHandler parameter to the Connector parameter with the same name.
7. If the enabled EventHandler parameter is true, add your AssemblyLine to the "Config -> AutoStart" folder in the Config Editor; thus the Tivoli Directory Integrator server will start your AssemblyLine on startup.
8. If the **debug** EventHandler parameter is true, set the Connector parameter with the same name to true.

JMX Connector

Existing configurations that use the Tivoli Directory Integrator v6.0 JMX EventHandler can be transformed into Tivoli Directory Integrator v7.0 configurations that use the JMX Connector in the following way:

1. Create a new AssemblyLine and insert the JMX Connector in it.
2. Copy the contents of the **eventTypes** JMX EventHandler parameter to the JMX Connector parameter with the same name.
3. Select "local" for the **mode** Connector parameter.
4. Leave the **url** Connector parameter blank.
5. Set the **allMBeans** Connector parameter to true.
6. Leave the **mBeanTypes** Connector parameter blank.

SNMP Server Connector

The Tivoli Directory Integrator v7.0 SNMP Server Connector provides all features of the Tivoli Directory Integrator v6.0 SNMP EventHandler except support for single-threaded mode. The Tivoli Directory Integrator v7.0 SNMP Server Connector works in multi-threaded mode only. If you need to migrate an existing Tivoli Directory Integrator v6.0 configuration using the SNMP EventHandler to a Tivoli Directory Integrator v7.0 configuration, which uses an AssemblyLine with the SNMP Server Connector, you need to do the following:

1. Create a new AssemblyLine.
2. Insert into the AssemblyLine an instance of the SNMP Server Connector.
3. Set the **udp.port** Connector parameter to the value this parameter has in your SNMP EventHandler configuration.
4. Set the **snmp.community** Connector parameter to the value this parameter has in your SNMP EventHandler configuration.
5. If your SNMP EventHandler used to be configured to be "Auto-started" by the TDI Server, add your new AssemblyLine to the "Config -> AutoStart" folder of the Config Editor.

IBM Directory Server Changelog Connector

Existing configuration that use the IBM Directory Server EventHandler can be migrated to use the IBM Directory Server Changelog Connector as follows:

1. Set the following Connector parameters to the values of the EventHandler parameters with the same names: **ldapUrl**, **ldapUsername**, **ldapPassword**, **ldapAuthenticationMethod**, **ldapUseSSL**, **ldapSearchBase**.
2. Leave the **jndiExtraProviderParams** Connector parameter empty.
3. Set the **iteratorStateKey** Connector parameter to some unique identifier, one that has no corresponding state saved in the System Store.
4. Set the **nsChangenum** Connector parameter to the next change number that the EventHandler would process. The last change number that the EventHandler has processed is normally stored in an external properties file, referenced by its **ldapChangeNumberFileName** parameter.
5. Set the **stateKeyPersistence** Connector parameter to "After read" (the EventHandler writes the last received change number to its file backend after it reads a changelog entry and before it dispatches it for processing).
6. Set the **mergeMode** Connector parameter to "Merge changelog and changed data". This will ensure that the changelog attributes (changenumber, targetdn, ...) appear as attributes of the Entry.
7. Set the **useNotifications** Connector parameter to true.
8. Set the **batchRetrieval** Connector parameter to false.

Note: As opposed to the EventHandler, the Connector does not let you select a part of the directory tree, for whose notifications it will listen – it subscribes for changes in the whole directory tree (the

Connector does not have equivalents of the **ldapEventBase** and **ldapSearchScope** EventHandler parameters). If this is critical for you, you can implement some custom filtering in your solution to overcome this limitation of the Connector.

HTTP Server Connector

A configuration that uses the HTTP EventHandler can be migrated to use the HTTP Server Connector like this:

1. Set the **tcpPort** Connector parameter to the value of the **Port** parameter of the EventHandler.
2. Leave the **backlog** Connector parameter empty.
3. Set the **contentType** Connector parameter to "text/html".
4. Set the **tcpDataAsProperties** Connector parameter to true (the EventHandler always returns the TCP information as properties).
5. Set the **headersAsProperties** Connector parameter to the value of the **headersAsProperties** of the EventHandler.
6. Set the **httpAuth** Connector parameter to true, if the EventHandler uses HTTP basic authentication (that is if it has a configured authentication Connector).
7. If the EventHandler uses HTTP basic authentication, set the **authRealm** Connector parameter to the value of the **authrealm** EventHandler parameter. If the **authrealm** EventHandler parameter is missing or empty, set the **authRealm** Connector parameter to "IBM-Directory-Integrator".
8. Set the **authConnector** Connector parameter to the value of the **AuthConnector** parameter of the EventHandler.
9. Set the **useSSL** Connector parameter to the value of the **useSSL** parameter of the EventHandler.
10. Set the **needClientAuth** Connector to false (the EventHandler does not support SSL client authentication).
11. Set the **msgChunked** Connector parameter to false (the EventHandler does not support chunking of HTTP responses).

LDAP Server Connector

A configuration that uses the LDAP Server EventHandler can be migrated to use the LDAP Server Connector as follows:

1. Set the **ldapPort** Connector parameter to the value of the **tcp.port** parameter of the EventHandler.
2. Leave the **backlog** Connector parameter empty.
3. Set the **ldapUseSSL** Connector parameter to the value of the **ldapUseSSL** parameter of the EventHandler.
4. Set the **charset** Connector parameter to the value of the **charset** parameter of the EventHandler.
5. Set the **ldapBinaryAttributes** Connector parameter to the value of the **binary** parameter of the EventHandler.

Netscape/iPlanet/Sun Directory Changelog Connector

The LDAP EventHandler catches notifications about changes in a directory tree. The EventHandler does not use a changelog, so it receives only real-time notifications. The Netscape/iPlanet/Sun Directory Changelog Connector offers basically the same functionality when run in real-time delivery mode. There are a few differences though:

The Connector does not have equivalents for the **ldapSearchFilter** and **ldapSearchScope** EventHandler parameters. To achieve the same functionality as in the EventHandler, you should implement some custom filtering that limits the set of received notifications.

The schema of the returned data differs between the Connector and the EventHandler. The Connector applies delta tagging to each Entry it returns, while the EventHandler provides the type of the change in the "ldap.operation" property. For details on the schema consult the documentation of each component.

Once the considerations above are resolved, you can migrate an existing configuration with the LDAP EventHandler to use the Netscape/iPlanet/Sun Directory Changelog Connector like this:

1. Set the following Connector parameters to the values of the EventHandler parameters with the same names: **ldapUrl**, **ldapUsername**, **ldapPassword**, **ldapAuthenticationMethod**, **ldapUseSSL**, **ldapSearchBase**.
2. Leave the **jndiExtraProviderParams** Connector parameter empty.
3. Set the **deliveryMode** Connector parameter to "Realtime" (the EventHandler does not use a changelog, it only catches real-time notifications).
4. Set the **mergeMode** Connector parameter to "Return only changed data" (no changelog is utilized in real-time delivery mode by the Connector).

Active Directory Change Detection Connector

The migration from the AD Changelog EventHandler to the Active Directory Change Detection Connector is straight forward in the most aspects since the EventHandler itself has incorporated the older version this connector – Active Directory Changelog Connector in order to obtain changes from the AD.

Similar to the EventHandler the corresponding Connector can also be interrupted any time during the synchronization process, in that case it will store its state in the User Property Store. Both the EventHandler and the Connector rely on the uSNChanged mechanism in this process, by storing the USN number in the property store. They also offer an API for retrieving the current USN synchronization values. The difference is that the EventHandler `getUSNvalues` method returns an Entry with Attributes:

```
START_USN  
END_USN  
CURRENT_USN_CREATED  
CURRENT_USN_CHANGEDT
```

whereas the Connector returns the current synchronization value as *long*.

Another difference is that the AD EventHandler initializes internally an LDAP Connector in order to block and receive change notifications. This behavior can also be simulated in the ADCD Connector by enabling the **useNotifications** parameter.

The following steps should be performed in order to migrate from an EventHandler-based solution to Connector-based one:

1. Create a new AssemblyLine with an instance of an Active Directory Change Detection Connector in Iterator mode.
2. Set the **ldapUrl**, **ldapUsername**, **ldapPassword** and **ldapAuthenticationMethod** to the values these connection parameters have in the EventHandler configuration.
3. Specify whether SSL connection is used according to the value in the old configuration.
4. Copy the content of the **ldapSearchBase** EH parameter to the same in the Connector configuration
5. Copy the content of the **persistentParameterName** EH parameter to the **persistentStateKey** Connector parameter.
6. Set the parameter **useNotifications** to true.
7. Set the **startAt** parameter according to the value in EH.
8. Leave the other Connector parameters as they are.
9. Transfer any logic in the Action Map section of the EventHandler to be invoked from the new AL.

z/OS LDAP Changelog Connector

The migration from the z/OS LDAP Changelog EventHandler to the z/OS LDAP Changelog Connector is facilitated by the fact that this changelog connector is explicitly developed to replace the EventHandler. Both the EventHandler and the Connector rely on a poll mechanism for extracting the changelog information since they do not support the Unsolicited Event Notification.

The main difference consists in storing the state key. The EventHandler uses a plain file passed as parameter - **ldapChangeNumberFileName** in its configuration to keep track of the last used changenumber, whereas the Connector takes advantage of the "Iterator State Key" mechanism in Tivoli Directory Integrator. Therefore, in order to migrate a solution from the z/OS LDAP Changelog EventHandler to the z/OS LDAP Changelog Connector, the value of the changenumber in the **ldapChangeNumberFileName**-specified file must be passed to the Connector for example with script before the initial start of the iteration, like this script snippet:

```
Com.ibm.di.store.StoreFactory.getDefaultPropertyStore().setProperty("changenumber",  
    new Long(changeNumber));
```

where **changenumber** is the name of the Iterator State Key and *changeNumber* is the value obtained from the store file used by the EventHandler.

An alternative to this solution is to pass the *changeNumber* value to the **nsChangenumber** parameter in the configuration of the Connector.

The following steps should be performed in order to migrate from an EH-based solution to Connector-based one:

1. Create a new AssemblyLine with an instance of the z/OS LDAP Changelog Connector in Iterator mode.
2. Set the **ldapUrl**, **ldapUsername**, **ldapPassword** and **ldapAuthenticationMethod** to the values these connection parameters have in the EventHandler configuration.
3. Specify whether SSL connection is used according to the value in the old configuration.
4. Copy the content of the **ldapSearchBase** EH parameter to the same in the Connector configuration.
5. Apply one of the solutions described above to supply the Iterator State Key with the last "changenumber" value used for synchronization by the EH.
6. Copy the content of the **pollInterval** EH parameter to the **nsSleepInterval** Connector parameter.
7. Leave the other Connector parameters as they are.
8. Transfer any logic in the Action Map section of the EventHandler to be invoked from the new AL.

DSMLv2SOAPServerConnector

The migration from the DSMLv2 EventHandler to the DSML v2 SOAP Server Connector requires rework of the AssemblyLines that are previously used with the EventHandler, so that they can be integrated in the solution with the DSMLv2 SOAP Server Connector. This is because the core architecture as changed; now a single AssemblyLine processes all operations. Therefore, all the old AssemblyLines logic responsible for handling the different types of DSMLv2 operations should be incorporated into the new AssemblyLine containing the DSMLv2 Soap Server Connector, or should be invoked using a AssemblyLine Connector. For this purpose branching components can be used in order to separate the logic for the specific DSMLv2 operations (available in the `dsm1.operation` Attribute).

The migration of a configuration with the DSMLv2 EventHandler to a similar one with the DSMLv2 SOAP Server Connector consists of the following steps:

1. Create a new AssemblyLine with an instance of a DSMLv2 SOAP Server Connector in Server mode.
2. Copy the content of the EH **port** parameter to the **dsm1Port** Connector parameter.
3. Set the **authRealm**, **useSSL**, **binaryAttributes** and **msgChunked** to the values these connection parameters have in the EventHandler configuration.

4. Create a branch component for each of the DSMLv2 operations listed as parameters in the EH configuration and apply in the branches the logic implemented in the corresponding old AssemblyLine, either by transferring there the appropriate AL components or by invoking the old AL itself using an AssemblyLine connector. In both cases the naming context will no longer be needed.
5. Copy the content of the twoEH **WaySSL** parameter to the **needClientAuth** Connector parameter.
6. The EH Attribute **headerAsProperties** cannot be passed to the Connector, since the HTTP parser it initializes internally is configured to always set this value to "false". Therefore, in case the solution accesses headers as properties, it should be modified to use Attributes for this purpose (`getAttribute()` instead of `getProperty()`).
7. For compliance, the **soapbinding** Connector Attribute should be set to "false" since the DSMLv2 parser internally used by the EH does not take advantage of it.
8. In case an **authConnector** is specified in the configuration of the DSMLv2 EventHandler, then the HTTP basic authentication of the Connector must be enabled and the appropriate logic must be implemented in the "After Accepting connection" hook (for example, initialize the authenticator Connector and call its `lookup()` method using an Entry with Attributes "username" and "password" as search criteria. Similar to the EventHandler the authentication is to be considered successful in case an Entry is returned).
9. The **indentoutput** parameter of the DSMLv2 parser internally used by the Connector cannot be set in contrast to the one used by the EH.

Migrating BTree tables and BTree Connector to System Store

The BTree Connector is deprecated, and is now only provided as an unsupported example. Therefore, you might decide to move the way your Delta information is maintained from the old Btree objects to Delta Tables in the System Store. The best strategy for doing this is engineering a situation where your Delta information is empty (for example, establishing a new baseline) and then switch from the Btree objects to the System Store Delta Tables. Note that the parameter that used to hold the filename of the Btree objects now indicates a table name in a database, so some editing of this value might be required.

Changing a solution to use the System Store Connector instead of the BTree Connector for storing Tivoli Directory Integrator Entries is straight forward since both connectors follow the same logic when specifying Key Attribute Name and Selection Mode attributes. The only difference is that instead of the underlying BTree database, the System Store Connector has to use predefined a database (for example the embedded Derby database) and specify a table to store into.

Storing other Java objects using the System Store Connector differs significantly from storing them with BTree and will require more elaborate transformation. The following solution, which puts Java objects in the underlying BTree database, cannot be directly applied to the System Store Connector, since it does not provide direct access to the backend database:

```
scripts var bt = system.getConnector("btreedb");
bt.initialize (null); var db = bt.getDatabase();
db.insert ("my key", new java.lang.String("my value"));
var value = db.search ("my key"); value = value + " - modified";
db.replace ("my key", value);
```

Instead of this the standard methods (`put()`, `find()` and `modify()`) from the Connector API can be used, but the object should be first wrapped into an Entry object, which subsequently can be stored in the System Store.

Migrating Cloudscape database to Derby

IBM Tivoli Directory Integrator 7.0 uses Derby v10.3 as its bundled database, used by default by the System Store. You will need to migrate your existing Cloudscape or Derby databases (created using previous versions of Tivoli Directory Integrator) to be able to use Tivoli Directory Integrator 7.0. Derby v10.3 drivers that are shipped with Tivoli Directory Integrator 7.0 cannot be used to communicate with older versions of Cloudscape.

For details, and information on differences between Cloudscape/Derby v10 and its prior versions, refer to the following web page:<http://publibfp.boulder.ibm.com/epubs/html/c1894710.html>.

Notable differences that have an immediate impact are as follows:

- The long varbinary data type is no longer supported. Instead, BLOB datatype has been introduced (making Derby compatible with DB2®). For this reason, all SQL Statements that made use of long varbinary datatype must now be modified to use BLOB.
- JDBC Java package names have changed from `com.ibm.db2j.*` in previous releases to `org.apache.derby.*` in Derby v10.
- The JDBC URL for Derby (embedded/network mode access) v10 is different from Cloudscape v5.1. Hence the JDBC properties mentioned in `global.properties` / `solution.properties` have also been modified for the current version of Tivoli Directory Integrator.

Table 8. JDBC URL differences

Connection type	Cloudscape v5.1	Derby v10
Embedded Derby / Cloudscape	<code>jdbc:db2j:</code>	<code>jdbc:derby:</code>
DB2 JDBC Universal Database Driver (Network mode)	<code>jdbc:db2j:net</code>	<code>jdbc:derby:net</code> (Not recommended to use)
DerbyClient Driver	-	<code>jdbc:derby</code> (Recommended)

Fortunately, the Derby team have provided a migration utility that migrates a Cloudscape v5.1 database to a new Derby v10 database. It migrates all the tables and their corresponding data into a newly generated Derby v10 database. It modifies all tables with varbinary datatype to BLOB datatype, hence making the migration process quite painless.

This utility is bundled with Tivoli Directory Integrator 7.0, in the `TDI_install_dir/tools/CSMigration` folder, along with a wrapper script that invokes the migration tool, called `migrateCS.bat(sh)`. To migrate a Cloudscape 5.1 System Store Database created using TDI v 6.0 to Derby v10, you have to invoke the migrate script in the following manner:

```
migrateCS [Path_of_CloudscapeV51_Database] [Path_of_new_DerbyV10_Database]
```

You may need to give some thought to the location of the new Derby database. In Tivoli Directory Integrator v6.0 and v6.1.x, the System Store database often was located in the installation directory of Tivoli Directory Integrator; this is an unfortunate location for many reasons. For Tivoli Directory Integrator 7.0 we strongly recommend you use a Solution Directory, away from the installation directory.

Besides migration of data, you also need to modify your `global.properties` / `solution.properties` files (using the migration tool or manually) to incorporate the new JDBC URL parameters.

Migrating global and solution properties files using the migration tool

Use the `tdimiggb1` tool located in the `TDI_install_dir/bin` directory to migrate any `global.properties` file starting with Tivoli Directory Integrator 6.x to 7.0. The filename is `tdimiggb1.bat` on Windows and `tdimiggb1.sh` on UNIX/LINUX. Use the `tdimiggb1-4log4j.properties` file to control logging for `tdimiggb1.bat(sh)`.

The usage if the command is as follows:

```
tdimiggb1 -f propfile [-b backfile] [-n newfile] [-v] [-?]
```

where:

- f propfile - The name of the file to migrate
- b backfile - Backup the original file with the specified name
- n newfile - Name to give the file that is migrated
- s dir - Working directory where the solution directory is located.
- v - Enable verbose mode
- ? - Prints the usage statement

During the installation of Tivoli Directory Integrator, the installer backs up the existing `global.properties` file; and then calls this command, in order to migrate the `global.properties`.

The migration tool tries to migrate a `global.properties` file (or `solution.properties` file if required) up to the latest Tivoli Directory Integrator version. The tool (`tdimiggb1`) makes no assumptions about which release the `global.properties` files starts from and can handle `global.properties` files starting at Tivoli Directory Integrator version 6.0. The tool also tries to apply all migration changes unless a particular migration step is specifically declared inappropriate for migration by the migration tool. For these cases, perform the migration steps manually.

The activities of the migration tool are broken down into stages. In sequence, the tool:

1. Checks whether you have to migrate your Derby (Cloudscape) database (Tivoli Directory Integrator 6.0 migration).
2. Performs all of the migration actions in the following order:
 - a. Delete actions.
 - b. Add actions.
 - c. Derby (Cloudscape) migration file changes (only if necessary and only for Tivoli Directory Integrator 6.0 migrations).
 - d. Migration modify actions.
3. Calls the Derby (Cloudscape) migration tool `migrateCS` to migrate the database up to the current Derby version (only for Tivoli Directory Integrator 6.0 migrations).

For each action set (migration modify actions for example), the migration tool tries to perform the migration actions starting from the earliest release to the latest release. For migration from Tivoli Directory Integrator 6.0, the caller must separately invoke the Derby (Cloudscape) migration tool to migrate the database up to the current Derby version. The `tdimiggb1` tool only makes the required Derby (Cloudscape) modifications to the properties file itself.

4. Uses `log4j` logging APIs for logging error messages.

The `log4j` configuration file is specified in the startup script (the `bat` or `sh`) file. The command uses a file called `tdimiggb1-log4j.properties` to set up the `log4j` logging. The command changes directory to the solution directory and therefore uses the `tdimiggb1-log4j.properties` file in the solution directory if the Tivoli Directory Integrator installation directory is not specified.

Chapter 6. Security and TDI

Introduction

Security features are found throughout IBM Tivoli Directory Integrator (Tivoli Directory Integrator). Some features secure access into remote systems from Tivoli Directory Integrator, others protect access into Tivoli Directory Integrator from remote systems, and yet others provide mechanisms to secure data, such as user credentials into remote systems.

Many of the features described in this chapter are not necessary when running Tivoli Directory Integrator in a stand-alone mode in a secured environment. However, the features come in handy when other systems must communicate with Tivoli Directory Integrator, such as through the remote Web Admin Console (AMC) management tool or the Tivoli Directory Integrator Remote Server API. Furthermore, if multiple people have access to the Tivoli Directory Integrator server it could be necessary to protect access to confidential data, as well as maintain the integrity of the integration rules that Tivoli Directory Integrator executes.

This chapter explains the following features:

1. "Secure Sockets Layer (SSL) Support"
2. "Remote Server API" on page 72
3. "TDI Server Instance Security" on page 92
4. "Miscellaneous Config File features" on page 102
5. "Web Admin Console Security" on page 108
6. "Summary of configuration files and properties dealing with security" on page 105
7. "Miscellaneous security aspects" on page 108

This guide does not describe all the security capabilities of the individual Tivoli Directory Integrator components. Some common elements are described in "Miscellaneous security aspects" on page 108, however for individual elements of security configuration in the individual TDI components, consult the *IBM Tivoli Directory Integrator V7.0 Reference Guide*.

Secure Sockets Layer (SSL) Support

SSL is an important foundation for many Tivoli Directory Integrator security features. You need a working-level knowledge of SSL in order to fully exploit the capabilities in TDI.

The following Connectors support SSL with properly configured IBM Tivoli Directory Integrator Servers:

- Connectors
 - JMS Connector
 - LDAP Connector
 - LDAP Server Connector
 - Sun Directory Change Detection Connector
 - IBM Directory Server ChangeLog Connector
 - Active Directory Change Detection Connector
 - Lotus® Notes® Connector
 - Axis Easy Web service Server Connector
 - Web service Receiver Server Connector
 - DSMLv2 SOAP Server Connector

SSL provides for encryption and authentication of network traffic between two remote communicating parties. Most production deployments of TDI make use of SSL. That is why SSL support is one of the major security features of TDI. More information on SSL as well as information on using SSL in Java programs from a development point of view can be found at <http://java.sun.com/j2se/1.5.0/docs/guide/security/jsse/JSSERefGuide.html>

TDI can be used as a client, as a server or as both at the same time. Configuring TDI for SSL when used as a client is different from configuring TDI when used as a server. That is why this section has been divided in two sub-sections – “Server SSL configuration of TDI components” and “Client SSL configuration of TDI components” on page 65.

Server SSL configuration of TDI components

When a TDI component is used as a server (for example a Server mode Connector) SSL mandates that a keystore to be used by IBM Tivoli Directory Integrator must be defined. For information on keystores and truststores, see the guide at <http://java.sun.com/j2se/1.5.0/docs/guide/security/jsse/JSSERefGuide.html>. The following steps are required to enable SSL support for IBM Tivoli Directory Integrator as a server:

Note: RMI is enabled by default in the Tivoli Directory Integrator server. Properties for server authentication carry the default keystore property values.

1. If you don't have a java (jks) keystore file already in IBM Tivoli Directory Integrator create a keystore file using *keytool* (found in *TDI_Install_dir/jvm/jre/bin*, or *TDI_Install_dir/jvm/bin* depending on your platform). If you don't have a personal key to be used in IBM Tivoli Directory Integrator get one from a Certificate Authority or create a self-signed key.
2. If the certificate in the IBM Tivoli Directory Integrator is a self-signed certificate, export the certificate.
3. If the IBM Tivoli Directory Integrator certificate is a self-signed certificate, using a key tool, import the exported IBM Tivoli Directory Integrator certificate to the keystore file in the client as a root authority certificate.
4. Edit *TDI_Install_dir/etc/global.properties* file for the keystore file location, keystore file password and keystore file type. In the current release, we support jks-type only.

```
## client authentication
javax.net.ssl.keyStore=serverapi\testadmin.jks
{protect}-javax.net.ssl.keyStorePassword=administrator
javax.net.ssl.keyStoreType=jks
```
5. Enable SSL for the clients (for example, using https in the Web browser).
6. Restart IBM Tivoli Directory Integrator

Notes:

1. The TDI server does not manage the keystores/truststores. All that the TDI server provides to the TDI components in terms of keystore support is the *global.properties* or *solution.properties* files, in which the standard Java keystore/truststore properties can be specified.
2. A TDI component can choose to use the default configured keystore/truststore in *global.properties* or *solution.properties*, or it can choose to implement its own handling of SSL sockets (for example implementing a custom *SSLServerSocket* Java class) so that it can use keystores/truststores different from the default.
3. If TDI needs to use both a client and a server certificate only the default certificate configured in *global.properties* or *solution.properties* is used, then this must be the same certificate. An alternative would be to write a custom implementation of the *SSLSocket* or the *SSLServerSocket* Java class and make it use a certificate different from the default.
4. See section “Certificates for the TDI Web service Suite” on page 109 for specifics on the certificates for TDI Web service components.

Client SSL configuration of TDI components

When a TDI component is used as a client (for example the LDAP Connector) SSL mandates that a truststore to be used by TDI must be defined. For information on keystores and truststores, see the guide at <http://java.sun.com/j2se/1.5.0/docs/guide/security/jsse/JSSERefGuide.html>.

The following steps are required to enable SSL support for IBM Tivoli Directory Integrator as a client:

1. Configure a server (such as IBM Tivoli Directory Server) to enable SSL.
2. If the certificate in the server is a self-signed certificate, export the certificate.
3. If you don't have a Java (jks) keystore file already, create a keystore file using *keytool* (found in *root_directory/jvm/jre/bin*, or *root_directory/jvm/bin*, depending on your platform) for IBM Tivoli Directory Integrator.
4. If the server certificate is a self-signed certificate, import the server certificate to the IBM Tivoli Directory Integrator keystore file as a root authority certificate using *keytool*.
5. Edit *root_directory/etc/global.properties* file for the keystore file location, keystore file password and keystore file type. Tivoli Directory Integrator 7.0 supports Java keystore (jks) type only.

Note: These four lines (comments starting with #) are no longer needed for client and server authentication to the Tivoli Directory Integrator server. Stores that belong to Tivoli Directory Integrator are set up to be used by default. This is part of enabling Remote Method Invocation (RMI) by default.

```
# Keystore file information for the server TDI authentication.  
# It is used to provide the public key of the TDI to the SSL enabled client.  
# javax.net.ssl.keyStore=D:\test\clientStore.jks  
# javax.net.ssl.keyStorePassword=secret  
# javax.net.ssl.keyStoreType=jks
```

6. Enable SSL for the Connectors.
7. Restart IBM Tivoli Directory Integrator.

Note: TDI truststore and keystore do not play any part in SSL configuration for the Domino Change Detection connector. See section “Lotus Domino SSL specifics” on page 108 for more information.

SSL client authentication

If a TDI component is used as a client and the server with which it communicates requires SSL client authentication, then apart from a truststore, Tivoli Directory Integrator needs a keystore as well. In this case the keystore can be defined just like it is defined when TDI is used a server – see the section “Server SSL configuration of TDI components” on page 64.

Note: Client TDI components which support SSL client authentication do not normally need a “SSL client authentication” check box as do TDI server components. All such a client TDI component needs in order to prove its identity to the server is to have its keystore generated and configured in *global.properties* or *solution.properties*. If the server requires an SSL client certificate then the client SSL library automatically sends the client’s certificate from the keystore configured in *global.properties* or *solution.properties*.

IBM Tivoli Directory Integrator and Microsoft Active Directory SSL configuration

Do the following steps to configure SSL for IBM Tivoli Directory Integrator and Microsoft Active Directory:

1. Install Certificate Services on the Windows Server and an Enterprise Certificate Authority in the Active Directory Domain. Details are available at <http://windowsitpro.com/article/articleid/14923/how-do-i-install-an-enterprise-certificate-authority.html>. Make sure you install an **Enterprise Certificate Authority**.

2. Start the Certificate Server Service. This creates a virtual directory in Internet Information Service (IIS) that enables you to distribute certificates.
3. Create a Security (Group) Policy to direct Domain Controllers to get an SSL certificate from the Certificate Authority (CA).
 - a. Open the **Active Directory Users and Computers Administrative** tool.
 - b. Right-click, under the domain, **Domain Controllers**.
 - c. Select **Properties**.
 - d. Select the **Group Policy** tab, and click to edit the **Default Domain Controllers Policy**.
 - e. Go to **Computer Configuration**→**Windows Settings**→**Security Settings**→**Public Key Policies**.
 - f. Right click **Automatic Certificate Request Settings**.
 - g. Select **New**.
 - h. Select **Automatic Certificate Request**.
 - i. Run the wizard. Select the **Certificate Template for a Domain Controller**.
 - j. Select your **Enterprise Certificate Authority** as the CA. Selecting a third-party CA works as well.
 - k. Complete the wizard.

Note: All Domain Controllers automatically request a certificate from the CA, and support LDAP using SSL on port 636.

4. Retrieve the Certificate Authority Certificate to the computer on which you installed IBM Tivoli Directory Integrator.

Note: You must install IIS before installing the certificate server.

- a. Open a Web browser on the computer on which you installed IBM Tivoli Directory Integrator.
- b. Go to `http://server_name/certsrv/` (where *server_name* is the name of the Windows 2000 server). You are asked to log in.
- c. Select the task **Retrieve the CA certificate or certificate revocation list**.
- d. Click **Next**.
The next page automatically highlights the CA certificate.
- e. Click **Download CA certificate**
A new download window opens.
- f. Save the file to the hard drive.
5. Create a certificate store using keytool. Use keytool.exe to create the certificate store and import the CA certificate into this store.

Note: keytool.exe is found in *root_directory/jvm/jre/bin*, or *root_directory/jvm/bin*, depending on your platform.

Use the following command:

```
jvm\jre\bin\keytool -import -file
certnew.cer -keystore keystore_name.jks
-storepass password-alias keyalias_name
```

For example, assume the following values:

```
Keystorename = idi.jks
Password = secret
Keyalias name = AD_CA
```

The command looks like this script:

```
C:\Program Files\IBM\TDI\V7.0\jvm\jre\bin\keytool -import
-file certnew.cer -keystore idi.jks -storepass secret -alias AD_CA
```

To verify the contents of your keystore, type the following script:

```
C:\Program Files\IBM\TDI\V7.0\jvm\jre\bin\keytool
-list -keystore idi.jks -storepass secret
```

The following lines result:

```
Keystore type: jks
Keystore provider: SUN
```

Your keystore contains 1 entry:

```
ad_ca, Mon Nov 04 22:11:46 MST 2002, trustedCertEntry,
Certificate fingerprint (MD5): A0:2D:0E:4A:68:34:7F:A0:21:36:78:65:A7:1B:25:55
```

6. Configure IBM Tivoli Directory Integrator to use the keystore created in the previous step. Edit `root_directory/global.properties` file for the keystore file location, keystore file password and keystore file type. In the current release, only jks-type is supported.

```
#server authentication
#example
javax.net.ssl.trustStore=c:\test\idi.jks
javax.net.ssl.trustStorePassword=secret
javax.net.ssl.trustStoreType=jks
#client authentication
#example
javax.net.ssl.keyStore=c:\test\idi.jks
javax.net.ssl.keyStorePassword=secret
javax.net.ssl.keyStoreType=jks
```

7. Enable SSL for your LDAP connector.
 - a. Go to the LDAP Connector configuration window.
 - b. Change **LDAP URL** to port 636.
 - c. Check **Use SSL**.
8. Restart IBM Tivoli Directory Integrator.

Note: The Tivoli Directory Integrator Windows service wrapper permits you to start TDI as multiple service instances.

Summary of properties for enabling SSL and PKCS#11 support

You can configure SSL properties for server authentication, client authentication, and PKCS#11 support. See “Using cryptographic keys located on hardware devices” on page 135 on Public Key Cryptography Standards (PKCS).

Table 9. SSL Server Authentication

Property	Default value	Description
<code>javax.net.ssl.trustStore</code>	<code>serverapi\testadmin.jks</code>	Location of the truststore files.
<code>{protect}-javax.net.ssl.trustStorePassword</code>	<code>administrator(encrypted by default)</code>	truststore password.
<code>javax.net.ssl.trustStoreType</code>	<code>jks</code>	Type of the truststore.

Table 10. SSL Client Authentication

Property	Default value	Description
<code>javax.net.ssl.keyStore</code>	<code>serverapi\testadmin.jks</code>	keystore files location.
<code>{protect}-javax.net.ssl.keyStorePassword</code>	<code>administrator(encrypted by default)</code>	keystore password.
<code>javax.net.ssl.keyStoreType</code>	<code>jks</code>	Keystore type.

Table 11. PKCS#11 Support

Property	Default value	Description
com.ibm.di.pkcs11cfg	etc\pkcs11.cfg	Use this to specify the path of the configuration file required to initialize the IBM PKCS11 implementation provider. Added in TDI 7.0.
com.ibm.di.server.pkcs11	false	Use pkcs11 compliant crypto devices for ssl. Added in TDI 7.0.
com.ibm.di.server.pkcs11.library	none	Specify the path to the PKCS11client library. Added in TDI 7.0.
com.ibm.di.server.pkcs11.slot	none	Specify the slot number of the device.
{protect}- com.ibm.di.server.pkcs11.pass	none	Access the pkcs11 compliant crypto device using this password. Encrypted by default. Added in TDI 7.0.

Keystore and truststore management

In order to create, edit, export and overall manage keystores and truststores the "Ikeyman" GUI utility or the "keytool" command line utility can be used. The executable file, keytool.exe is found in *root_directory\jvm\jre\bin*, or *root_directory\jvm\bin*, depending on your platform.

Managing a CA-signed certificate using keytool

Normally the process of acquiring and using CA-signed certificates goes like this: First a key pair is generated. After that a certificate for the public key is requested from a Certification Authority. When the Certification Authority sends back the signed certificate, the certificate is imported into the appropriate truststore. Below is an example of this process using the *keytool* Java utility:

1. `keytool -genkey -dname cn=<server_ip_address> -keystore server.jks -storepass secret -keypass secret` This command creates a private/public key pair and stores it into the 'server.jks' keystore file.
2. `keytool -certreq -file myRequest.csr -keystore server.jks -storepass secret -keypass secret` This command creates a Certificate Signing Request in the 'myRequest.csr' file for the public key created in step 1. The created Certificate Signing Request now can be sent to a Certification Authority.
3. `keytool -import -trustcacerts -file server.cer -keystore mytruststore.jks -storepass secret` This command reads the certificate (public key) stored in the 'server.cer' certificate file (possibly the response of a Certificate Signing Request) and imports it into the 'mytruststore.jks' keystore file.

After these steps are executed, the keystore 'server.jks' contains the created public/private key pair. On the other hand, the 'mytruststore.jks' keystore represents a truststore which trusts the public key.

Creating a self-signed certificate using keytool

The standard Java utility *keytool* can be used to create self-signed certificates. The following commands can be entered at the command prompt in the JVM bin directory:

1. `keytool -genkey -dname cn=server_ip_address -validity 18263 -keystore server.jks -storepass secret -keypass secret`
This command creates a private/public key pair and stores it into the //server.jks keystore file.
2. `keytool -export -alias mykey -file server.cer -keystore server.jks -storepass secret`
This command extracts the public key (certificate) from the server.jks //keystore file and stores it into the server.cer certificate file.
3. `keytool -import -trustcacerts -file server.cer -keystore server.jks -storepass secret -alias mytrustedkey`
This command reads the certificate (public key) stored in the server.cer //certificate file and imports it into the server.jks keystore file.

Answer "yes" and press Enter to the "Certificate already exists in keystore under alias <mykey> Do you still want to add it? [no]:" question.

Note: The number 18263 is the validity period of the certificate in days (18263 days is roughly equal to 50 years).

After these steps are executed a server.jks keystore file is created which contains a key for the server. This file is also a truststore which contains the server public key, that is, trusts the server key. In this way this server.jks file can be used as both the server keystore and client truststore file.

Extending the validity of a certificate using keytool

Suppose you have a keystore called "teststore.jks" which includes a expired (or about to expire) self-signed certificate whose alias name is "mycertalias". The keystore has the associated private key in it. Assume that the password for the keystore as well as the private key is "mypassword". Now, if you want to extend the validity of this certificate by another 365 days, you can run the following command using keytool:

```
keytool -selfcert -v -alias mycertalias -validity 365 -keystore teststore.jks -storepass mypassword
```

The above operation generates a new self-signed certificate, that has the same DN, SIGALG, KEYS as the original certificate (mycertalias) but has a new SERIAL NUMBER and VALIDITY period.

Note: The generated new certificate automatically replaces the original one. So if you need the original one later for reference or for any reason, you must keep a copy of the original keystore before doing the certificate extension explained above.

Creating a self-signed certificate using Ikeyman

Ikeyman provides a graphical user interface for managing keystores and truststores. The Ikeyman tool can be launched from within the Configuration Editor by choosing **KeyManager** from the Tools menubar. In order to create a keystore which contains a self-signed certificate you have to follow these steps:

1. Start the Ikeyman GUI tool
2. From the **Key Database File** menu click **New....**
3. In the **New** window:
 - a. set the **Key database type** to **JKS** (the default)
 - b. set the **File Name** to **server.jks**
 - c. set the **Location** to the appropriate value
4. In the **Password Prompt** field:
 - a. enter the keystore password you have chosen
 - b. confirm the password value
 - c. click **OK**
5. From the **Create** menu, click **New Self-Signed Certificate....**
6. In the **Create New Self-Signed Certificate** window:
 - a. Set the **Key Label** to a name (label) that is used to identify the key and certificate in the database, for example, my self-signed certificate
 - b. Set the **Key Size**
 - c. Set the **Common Name** to the fully qualified host name of the server as the common name, for example, **www.myserver.com**
 - d. Set the **Organization** to your organization's name
 - e. Fill in any of the optional fields if you have to
 - f. Specify the **Country or region**
 - g. Specify the **Validity Period** in days.
 - h. Click **OK**
7. Click the **Extract Certificate...** button.

8. In the **Extract Certificate to a File** window:
 - a. Set the **Data type** to Base64-encoded ASCII data (the default)
 - b. Set the **Certificate file name** to `server.arm`
 - c. Set the "Location" to the appropriate value.
 - d. Click **OK**
9. From the drop down list select **Signer Certificates**
10. Click the **Add...** button.
11. In the **Add CA's Certificate from a File** window:
 - a. For the **Data Type** select **Base64-encoded ASCII data**
 - b. Set the **Certificate file name** to `server.arm`
 - c. Input the **Location**.
 - d. Click **OK**.


After these steps are executed a `server.jks` keystore file is created which contains a key for the server. This file is also a truststore which contains the server public key, that is, trusts the server key. In this way this `server.jks` file can be used as both the server keystore and client truststore file.

Exporting a key from a keystore to a PKCS#12 file using Ikeyman

1. Start the Ikeyman GUI tool
2. From the **Key Database File** menu click **Open....**
3. In the **Open** dialog select the path to the keystore. Make sure the **Key database type** is set properly.
4. In the **Key database content** section select **Personal Certificates**.
5. Now all keys available in the keystore should be displayed in the **Key database content** section.
6. Select one of them and click the **Export/Import...** button.
7. In the **Export/Import...** window:
 - a. Select the **Export Key** radio button.
 - b. Select **PKCS12** in the **Key file type** field.
 - c. Enter file name and location where the PKCS12 file is created.
 - d. Click **OK**.

Importing a key from a PKCS#12 file into a keystore using Ikeyman

1. Start the Ikeyman GUI tool
2. From the **Key Database File** menu click **Open...**
3. In the **Open** dialog select the path to the keystore. Make sure the **Key database type** is set properly.
4. In the **Key database content** section select **Personal Certificates**.
5. Now all keys available in the keystore should be displayed in the **Key database content** section.
6. Click the **Export/Import...** button.
7. In the **Export/Import...** window:
 - a. Select the **Import Key** radio button.
 - b. Select **PKCS12** in the **Key file type** field.
 - c. Enter file name and location of the PKCS12 file, which contains the key to be imported.
 - d. Click **OK**.



SSL example

In order to demonstrate how TDI can be configured for SSL when used as a server and also when used as a client, two examples are provided – one deploying the LDAP Server Connector and one deploying the LDAP Connector.

TDI component as a server

This example uses the LDAP Server Connector. The LDAP Server Connector listens for LDAP requests. When an LDAP request arrives the Connector parses the request and provides the request data to the hosting AssemblyLine. The AssemblyLine then processes the request and provides the data for the response to the LDAP Server Connector, so that it can build the LDAP response and send it back to the LDAP client. What follows is a step by step guide how to configure TDI for SSL when the LDAP Server Connector is used:

1. Obtain the server keystore either requesting it from a Certification Authority (CA) or creating a self-signed certificate as explained in either the "Creating a self-signed certificate using keytool" on page 68" section or the "Creating a self-signed certificate using Ikeyman" on page 69" section.
2. Set the keystore details in global.properties or solution.properties as described in the "Server SSL configuration of TDI components" on page 64" section.
3. Select **Use SSL** on the Connector GUI configuration window. You may need to expand the Advanced section to make the parameter visible.

Mode **Server** ☐ For-Each Loop

Initialize **at startup**

LDAP Port **389**

Comment

Detailed Log ☐

▼ **Advanced**

Connection Backlog

Use SSL ☐

Character Encoding **UTF-8**

Binary Attributes

TDI component as a client

This example uses the LDAP Connector. The LDAP Connector connects to an LDAP Server and sends an LDAP request. After the Server returns the LDAP response the LDAP Connector provides that response to the AssemblyLine for further processing. What follows is a step by step guide how to configure TDI for SSL when the LDAP Connector is used:

1. Generate the client truststore.
2. Import the LDAP server certificate into the client truststore.
3. Set the truststore details in `global.properties` or `solution.properties` as described in the "Client SSL configuration of TDI components" on page 65" section.

The following command line imports an existing certificate into a keystore (the keystore is created if not already existing):

```
keytool -import -trustcacerts -file myLDAPServerCert.cer -keystore myClientTruststore.jks -storepass myclientTruststorePassword -alias myTrustedLDAPServerAlias
```

This command line imports a the `myLDAPServerCert.cer` certificate under alias `myTrustedLDAPServerAlias` into the `myClientTruststore.jks` keystore. The password to access the keystore is `myclientTruststorePassword`.

Remote Server API

Introduction

This section does not cover securing an instance of a Tivoli Directory Integrator Server; this is discussed in "TDI Server Instance Security" on page 92. Instead, this section discusses how client applications can contact a server.

IBM Tivoli Directory Integrator supports the concept of a Remote 251 (also known as just "Server API"), where client tasks can invoke tasks on a remote Tivoli Directory Integrator Server by means of an access layer called 260.

Note: The "remote Server" could very well be running on the same computer as the client application, for example if you start up a Server instance on your local computer and then access it using the Remote API through the loopback address, 127.0.01. All concepts discussed below are still valid, even though the remote Server runs locally.

The Server API calls address the following areas:

- Getting Server information
- Getting information for components installed on the Server
- Reading and writing to configuration(s) loaded by the Server
- Loading new configurations into the Server
- Starting, querying and stopping AssemblyLines
- Cycling through AssemblyLines

Note: Increasing needs for remote server access for each running Tivoli Directory Integrator server have resulted in a change from local access by default to remote access by default. As of 7.0, the remote server API is enabled by default. Prior to 7.0, the server API was enabled only for local access by default, where local access means access from the same Java Runtime Environment (JRE). To ensure security, remote access requires SSL client authentication. SSL access using client authentication is provided with the sample keystore and truststore deployed with TDI.

The Server API itself is documented in the Tivoli Directory Integrator Java API documentation (*TDI_install_dir/docs/api*; you can launch a browser to display this documentation by selecting **Help -> Welcome -> JavaDocs** in the CE). The package of interest in this context is `com.ibm.di.api`. Also see the chapter called "Using the Server API" in *IBM Tivoli Directory Integrator V7.0 Reference Guide*.

The Configuration Editor uses the Remote API to talk to the server you use to test-run your solutions. If this Tivoli Directory Integrator server is running on the same machine, it is often called the "local development server". For setups where the deployment platform does not support the Configuration Editor (for example, z/OS or i5/OS), you can run the development server on the deployment server, and the Configuration Editor on a supported platform like Windows (this way of running we call the "Remote Configuration Editor"). This design provides a uniform interface for both remote and local Config files. For some aspects of the Configuration Editor talking to a remote deployment server, see "Using the Remote Configuration Editor" on page 104.

The Server API "Configuring the Server API" is configured through a set of server properties. These properties are specified in the `global.properties` configuration file of the TDI Server. Some of the properties, in turn, point to additional configuration files and keystore files.

The Server API provides a number of security-related features (which both TDI Solution-based clients as well other client applications have to deal with). There are three aspects to Server API Access Security:

1. "Server API SSL remote access" on page 76 (which secures the transport channel to a remote TDI Server),
2. "Server API authentication" on page 77 (which handles the client authentication to a TDI Server),
3. "Server API Authorization" on page 85 (which handles the client authorization to a TDI Server, that is, what the client is allowed to do once authenticated).

Configuring the Server API

The relevant properties are:

Property	Description
<code>api.on</code>	If set to true , the Server API is initialized on startup and can be used; otherwise the Server API is not initialized and cannot be used. All other properties whose names start with "api." are only taken into account if <code>api.on</code> is set to true .

Property	Description
<i>api.audit.on</i>	If set to true , the audit feature is turned on. If it is set to true , an audit entry is created at each audit point, even if the audit notifications are suppressed.
<i>api.config.folder</i>	<p>If you place configuration files in this folder, you can edit them through the Server API. At startup, the TDI server scans this folder for the Solution Names of the config files located there.</p> <p>Configuration files placed in other folders cannot be edited through the Server API.</p>
<i>api.config.timeout</i>	If specified, this property contains the serverapi config load time-out value in minutes. Default timeout is two minutes. Added in TDI 7.0.
<i>api.notification.suppress</i>	<p>If set to true, the suppress notification property is turned for each property set. The property contains a list of server notification types that are to be suppressed. Notifications of suppressed types are not propagated by the notifications framework. The notification types in the list are separated by spaces. You can include wildcards. For example:</p> <p><code>api.notification.suppress=di.al.* di.ci.start</code></p> <p>This example suppresses all AssemblyLine related notifications as well as notifications for starting a configuration instance. If the property is missing or is empty, no notifications are suppressed. Added in TDI 7.0.</p>
<i>api.user.registry</i>	If configured, specifies the Server User Registry file name.
<i>api.user.registry.encryption.on</i>	If set to true , the Server API decrypts the Server User Registry file on startup.
<i>api.remote.on</i>	If set to true (default), the remote RMI part of the Server API is initialized and can be used; otherwise, the remote RMI part of the Server API is not initialized and cannot be used. Added in TDI 7.0.
<i>api.remote.ssl.on</i>	If set to true , SSL with client and server authentication is used on RMI connections of the Server API and its JMX layer; and the Server API uses the Server certificate and private key (the one specified through the <code>api.keystore</code> and <code>api.key.alias</code> properties) for SSL connections. RMI clients must trust that certificate. If set to false , no SSL is used for client connections and no authentication and authorization is performed; connections are accepted from the local host and from hosts listed in the <code>api.remote.nonssl.hosts</code> property; if <code>api.remote.nonssl.hosts</code> is empty, only connections from the local host are accepted.
<i>api.remote.ssl.client.auth.on</i>	If set to true , the SSL client authentication for remote Server API is switched on.
<i>api.remote.nonssl.hosts</i>	If specified, shows a list of IP addresses to accept non-SSL connections from (host names are not accepted). Use space, comma or semicolon as a delimiter between IP addresses. This property is only taken into account when <code>api.remote.ssl.on</code> is set to false .
<i>api.remote.naming.port</i>	If specified, the port on which the RMI registry listens for requests.
<i>api.truststore</i>	If specified, the keystore file that contains the public certificates of all remote users of the Server API.
<i>{protect}-api.truststore.pass</i>	If specified, the password for the keystore file named through the <code>api.remote.server.truststore</code> property.
<i>api.jmx.on</i>	If set to true , the JMX layer of the Server API is initialized on startup and can be used; otherwise, the JMX layer is not initialized and cannot be used.
<i>api.jmx.remote.on</i>	If set to true , the remote JMX interface (as defined by JSR160) is initialized and can be used; otherwise the remote JMX interface is not initialized and cannot be used.
<i>api.config.folder</i>	If set to <code>TDI_root/configs</code> , only the configuration files placed in this folder can be edited through the Server API.

Property	Description
<i>api.config.lock.timeout</i>	If set to 0 , there is no timeout.
<i>api.custom.method.invoke.on</i>	The ability to use <code>invokeCustom()</code> methods can be turned on or off (the default is false, or off). If the value of this property is set to true then users can use these methods.
<i>api.custom.method.invoke.allowed.classes</i>	If specified, gives the list of classes that can be invoked directly by the Server API methods for custom method invocation (<code>Session.invokeCustom(...)</code>). This property is only taken into account if <code>api.custom.method.invoke.on</code> is set to true . The classes in this list must be separated by a space, a comma or a semicolon.
<i>api.custom.authentication.</i>	If specified, points to a JavaScript text file that contains custom authentication code. To enable the built-in LDAP/JAAS Authentication mechanism, set this property to <code>[ldap]/[jaas]</code> .
<i>com.ibm.di.server.id</i>	If specified, contains the server ID. Assign a unique value for each server from the set of servers that are running on the same IP and Port.
<i>api.config.load.timeout</i>	If specified, contains the serverapi config load timeout value in minutes. Added in Tivoli Directory Integrator 7.0.
<i>api.notification.suppress</i>	<p>If specified, gives a list of server notification types that you want to suppress. Notifications of suppressed types are not propagated by the notifications framework. Notification types in the list are separated by spaces. You can include wildcards.</p> <p>Example:</p> <pre>api.notification.suppress=di.al.* di.ci.start</pre> <p>The above example suppresses all AssemblyLine related notifications as well as notifications for starting a configuration instance. If the property is missing or is empty, no notifications are suppressed. Added in Tivoli Directory Integrator 7.0.</p>

Note: The Java system properties that the Server API uses for its configuration are the same, regardless of whether the client is a Java program or a different instance of the Tivoli Directory Integrator Server. What should be noted though is that the way these Java system properties are set might be different. In Tivoli Directory Integrator these properties are normally set by editing the `global.properties` or `solution.properties` files, whereas in a Java program they can be specified either at the command line using the `-D` Java command line switch or by using Java code within the Java program using the `java.lang.System.setProperty(key,value)` standard Java method.

Remote Server API access on a Virtual Private Network

When the Remote Server API is accessed from a client on a Virtual Private Network (VPN), the VPN assigns an IP address to the Server API client computer. This VPN-assigned IP address needs to be specified in an RMI Java system property. If the Server API client is the Remote Configuration Editor, then this property can be set in `global.properties` or `solution.properties` by adding the following line to the properties files:

```
java.rmi.server.hostname=<IP_address>
```

Where *IP_address* is the VPN-assigned IP address.

If the Server API client is a custom Java program, then this property can be set from within the Java code in the following way:

```
java.lang.System.setProperty("java.rmi.server.hostname", "IP_Address");
```

where *IP_address* is the VPN-assigned IP address.

Note that the RMI Java system property needs to be set before any Server API related RMI code.

Server API access options

The Server API can be used in a variety of ways:

- Access the Server API from the Remote Configuration Editor through a network connection
- Access the Server API from Tivoli Directory Integrator components running in a remote Tivoli Directory Integrator server (remote Server API access). Examples of such components are:
 - System Queue Connector
 - Server Notifications Connectorand so on.
- Access the Server API from within the same Java Virtual Machine of the TDI Server (local Server API access); in this case the Server API can be reached from JavaScript in hooks or from the Script Component in addition to the options above.
- Access the Server API from non-TDI Java applications. For this to work:
 - Java 5.0 or higher is required on the client side.
 - The following jar files must be included in the CLASSPATH of the remote side:
 - jars/common/diserverapi.jar
 - jars/common/diserverapirmi.jar
 - jars/3rdparty/others/log4j-1.2.15.jar
 - jars/common/miconfig.jar
 - jars/common/miserver.jar
 - jars/common/mmconfig.jar
 - jars/common/tdiresource.jar
 - jars/3rdparty/IBM/icu4j_4_0.jar
 - jars/3rdparty/IBM/ITLMTToolkit.jar
 - jars/3rdparty/IBM/jlog.jar

You can copy these jar files from the Tivoli Directory Integrator installation.

- If custom non-TDI objects are used in the solution being implemented with the Server API (for example as Attribute values of an Entry that is transferred over the wire) the corresponding Java classes have to be available on the client side as well. These classes must be serializable and they have to be included in the CLASSPATH of the client JVM.

Server API SSL remote access

The Server API provides two sets of interfaces – local and remote. It is only the remote interfaces that can use SSL. The local interfaces do not use SSL as the access is within the boundaries of the Java Virtual Machine. TDI can act as a server, as a client; as well as both as a client and as a server in a Server API access scenario. When SSL is used with the Server API, then a keystore and a truststore must be configured. There are two options for configuring these. Which of them is used depends on whether the Java System property `api.client.ssl.custom.properties.on` exists and on its value.

Using Server API specific SSL properties

When the Java System property `api.client.ssl.custom.properties.on` is set to "true", then SSL is configured through the following TDI Server API-specific Java System properties:

- **api.client.keystore** – specifies the keystore file containing the client certificate
- **api.client.keystore.pass** – specifies the password of the keystore file specified by `api.client.keystore`
- **api.client.key.pass** – specifies the password of the private key stored in the keystore file contained in `api.client.keystore`; if this property is missing, the password specified by **api.client.keystore.pass** is used instead.
- **api.truststore** – specifies the keystore file containing the TDI Server public certificate.

- **api.truststore.pass** – specifies the password for the keystore file specified by api.truststore.

Use the Server API specific SSL properties when your client application is using the standard Java SSL properties. The standard Java SSL properties are properties used to configure another SSL channel used by the same application.

You can specify these properties as JVM arguments on the command line, for example:

```
java MyTDIServerAPIClientApp
-Dapi.client.ssl.custom.properties.on=true
-Dapi.truststore=C:\TDI\serverapi\testadmin.jks
-Dapi.truststore.pass=administrator
-Dapi.client.keystore=C:\TDI\serverapi\testadmin.jks
-Dapi.client.keystore.pass=administrator
```

This example refers to the sample "testadmin.jks" keystore file shipped with TDI. Note that it contains both the client private key and also the public key of the TDI Server, so we use it both as a keystore and truststore.

You can specify these properties in global.properties or solution.properties when the client is a Tivoli Directory Integrator server.

Using the standard SSL Java System properties

When the Java System property api.client.ssl.custom.properties.on is missing or when it is set to "false", then the standard JSSE system properties are used for configuring the SSL channel. Follow the standard JSSE procedure for configuring the keystore and truststore used by the client application.

You can specify these properties as JVM arguments on the command line, for example:

```
java MyTDIServerAPIClientApp
-Djavax.net.ssl.keyStore=C:\TDI\serverapi\testadmin.jks
-Djavax.net.ssl.keyStorePassword=administrator
-Djavax.net.ssl.trustStore=C:\TDI\serverapi\testadmin.jks
-Djavax.net.ssl.trustStorePassword=administrator
```

Also these properties can be specified in global.properties or solution.properties when the client is a TDI server.

Server API authentication

Server API authentication is usually referred to in the context of a remote Server API client establishing a Server API session. This scenario represents the substance of the Server API authentication logic as the Server API provides several different kinds of client authentication. But before diving into the different authentication mechanisms let us discuss the scenario in which a local client establishes a local Server API session.

Local client session

A local client session is a session established by a client which runs in the same Java Virtual Machine as the TDI server. Examples of such sessions are local sessions for access to the local Server API established from JavaScript code in hooks or in a Script component, from Connectors and Function Components which are executed as part of an AssemblyLine which runs in the same TDI server, and so on. When a local client establishes a local Server API session, the client has two options:

- Do not provide a username and password pair – in this case the local Server API session is established and the client is authorized as having the "admin" role. For more information about Server API roles, see "Server API Authorization" on page 85.
- Provide a username and password pair – in this case the Server API session is established only after the "username" supplied in the username and password pair is authorized according to the Server API Authorization logic described in the "Server API Authorization" on page 85 section. This option would normally be used when a specific user ID is needed for authentication – for demos, prototyping, and so on.

Remote client session

A remote client session is a session established by a client which does not run in the same Java Virtual Machine as the TDI server. Examples of such sessions are sessions for access to a remote Server API established from the Configuration Editor, or a Java application wishing to access a TDI Server. For access of this kind there are the following methods of authentication to the TDI Server:

JAAS authentication

The Java Authentication and Authorization Service (JAAS) is supported as an authentication module for Tivoli Directory Integrator Server APIs. JAAS is a set of APIs that enables services to authenticate and enforce access controls upon users. The JAAS authentication is facilitated by the Tivoli Directory Integrator Server API. No changes are required on the TDI Server API clients such as CLI and AMC in order to use the JAAS authentication module.

In order to use JAAS authentication, you must configure the appropriate properties in `global.properties` or `solution.properties` and the JAAS Logon should be installed.

SSL-based authentication

This is the only authentication mechanism available in TDI 6.0. SSL-based authentication is based on a two-stage verification of the client's credentials.

1. First the TDI server verifies that a client (represented by its SSL certificate) has the right to access the TDI server by checking whether the client's SSL certificate is contained in the TDI server's truststore, that is, checks whether the TDI server trusts this client. Checking whether the client's certificate is contained in the server's truststore is part of the SSL handshake sequence.

Attention: A client certificate example, corresponding to the Server certificate example in file `testserver.jks` is provided in file `serverapi/testadmin.jks`; the certificate's password is "administrator". As with all default security parameters you should not rely upon these and generate your own client/server certificates and specify these in the properties files. See "Certificates for the TDI Web service Suite" on page 109.

The truststore is kept in the file indicated by the `api.truststore` property.

2. If the truststore check is successful then the server verifies that the client SSL certificate distinguished name (DN) matches a user ID in the "Server API User Registry" on page 87. If the client certificate's DN does not match any of the user IDs in the Server API User registry file the connection request from the client is denied. This second step could be regarded as part of the authorization sequence as well.

The SSL-based authentication mechanism can be switched off in Tivoli Directory Integrator 7.0. An additional property is available in the TDI Server configuration file `global.properties` or `solution.properties`: `api.remote.ssl.client.auth.on`. When this property is set to "true", the TDI Server requires client authentication within the SSL handshake (the TDI 6.0 mechanism for SSL-based authentication). SSL client authentication for Tivoli Directory Integrator Server API does not depend on whether a username and password pair is supplied. This means that if no username and password pair is supplied, the TDI 6.0 mechanism for SSL-based authentication is used. And if a username and password pair is supplied then the client still needs to send its SSL certificate for authentication, but the User ID for authentication (and at a later step authorization) is taken from the username supplied.

When `api.remote.ssl.client.auth.on` is set to "false", SSL-based authentication cannot be used. When the property is not specified a value of "false" is assumed.

Username/password based authentication

This mechanism requires a client to supply a username and password on the opening of his Server API connection to the Tivoli Directory Integrator server. In order to configure this authentication method an authentication hook is used.

Authentication hook: This hook allows the provision of custom JavaScript code that performs username and password based authentication. This hook allows bundlers/deployers to write customized JavaScript code, which given a username and password pair determines whether the authentication should succeed or not.

The property allowing for this custom JavaScript authentication is specified in the TDI Server configuration file `global.properties` or `solution.properties`: **`api.custom.authentication`**. The `api.custom.authentication` property points to a JavaScript text file on the disk that contains custom authentication code. If this property is not specified then the TDI 6.0 SSL-based authentication mechanism is used. When the `api.custom.authentication` property is specified, the JavaScript code contained in the specified file is executed for each username and password based authentication request.

The authentication script has access to the predefined script object `userdata`. This object provides the following two public members:

- **`userdata.username`** – contains the name of the user requesting authentication
- **`userdata.password`** – contains the password provided by the user

The script is free to perform whatever checks and authentication actions it needs. It returns whether the authentication is successful through the **`ret`** object:

- set **`ret.auth = true`** to specify that the authentication is successful
- set **`ret.auth = false`** to specify that the authentication is not successful; in this case the authentication script can provide additional information for why the authentication failed through the **`ret.errordescr`** attribute (for example `ret.errordescr = "Invalid user name"`) and **`ret.errorcode`** (for example `ret.errorcode = 1`).

The description and error code fields is provided by the `AuthenticationException` thrown by the `ServerAPI` on unsuccessful authentication.

The authentication script has access to the main script object. It can be used for logging custom messages in the Tivoli Directory Integrator Server log file (for example `main.logmsg("Authentication failed for user : " + userdata.username)`).

An example authentication hook: An example authentication hook JavaScript file is available (in `TDI_install_dir/examples`) in order to demonstrate what the JavaScript of an authentication hook could look like. This example JavaScript can also be used as the basis of real-world TDI authentication hooks. The example JavaScript demonstrates how an authentication hook can use an LDAP server (Tivoli Directory Server, Active Directory, and so on) for authenticating client requests.

The JavaScript file is named `"ldap_auth.js"` and is installed in the `examples/auth_ldap` TDI Server folder. To deploy this sample LDAP authentication mechanism users can copy that file to the TDI solution folder and specify `api.custom.authentication=ldap_auth.js` in `global.properties` or `solution.properties`. The JavaScript code in `"ldap_auth.js"` tries to bind to an LDAP Server with the specified username and password. If the bind operation is successful, the script indicates a successful authentication, otherwise the authentication is rejected. The details for connecting to the LDAP Server like the server URL are specified in the `"ldap_auth.js"` script – this means that users have to edit this file and set the proper connection parameters before using the script. Here is the sample `"ldap_auth.js"` script:

```
env = new Packages.java.util.Hashtable();
env.put("java.naming.factory.initial", "com.sun.jndi.ldap.LdapCtxFactory");
env.put("java.naming.provider.url", "ldap://192.168.113.54:389");
env.put("java.naming.security.principal", userdata.username);
env.put("java.naming.security.credentials", userdata.password);
env.put(Packages.java.naming.Context.SECURITY_AUTHENTICATION, "simple");
```

```
main.logmsg("Authentication request for user: " + userdata.username);
```

```
try
{
    mCtx = new Packages.java.naming.directory.InitialDirContext(env);
```

```

    ret.auth = true;
}
catch(e)
{
    ret.auth = false;
    ret.errordescr = e.toString();
    // ret.errorcode = "49";
}

```

LDAP Authentication support

The TDI Server API provides support for LDAP Authentication. This allows you to leverage your existing LDAP infrastructures that already hold User IDs and Passwords.

LDAP Authentication Configuration: In order to use LDAP authentication the appropriate properties must be configured in `global.properties` or `solution.properties`. The list of these properties along with their descriptions follows:

api.custom.authentication

This is the same property used for username and password authentication. For more information on username and password authentication see the "Username/password based authentication" section. This property points to a JavaScript text file on the disk that contains custom authentication code. The user may not specify this property, in which case he can only use the TDI 6.0 SSL-based authentication mechanism. The Tivoli Directory Integrator 7.0 username and password authentication does not work. Set this property to "[ldap]" to enable the Tivoli Directory Integrator 7.0 built-in LDAP Authentication mechanism, like this:
`api.custom.authentication=[ldap]` All properties starting with "*api.custom.authentication.ldap.*" are only be taken into account when *api.custom.authentication* is set to *[ldap]*.

api.custom.authentication.ldap.critical

This parameter specifies the Server API behavior when the LDAP Authentication module cannot be initialized on startup. If this parameter is set to "true" the Server API initialization fails and the Server API is not started.

If this parameter is missing or is set to "false" the Server API logs the LDAP Authentication initialization error but the Server API is started. An attempt to initialize the LDAP Authentication module is made on each authentication request received by the Server API until the LDAP Authentication module is initialized.

api.custom.authentication.ldap.hostname

The LDAP Server hostname. If LDAP custom authentication is used, this is a required property.

api.custom.authentication.ldap.port

The LDAP Server port number. For example, 389 for non-SSL or 636 for SSL. If LDAP custom authentication is used, this is a required property.

api.custom.authentication.ldap.ssl

Specifies whether SSL is used to communicate with the LDAP Server. When set to "true" SSL is used, otherwise SSL is not used.

api.custom.authentication.ldap.searchbase

Specifies the LDAP directory location where user searches is preformed. When this property is not specified user searches is not performed.

api.custom.authentication.ldap.admindn

Specifies an LDAP Server administrator distinguished name that is used for user searches. When this property is not specified anonymous bind is used for user searches.

api.custom.authentication.ldap.adminpassword

Password for the LDAP Server administrator distinguished name.

api.custom.authentication.ldap.userattribute

Specifies the user id attribute to be used in searches. When this property is not specified user searches are not performed. An example setting of this property would be:
api.custom.authentication.ldap.userattribute=cn

If a required property is missing an exception is thrown on initialization.

If the value of either **api.custom.authentication.ldap.searchbase** or **api.custom.authentication.ldap.userattribute** is missing no search context is initialized and no searches are performed during the actual user authentication. (No search means that the bind to the LDAP Server is attempted directly with the username and password provided for authentication.)

When **api.custom.authentication.ldap.adminidn** is provided a search context is created using "simple" authentication. If an error occurs during the search context initialization, the initialization of the LDAP Authentication module fails and an exception is thrown.

When **api.custom.authentication.ldap.adminidn** is not provided a JNDI search context is created using JNDI "anonymous" bind.

Note: If the search context cannot be initialized using **api.custom.authentication.ldap.adminidn**, authentication fails directly – no anonymous bind is attempted.

LDAP Authentication Logic: On each attempt to authenticate a user the LDAP Authentication module is passed the username and the password for the user to be authenticated. The following authentication paths are possible:

- Both **api.custom.authentication.ldap.searchbase** and **api.custom.authentication.ldap.userattribute** properties are specified :
 - If the username given for authentication ends with the value of the **api.custom.authentication.ldap.searchbase** property it is assumed that a full distinguished name is provided and no user search is performed. A bind to the LDAP Server is attempted directly with the username and password provided for authentication. If the bind succeeds the authentication is considered successful, otherwise the authentication is considered failed.
 - If the username does not end with the value of the **api.custom.authentication.ldap.searchbase** property, a search with a subtree search scope is executed against the search context created on initialization. The search query used is "(<LDAPUserIDAttribute>=<username>)" where *LDAPUserIDAttribute* is the value of the **api.custom.authentication.ldap.userattribute** property and *username* is the username given for authentication. If exactly one search result is returned, a bind to the LDAP Server is performed with the distinguished name of the returned entry and the password provided for authentication. The authentication succeeds only if the bind to the LDAP Server is successful. In all other cases it is considered that the authentication has failed. If multiple search results are returned, authentication fails.
- At least one of **api.custom.authentication.ldap.searchbase** or **api.custom.authentication.ldap.userattribute** properties is not specified.

In this case no searches are performed and a bind to the LDAP Server is attempted directly with the username and password provided for authentication. If the bind succeeds the authentication is considered successful, otherwise it is considered that the authentication failed.

LDAP Group Support: To ease administration, Tivoli Directory Integrator allows permissions to be configured for groups the same way as they are configured for users. You can set permissions in the User Registry using groups exactly the same syntax as you would for a user. The fact is that the User Registry does not care whether a security entity is a group or a user. The distinction between users and groups is drawn during the authentication process.

Group membership is configured in the LDAP directory, against which Tivoli Directory Integrator authenticates users. If a user is a member of some LDAP group, all permissions for that group are automatically inherited by the user when the user is authenticated. Group support is disabled by default, so you must turn it on.

The system properties that are related to LDAP group support are:

api.custom.authentication.ldap.groupsupport

This is an optional property – a boolean flag. If this property is missing, the default value "false" is used. Specifies whether group membership is resolved when authenticating users. If the group membership is resolved, it is taken into account during authorization.

api.custom.authentication.ldap.usermembershipattribute

This property is required only if **api.custom.authentication.ldap.groupsupport** is set to true. Specifies the name of the attribute of a user in LDAP that contains a list of the groups of which the user is a member.

api.custom.authentication.ldap.usermembershipattributecontent

This property is required only if **api.custom.authentication.ldap.groupsupport** is set to true. Specifies how groups are named in the membership attribute of a user. For example, if the user's membership attribute contains values that correspond to the "objectSID" attributes of groups, set this property to "objectSID". If the user's membership attribute contains distinguished names of groups, then set this property to "dn".

api.custom.authentication.ldap.groupnameattribute

This property is required only if **api.custom.authentication.ldap.groupsupport** is set to true. Specifies the name of a group's attribute in LDAP which corresponds to the way the group is named in the TDI User Registry. For example, if LDAP groups are addressed in the Tivoli Directory Integrator registry by their common name, then set this property to "cn". If the User Registry contains the distinguished names of the groups, then set this property to "dn".

api.custom.authentication.ldap.groupsearchbase

This property is required only if **api.custom.authentication.ldap.groupsupport** is set to true. Represents the LDAP directory context, where groups are searched.

api.custom.authentication.ldap.binaryattributes

This is an optional property – it represents a list of space-separated attribute names. Specifies attributes which have non-string syntax.

Active Directory example:

This example shows how to configure group support to work with an **Active Directory** server:

```
api.custom.authentication.ldap.groupsupport=true
api.custom.authentication.ldap.usermembershipattribute=tokenGroups
api.custom.authentication.ldap.usermembershipattributecontent=objectSID
api.custom.authentication.ldap.groupnameattribute=sAMAccountName
api.custom.authentication.ldap.groupsearchbase=DC=mytestadserver,DC=com
api.custom.authentication.ldap.binaryattributes=objectSID tokenGroups
```

The 'tokenGroups' attribute is a calculated attribute that exists for all users in Active Directory.

It contains a collection of the Security Identifiers (SIDs) for all security groups that the user is a member of.

This collection contains only security groups (distribution groups, used for e-mail, are not included) and it contains all security groups including nested and primary groups.

The Security Identifiers are binary attributes so they must be set in the **api.custom.authentication.ldap.binaryattributes** property.

In the above example, groups are named by their "sAMAccountName" LDAP attribute in the TDI User Registry.

Tivoli Directory Server example:

This example shows how to configure group support to work with **Tivoli Directory Server**:

```
api.custom.authentication.ldap.groupsupport=true
api.custom.authentication.ldap.usermembershipattribute=ibm-allGroups
api.custom.authentication.ldap.usermembershipattributecontent=dn
api.custom.authentication.ldap.groupnameattribute=dn
api.custom.authentication.ldap.groupsearchbase=ou=mytestou,c=mytestcountry
```

For a given user entry, the "ibm-allGroups" operational attribute enumerates all static, dynamic and nested groups, to which that user has membership.

Notes:

1. Tivoli Directory Integrator determines group membership by directly examining the LDAP user entry (as opposed to indirectly determining membership by scanning through all groups). For this approach to work correctly, the user entry must have an attribute that enumerates the groups, of which the user is a member. The group support works only with LDAP Servers that do support such a membership attribute on each user entry.
2. If you modify the group membership of a user, this does not affect existing Server API sessions. It is, however, reflected in sessions established after the modification.
3. Group support is currently provided only for LDAP authentication. There is no group support for JAAS authentication or authentication with custom JavaScript.
4. When SSL client authentication is enabled in the Server API, clients that do not specify a username are to be authenticated and authorized based on the owner of the SSL client certificate. If LDAP authentication with group support is also enabled (along with the SSL client authentication), group membership is resolved for the owner of the SSL client certificate.

Host based authentication

Host based authentication is used, when SSL is turned off by specifying *api.remote.ssl.on=false* in *global.properties* or *solution.properties* files. Host based authentication is configured using the *api.remote.nonssl.hosts* property. This property specifies the list of host IP addresses from which remote Server API clients can use the Server API without specifying a username and password.

The syntax of this list of hosts is: a list of IP addresses (host names are not accepted); use a space, a comma or a semicolon as a delimiter between IP addresses. An example value of this property would be:

```
api.remote.nonssl.hosts=192.168.111.222, 192.168.112.158
```

When a client using host based authentication is successfully authenticated, then the client is granted admin authorization authority. That is why adding IP addresses to this list must be done with great care. It is not advisable to use host based authentication in production environment because of its security issues. Host based authentication would normally be used while developing a solution or when doing a demo.

Summary of Server API Authentication options

The following authentication options are available:

SSL-based authentication (the mechanism available in TDI 6.0)

Only works when *api.remote.ssl.client.auth.on=true* (you also need *api.on=true*, *api.remote.on=true*, *api.remote.ssl.on=true*). The user is authorized by the rights assigned to the SSL certificate user ID in the Server API User Registry.

Note: When SSL is used and the remote client application uses Server API listener objects, then the client application must have its own certificate that is trusted by the TDI Server (this is

analogous to the setup for SSL client authentication). If there is no client certificate trusted by the TDI Server, the listener objects do not work and the remote client application cannot receive notifications from the TDI Server.

Username/password based authentication

Only works when *api.custom.authentication* is set to a JavaScript authentication file. This authentication method works regardless of whether SSL is used and whether SSL client authentication is used. The user is authorized as per the rights assigned to the *username* user in the “Server API User Registry” on page 87.

LDAP authentication

This was described in section “LDAP Authentication support” on page 80, and is dependent on a number of *api.custom.authentication* settings in the *global.properties* or *solution.properties*.

Host-based authentication

Only works when *api.remote.ssl.on=false*. Then opening of Server API sessions without username and password supplied from all hosts specified by the *api.remote.nonssl.hosts* property are successfully authenticated and granted admin authority. The *api.remote.nonssl.hosts* property can be specified in the *global.properties* or *solution.properties*.

Server API JMX layer does not support username and password authentication

The remote JMX layer of the Server API does not support username and password based authentication. It ignores the *api.custom.authentication* properties. Regardless of the value of these properties and whether custom authentication is enabled or not for the Server API, the remote JMX layer performs the following authentication:

- If SSL is turned on and SSL client authentication is turned on, the remote JMX layer performs SSL-based authentication (as in TDI 6.0).
- If SSL is turned on and SSL client authentication is turned off, the remote JMX layer does not work.
- If SSL is turned off, the remote JMX client is successfully authenticated only if its host is specified on the *api.remote.nonssl.hosts* property, that is, host-based authentication is assumed. In this case the client is granted admin authority.

The net result is that the Server API JMX layer does not support username and password authentication:

Server API authentication setup examples

Authentication configuration examples:

1. Non-SSL configuration and custom authentication:

```
api.remote.ssl.on=false
api.remote.nonssl.hosts=192.168.113.51, 192.168.113.52
api.custom.authentication=ldap_auth.js
```

SSL is not used.

- Authentication requests with no username and password supplied succeed only if they are invoked from the localhost or from 192.168.113.51 or 192.168.113.52.
- Authentication requests with username and password supplied succeed only if the *ldap_auth.js* successfully authenticates the user specified with the username and password parameters.
- Remote JMX clients are authenticated only when the request comes from the localhost or from 192.168.113.51 or 192.168.113.52.

2. SSL (without client authentication) and custom authentication:

```
api.remote.ssl.on=true
api.remote.ssl.client.auth.on=false
api.custom.authentication=ldap_auth.js
```

SSL is used for remote Server API communication.

- Authentication requests with no username and password supplied fail because neither SSL client authentication, nor host-based authentication is switched on.

- Authentication requests with username and password supplied succeed only if the `ldap_auth.js` successfully authenticates the user specified with the username and password parameters.
- Host-based authentication is not available in this case regardless of the value of the `api.remote.nonssl.hosts` parameter, because `api.remote.ssl.on` is set to `true`.
- Remote JMX layer is not be accessible. This is because SSL is turned on but SSL client authentication is not used.

3. SSL with client authentication and custom authentication:

```
api.remote.ssl.on=true
api.remote.ssl.client.auth.on=true
api.custom.authentication=ldap_auth.js
```

SSL is used for remote Server API communication and the Server requires SSL client authentication.

- Authentication requests with no username and password supplied succeed when the SSL certificate of the client is present in the Server's truststore (or verifiable using the certificates in the truststore).
- Authentication requests with username and password supplied succeed only when the SSL client authentication is successful (the SSL certificate of the client is present in the Server's truststore) and the `ldap_auth.js` script successfully authenticates the user specified with the username and password parameters. In this case, authorization is performed based on the username parameter from the username and password supplied and not with the user identity from the SSL client certificate.
- Host-based authentication is not available in this case regardless of the value of the `api.remote.nonssl.hosts` parameter, because `api.remote.ssl.on` is set to `true`.
- Remote JMX clients are authenticated when the SSL certificate of the client is present in the Server's truststore (or verifiable using the certificates in the truststore).

4. SSL with client authentication & no custom authentication:

```
api.remote.ssl.on=true
api.remote.ssl.client.auth.on=true
api.custom.authentication=
```

(as an alternative, the "`api.custom.authentication`" property may be missing entirely)

SSL is used for remote Server API communication and the Server requires SSL client authentication.

- Authentication requests with no username and password supplied succeed when the SSL certificate of the client is present in the Server's truststore (or verifiable using the certificates in the truststore).
- Authentication requests with username and password supplied do not succeed because custom authentication is not configured.
- Host-based authentication is not available in this case regardless of the value of the `api.remote.nonssl.hosts` parameter, because `api.remote.ssl.on` is set to `true`.
- Remote JMX clients are authenticated successfully only when the SSL certificate of the client is present in the Server's truststore.

Server API Authorization

After a client Server API session request is authenticated it needs to be authorized.

Users of the Remote API can be assigned several roles; a role defines a list of Server API calls that can be executed by the user and also defines in what context these calls can be executed. A Server API method can be executed if there is at least one role assigned to the user that allows the execution of this method in the context the user tries to execute it. For example, a role can grant the user rights to execute only specific AssemblyLines from a specific configuration. Refer to "Server API User Registry" on page 87 for details on how to create the file that holds these user rights.

Authorization is based on the user id. Depending on the authentication mechanism used the user id is retrieved in a different way:

- SSL based authentication – the user id is the distinguished name (DN) of the client’s SSL certificate.
- Username or password based authentication – the user id is the username supplied in the username and password pair.
- Host based authentication – no user id can be retrieved from the client using this authentication mechanism; in this case the client session is authorized with the *admin* role.

Authorization roles

Users of the Remote API are assigned roles; a role defines a list of Server API calls that can be executed by the user and also defines in what context these calls can be invoked. For example, a role can grant the user rights to invoke only specific AssemblyLines from a specific configuration.

Several roles can be assigned to a user, including assigning the same role several times with different parameters. A Server API method can be invoked if there is at least one role assigned to the user that allows the execution of this method in the context the user tries to execute it.

There are no deny semantics – actions cannot be explicitly forbidden. The following roles apply to the Server API security model:

<p>Read role: read [list_of(configuration)]</p>	<p>The <i>read</i> role allows the user read data from the Server’s configuration(s).</p> <p>If no list of configurations is specified or the list is empty, the user is not allowed to read any configuration.</p> <p>A special value * (asterisk) can be specified for the list of configurations and this means that the user is allowed to read (through Server API calls) all configurations currently loaded by the Server.</p> <p>When the list of configurations is not null/empty and does not specify * the user is allowed to read only the configurations specified.</p> <p>The <i>read</i> role does not grant permission to start processes (AssemblyLines) or apply any changes to the Server and its configurations. For example:</p> <pre>[ROLE]:read [CONFIG]:*</pre>
<p>Execute role: execute [list_of(configuration [list_of(AssemblyLines)])]</p>	<p>The <i>execute</i> role gives the user permissions to execute AssemblyLines.</p> <p>If no list of configurations is specified or the list is empty, the user is not allowed to execute any AssemblyLine from any configuration.</p> <p>A special value * (asterisk) can be specified for the list of configurations and this means that the user is allowed to execute all AssemblyLines from all configurations.</p> <p>When the list of configurations is present and does not specify * the user is only allowed to start the processes from the configurations specified in the list. For each configuration specified in the list:</p> <ul style="list-style-type: none"> • If a list of AssemblyLines is not specified, the user is not allowed to execute any AssemblyLine from this configuration. • If a special value * (asterisk) is specified for the list of AssemblyLines, the user is allowed to execute all AssemblyLines from this configuration. • If the list of AssemblyLines is present and does not specify * the user is allowed to execute only the AssemblyLines specified in the list. <p>For example:</p> <pre>[ROLE]:execute [CONFIG]:C:/TDI/rs.xml [AL]:* [CONFIG]:C:/TDI/prototype.xml [AL]:TestAssemblyLine</pre>

<i>Admin</i> role: admin	<p>The <i>admin</i> role allows the user to execute all Server API calls in every possible context.</p> <p>A user with <i>admin</i> role is allowed to read and modify configurations, to load new configurations, to execute AssemblyLines, to read and modify server parameters.</p> <p>For example: [ROLE]:admin</p> <p>Note:</p> <p>Admin role is required to use the Remote Configuration Editor. Also see “Using the Remote Configuration Editor” on page 104.</p>
---------------------------------	---

The values specified in a [CONFIG] tag can be either Config file names, or Solution Names if they have been specified in the Config file.

Server API User Registry

The User Registry, identified by the *api.user.registry* property in the *global.properties* or *solution.properties* file is a text file that maintains the information about all the users of the API and their roles. This file is encrypted with the Server's certificate specified by the *api.key.alias* property from the keystore specified by the *api.keystore* property. The encryption algorithm employed is Asymmetric RSA encryption or decryption; that is why the “Example Server certificate creation” on page 109 specifying the RSA algorithm, which is the default algorithm of the “The TDI Encryption utility” on page 96 provided with Tivoli Directory Integrator that you can use for this purpose. On startup, the Server API engine decrypts and reads this file into its memory structures.

Notes:

1. The entire user registry file is encrypted as it is, block by block, in a straightforward manner using the RSA algorithm and the server public key. A digital signature or some sort of hashing is not utilized.
2. The authorization against the user registry is not optional. Currently the Tivoli Directory Integrator Server has no concept of a plug-in authorization mechanism.

The contents of the Identity Registry text file is structured as follows:

```
[USER]
[ID]:<user_identifier>
[ROLE]:<role_identifier>
    [CONFIG]:<config_identifier>
        [AL]:<assembly_line_name>
        [AL]:<assembly_line_name>
    ...
    [CONFIG]:<config_id>
    ...
[ROLE]:<role_identifier>
...
[ROLE]:<role_identifier>
...
[ENDUSER]

[USER]
[ID]:<user_identifier>
[ROLE]:<role_identifier>
...
[ENDUSER]
...
```

Each tag must span a single line and each tag must be on a separate line. Tabs and spaces do not matter. Empty lines may appear anywhere. The tags in the Identity Registry file and their arguments are as follows:

Tag	Argument
[USER]	This tag takes no arguments, and serves as an opening bracket for the tags below; a [USER] and [ENDUSER] pair of tags, each placed on a single line, provide the definition of a single user in the registry file. There can be multiple pairs of this type, each of which specify a user of the Server API.
[ID]:<user_identifier>	This tag is the first tag after the [USER] tag and its argument <user_identifier> is the unique identifier of the user of the Server API. This ID value is the 89 from the truststore file. The tag and the argument of the tag are placed on a single line, and there can be only one [ID]: tag included in a [USER] and [ENDUSER] pair.
[ROLE]:<role_identifier>	This tag specifies a role for the user. Possible roles are: read , execute or admin . Everything after the [ROLE]: tag and its argument and before another [ROLE]: tag or an [ENDUSER] tag (whichever comes first) specifies details of this user role. The tag and the argument of the tag are placed on a single line, and there can be multiple [ROLE]: tags included in a [USER] and [ENDUSER] pair, specifying multiple roles for that user.
[CONFIG]:<config_id>	This tag specifies the identifier of a Tivoli Directory Integrator configuration, the absolute file path of the configuration. Relative file paths are not recognized. This tag is subordinate to a [ROLE]: tag, and the tag specifies a configuration for the role given by the [ROLE]: tag. This tag and its argument are placed on a single line, and there can be multiple [CONFIG]: tags, all belonging to the superior [ROLE]: tag. If no [CONFIG]: tag is associated with a [ROLE]: tag, the list of configurations for the corresponding role definition is empty.
[AL]:<assembly_line_name>	This tag specifies an AssemblyLine name. This tag is subordinate to a [CONFIG]: tag. The tag and its argument are placed on a single line, and there can be multiple [AL]: tags, all belonging to the superior [CONFIG]: tag. If no [AL]: tag is associated with a [CONFIG]: tag, the list of AssemblyLines for the corresponding configuration ID is empty.

The following text is an example of an Identity Registry file:

```

USER]
[ID]:CN=Stan, OU=TDI, O=IBM, C=US
[ROLE]:admin
[ENDUSER]

[USER]
[ID]:CN=John, OU=TDI, O=IBM, C=US
[ROLE]:read
    [CONFIG]:*
[ROLE]:execute
    [CONFIG]:C:/TDI/rs.xml
        [AL]:*
        [CONFIG]:C:/TDI/prototype.xml
        [AL]:TestAssemblyLine
[ENDUSER]

[USER]
[ID]:CN=Peter, OU=TDI, O=IBM, C=US
[ROLE]:execute
    [CONFIG]:C:/TDI/rs.xml
        [AL]:*
[ENDUSER]
```

This set of Identity Registry entries implies the following constraints:

- "Stan" is an administrator according to this registry file, and is allowed to perform each and every Server API operation.
- John is allowed to read all configurations loaded on the Server, but can only execute processes from two configurations:
 - From "rs.xml", John can execute all AssemblyLines.
 - From "prototype.xml" John is only allowed to execute the AssemblyLine named "TestAssemblyLine".
- Peter can only execute all AssemblyLines from the "rs.xml" configuration.

Note: The **keytool** and/or the **Ikeyman** utility can be used to obtain the user ID from the truststore file. The following command line prints all users from the truststore file:

```
keytool -v -list -keystore <trust_store_file> -storepass <trust_store_pass>
```

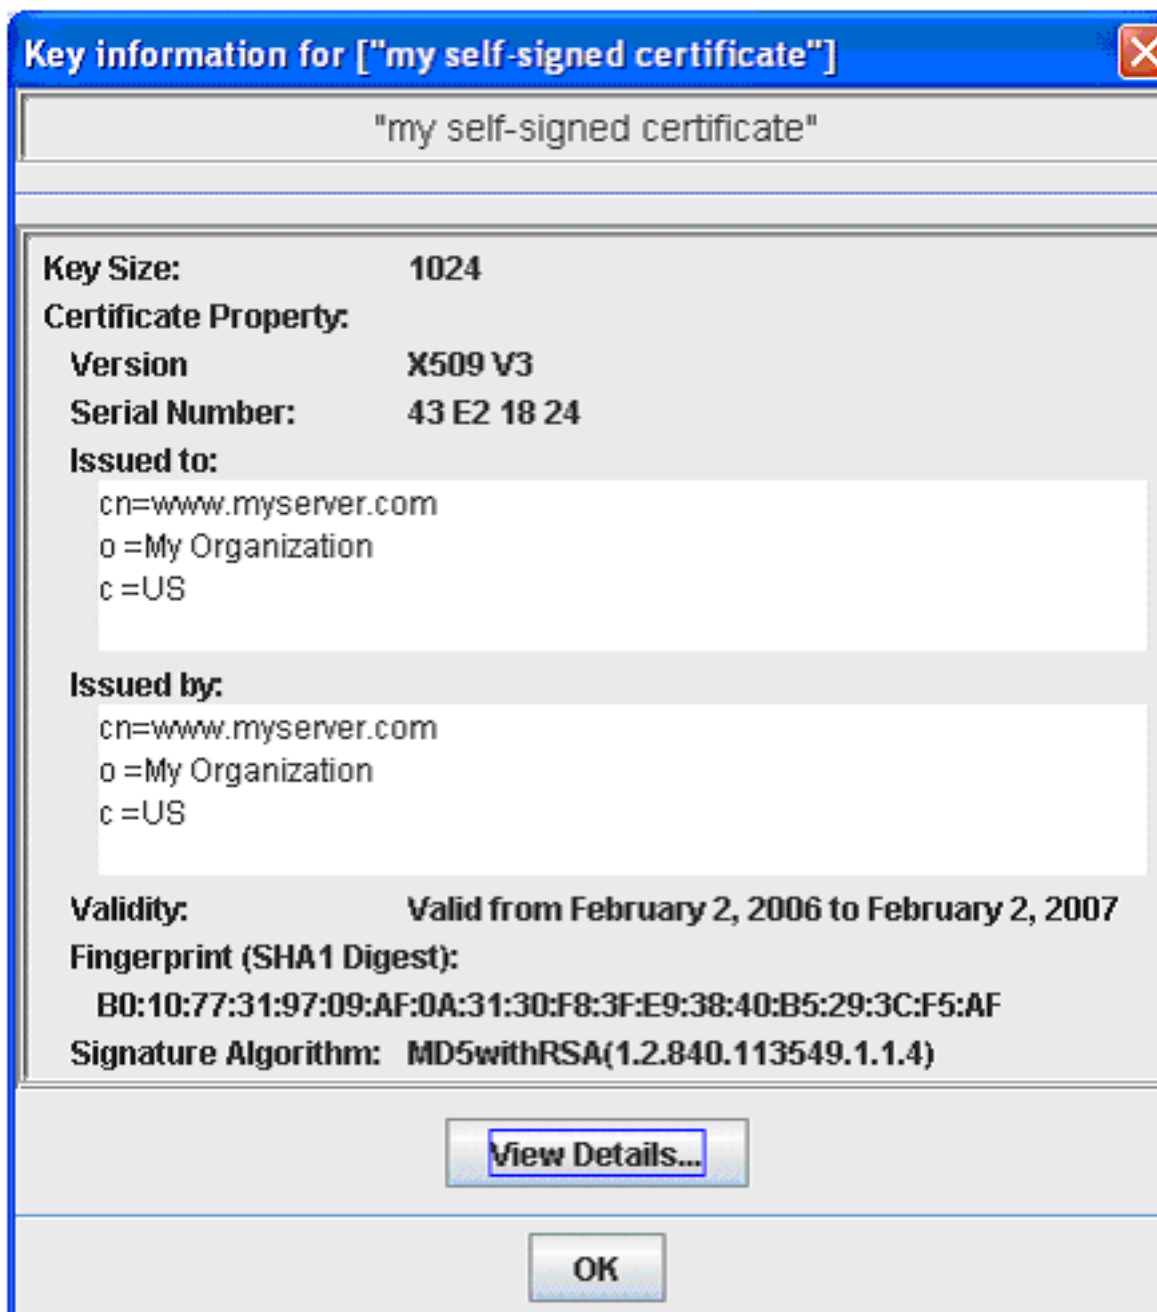
where <trust_store_file> is the keystore file that contains the certificates of all trusted users and <trust_store_pass> is the password for this keystore file. This command line prints something like the text below for each user certificate:

```
Owner: CN=Stan, OU=TDI, O=IBM, C=US
Issuer: CN=Stan, OU=TDI, O=IBM, C=US
Serial number: 408f6a34
Valid from: 4/28/04 11:24 AM until: 7/27/04 11:24 AM
Certificate fingerprints:
    MD5:  F6:EF:81:8B:4C:0F:10:E4:A0:16:99:AB:42:29:70:8B
    SHA1: FE:37:62:8B:42:2F:54:F8:F6:F3:FC:A1:DD:7D:2A:51:9A:85:09:02
```

The value of the **Owner** field **must** be specified as value for the [ID]: tag in the Identity Registry as is, including all white space and commas. For this example, the line with the ID tag looks like:
[ID]:CN=Stan, OU=TDI, O=IBM, C=US

An alternative way to obtain the user ID from the truststore file is to use Ikeyman in the following way:

1. Start Ikeyman.
2. From the **Key Database File** menu click **Open...**
3. In the **Open** field, set the appropriate values and click **OK**.
4. In the **Password** field, enter the password for the truststore file.
5. Click on the certificate you are interested in.
6. Click the **View/Edit...** button. A window opens which contains information on the subject's DN (user ID).



Server Audit Capabilities

The Tivoli Directory Integrator Audit Component enables the TDI Server to audit events such as authentication and authorization in the Server API.

Notifications are generated when authentication and authorization (auth*) events occur. Audit data is packaged into an Entry and provided as user data in the notification. The "Audit Service" consists of a separate Audit config that is automatically loaded by the TDI server. The Audit config contains auto-started Audit AssemblyLines. The Audit ALs iterate on the notification connector using suitable filters. Tivoli Directory Integrator users can even generate "user defined notifications" if they want to create audit events from within their own code.

Tivoli Directory Integrator auditing contains two main parts:

- A way for generating the necessary audit information
- An "Audit service" for handling existing audit data

Generating necessary audit information is implemented by creating Tivoli Directory Integrator Entries on each audit point in the Server API, and by broadcasting these Entries wrapped in a notification. For this purpose a new class is presented in the Server API (`com.ibm.di.api.APIAuditor`), that generates the Entry, attaches the Entry as `UserData` to a notification, and sends it to all interested listeners.

The "Audit Service" is the main consumer of the audit notifications. The Audit Service is a config consisting of several ALs that iterate on the Notification Connector. Using different filters can register to a variety of notification types.

Auditing scope

The TDI audit capability follows only what people do, and does not follow Server events in general. There is a difference between a user being authorized to perform a task (stop an AL) and the task actually being performed (AL is terminated). Being authorized is an authorization event and the performing of a legal action, like stopping an AL, is a Server event. When a user instructs an AL to stop and the AL terminates, an authorization event is paired with a Server event. At other times, a Server event occurs by itself, as when an AL completes naturally. Only events which involve direct user interaction are audited. This limits the default audit points to authentication and authorization events inside the Server API. Almost every method exposed by the Server API is protected by its own piece of authorization code. The Audit component does not try to send notifications for all authorization events, but selects a reasonable subset of authorization-guarded Server API methods. The principles for the selection are to audit all events that:

- Delete logs or tombstones
- Start or stop TDI entities such as configs, ALs, and the Server
- Replace the config instance configuration: replace the config instance configuration or the check-in configuration
- Allow the user to change vital TDI data: set external property, post a message in the System Queue, call custom Java code inside the Tivoli Directory Integrator JVM

Suppression of notifications

The Tivoli Directory Integrator Server API allows certain notification types to be suppressed for improved performance. The notification framework does not propagate suppressed events. If you try to broadcast an event of a type *suppressed*, the Server API does not issue an error. However, the suppressed event cannot reach any of the registered notification listeners. The list of suppressed event types is configured by a system property named:

```
api.notification.suppress
```

By default, all authentication and authorization events are suppressed:

```
api.notification.suppress=di.server.api.authenticate di.server.api.authorize*
```

The event types in the list are separated by spaces. Wildcards matching multiple event types are allowed. If the event type property is missing or is empty, no events are suppressed. You can suppress all custom notifications by typing:

```
api.notification.suppress=user
```

Note: Suppression affects the whole TDI Server and can result in suppression of all kinds of notifications. Even built-in notifications, such as an AssemblyLine starting or the Server shutting down, can be suppressed. Improper use of the suppression capabilities can interfere with the work of components that listen for notifications such as the Tombstone Manager and the Server Notifications Connector.

Sending notifications

Sending notifications uses a method in `com.ibm.di.api.APIEngine`:

```
public static void sendNotification (String type, String id, Object data, String configInstanceId)
```

This method creates a `DIEvent`. By this means, a notification is delivered to every Listener registered to receive the a particular type of notification. Notification delivery parameters include:

Table 12. Notification Delivery Parameters

Parameter name	Definition
type	Notification event type.
id	Notification event ID.
data	Notification event <code>UserData</code> object in the form of a Java object with additional information.
configInstanceId	Notification <code>ConfigInstance</code> ID to which the notification is bound.

The `com.ibm.di.api.APIEngine` method throws `DIException` if the type parameter is null. Calls to the method can be invoked either:

- Locally, from the Tivoli Directory Integrator Server JVM. This type of access includes scripting in `AssemblyLine` hooks and also uses the API from new Connectors implemented in Java and deployed on the Tivoli Directory Integrator Server
- Remotely, from another JVM (on the local or a remote network computer), through Remote Method Invocation (RMI). This type of access uses solutions that:
 - Connect remotely to Tivoli Directory Integrator
 - Manage processes within Tivoli Directory Integrator
 - Build business logic on top of Tivoli Directory Integrator
 - Are applications dedicated only to Tivoli Directory Integrator
 - Are applications that use Tivoli Directory Integrator to accomplish some of their goals

TDI Server Instance Security

This section does not deal with the specifics of client (TDI-based or other) access to a TDI Server, this is discussed in “Remote Server API” on page 72; instead, it focuses on the encryption algorithms used, and the miscellaneous configuration files needed to set up a server instance.

The TDI Server requires a keystore containing both its private key and associated certificate/public key that is used for PKI encryption of Config Files, properties in Properties files, Server User registry files and other objects, as well as being used for SSL communication.

The system properties `api.keystore` and `api.key.alias` specify the keystore and the key alias of the Server's certificate/key within the keystore. The password of the keystore and the password of the key itself (if different from the keystore password) are specified in the Server's stash file. Access to a keystore is guarded by a password, defined at the time the keystore is created, by the person who creates the keystore, and changeable only when providing the current password. In addition, each private key in a keystore can be guarded by its own password. For more information on the stash file of the server, see section “Stash File” on page 93.

The RSA algorithm is used for encryption of files and property values. It is used as a block cipher where the block size is determined by the modulus component of the RSA key. Encryption is done in ECB (Electronic Codebook) mode. PKCS#1 Padding is applied separately on each block. Note that the same RSA key-pair, which is used for encryption of files, is also used for SSL communication with the Server. Tivoli Directory Integrator uses the RSA implementation from the IBMJCE security provider. All key sizes supported by that provider are also supported by Tivoli Directory Integrator. From Tivoli Directory

Integrator v7.0, secret key ciphers can also be employed for encryption. RSA is used the default for backward compatibility reasons, but secret key ciphers are much faster and much more secure than public key ciphers.

DES and AES algorithms are used for encryption of password-protected configuration files. An encryption key (DES or AES) is derived from the UTF-8 binary representation of the password. The derived encryption key is 64 bit for DES and 128 bit for AES. ECB mode is used with no padding.

DES/AES keys are derived from passwords, when using password-protected configuration files. Apart from the above, the Tivoli Directory Integrator does not generate keys. Existing keys are loaded from an external key store. Key establishment and key store access are performed through the IBMJCE and IBMJSSE2 security providers. All key sizes and algorithms supported by those providers can be used with the Tivoli Directory Integrator.

Stash File

The stash file contains the Server keystore password values encrypted with AES128 with a fixed key. The Server stash file is named "idisrv.sth" (the name is not configurable) and it is loaded by the Server from the Solution Folder. A command line utility for creating a stash file is available in the TDI bin folder:

createstash.bat or *createstash.sh*:

```
createstash <keyStorePassword> [<keyPassword>] [<securityProviderClass>]
```

where *keyStorePassword* is the password of the keystore file specified by the *api.keystore* system property and *<keyPassword>* is the password of the Server's private key specified by the *api.key.alias* system property.

keyPassword is an optional parameter if no *<securityProviderClass>* parameter is specified. If *<keyPassword>* is not specified it is assumed that the Server's private key password is the same as the keystore's password. To use the utility with the *<securityProviderClass>* parameter, you must specify both previous parameters: *keyStorePassword* and *keyPassword*. If a security provider is specified then this provider is used for the cryptography.

The utility creates a stash file named "idisrv.sth" with the specified password(s) in the current directory.

Attention: IBM Tivoli Directory Integrator 7.0 comes bundled with a sample stash file, with a password of "server". For improved security, we strongly advise you to generate your own stash file using the aforementioned utility. Also, the stash file must be kept inaccessible, except for the actual Tivoli Directory Integrator Server that needs it.

Server Security Modes

The Tivoli Directory Integrator Server can run in two modes: **standard** and **secure**.

Standard mode

When run in standard mode, the Server does not PKI encrypt configurations saved on disk, unless a specific Server API call that requests PKI encryption is invoked. When in this mode the Server is able to read both encrypted and unencrypted configurations.

Secure mode

When run in secure mode the Server encrypts all configurations it saves on the disk using PKI encryption. In secure mode the Server can only read and load encrypted configurations. When the system property *com.ibm.di.server.securemode* is set to "true", the Server runs in secure mode. (A system property for the use of the TDI Server can be set by adding it in the *global.properties* or *solution.properties* file or directly specify it on the java command line when starting the TDI server: *-Dcom.ibm.di.server.securemode=true*)

If the command line option *-e* is specified on the java command line when starting the Server, it runs in secure mode regardless of the value of the *com.ibm.di.server.securemode* system property.

Note: Pre-TDI 6.0 password-based encryption of configuration files is supported for backward compatibility. Password-based encryption is used when the user specifies a password when creating the configuration. Pre-TDI 6.0 password-based configuration encryption cannot be combined with PKI encryption. If you specify a password when the Server is run in secure mode, an error message is displayed.

Working with encrypted TDI configuration files

Introduction

To provide confidentiality of data, Tivoli Directory Integrator (TDI) can encrypt configuration files, property values in properties files, server user registry files and JavaScript files.

TDI encryption involves a cryptographic transformation that uses a key or a key-pair. The key/key-pair needs to be hosted in a keystore file.

The cryptographic transformation can be either public-key encryption or secret key encryption. By default TDI uses public key encryption. (The secret key encryption option has been introduced in TDI 7.0. Before that only public key encryption was supported.)

See:

Public key encryption utilizes a key-pair that consists of a public key and a private key. The public key is used for encryption and the private key is used for decryption. Currently only the RSA cipher is supported for public key encryption. Public/private key pairs can be generated and managed using the standard JRE utilities `keytool` and `Ikeyman`. See “Keystore and truststore management” on page 68 for more information on managing certificates with associated public and private keys.

Secret key encryption relies on a secret key, which is used for both encryption and decryption. TDI provides a tool for generating and managing secret keys – see section “The TDI secret key tool” on page 98.

TDI data encryption is configured by the following system properties (these can be set in `global.properties` or `solution.properties`):

- `com.ibm.di.server.encryption.keystore` : the keystore file that contains the key/key-pair for encryption
- `com.ibm.di.server.encryption.keystoretype` : the type of the keystore file
- `com.ibm.di.server.encryption.key.alias` : the alias of the key/key-pair in the keystore
- `com.ibm.di.server.encryption.transformation` : the name of the encryption transformation; see remarks below

The password of the keystore and the password of the key/key-pair itself (if different from the keystore password) are specified in the Server's “Stash File” on page 93. (Access to a keystore is guarded by a password, defined at the time the keystore is created, by the person who creates the keystore, and changeable only when providing the current password. In addition, each private key in a keystore can be guarded by its own password.)

The name of the transformation can be either RSA or some secret key transformation (for example, AES/CBC/PKCS5Padding). More detailed discussion of what is in a transformation name can be found at http://www-128.ibm.com/developerworks/java/jdk/security/50/secguides/JceDocs/api_users_guide.html#trans; general information about Java Security (which is what Tivoli Directory Integrator uses) can be found at <http://www-128.ibm.com/developerworks/java/jdk/security/60/secguides/jsse2Docs/JSSE2RefGuide.html>.

Notes:

1. The “`com.ibm.di.server.encryption.*`” properties affect not only encryption of configurations, but also encryption of property files, JavaScript files and the Server API User Registry.

2. If you change the encryption key and/or the encryption transformation, the Server cannot decipher previously encrypted files. To work around this problem, decrypt the old files with the old key (you must have the old key available in order to do so) and encrypt them with the new key. Encryption and decryption of files can be done with the cryptoutils tool.
3. The standard RSA algorithm has a restriction on the length of data it can work on. TDI uses a custom scheme that splits input data into small enough equally sized blocks and encrypts each of them separately.
4. Data encrypted with RSA result in different cipher-texts on different encryption runs. This effect is a feature of the PKCS#1 padding scheme used with RSA.
5. A secret key (symmetric) cipher can be either a block cipher or a stream cipher. Stream ciphers encrypt the bits of the message one at a time, and block ciphers take a number of bits and encrypt them as a single unit (a block). Block ciphers (for example, AES) use a feedback mode (so that patterns in the plain-text are not preserved in cipher-text) and a padding scheme (to allow encryption of data, whose length is not multiple of the block size of the cipher). Stream ciphers (for example, RC4) do not use a feedback mode and a padding scheme.
6. If the transformation involves a block cipher, it must use some padding scheme (for example, "PKCS5Padding"), otherwise the Server not be able to encrypt data whose length is not multiple of the block size of the cipher. (Stream ciphers do not use padding, so they are not affected by this restriction.)
7. The algorithm of the key/key-pair must match the algorithm in the specified transformation. For example if the transformation is RSA then an RSA key-pair must be provided; if the transformation is DES/ECB/PKCS5Padding, you must provide a DES key. You can generate a new secret key using the skeytool TDI utility.
8. JKS keystores do not support secret keys, so you should use some other keystore type such as JCEKS if you want to use a secret key encryption.
9. When using a block cipher in a feedback mode that requires an initialization vector (IV), encrypted data is prefixed with the initialization vector as plaintext. The IV does not have to be kept secret but must be unpredictable. That is why a random IV is generated for each piece of data that is being encrypted. Generation of random data can sometimes be resource-intensive, so you may wish to consider a non-IV feedback mode (ECB) if performance is an issue.
10. Which secret key transformations are supported for encryption depends on the capabilities of the Java security provider. By default TDI uses the IBMJCE provider. Supported block ciphers are: DES, AES, DESede (Triple DES), Blowfish and RC2. They can be used in any of the following feedback modes – ECB, CBC, CFB, OFB, PCBC. The only available padding scheme is "PKCS5Padding". The MARS block cipher should not be used for encryption, because it does not support padding (<http://www-128.ibm.com/developerworks/java/jdk/security/50/secguides/JceDocs/api/com/ibm/crypto/provider/Mars.html>). Supported stream ciphers are RC4 and ARCFOUR (basically the same cipher under two different names). The SEAL stream cipher requires large keys (160 bit) so it can be used only after configuring unrestricted IBM SDK policy on the TDI JRE (<http://www-128.ibm.com/developerworks/java/jdk/security/50/#sdkpol>).

Creating an encrypted TDI configuration file from scratch

This is how you create an encrypted IBM Tivoli Directory Integrator configuration file from scratch.

Using the cryptoutils command line tool:

1. Create a normal un-encrypted TDI configuration file using the Configuration Editor.
2. Use the cryptoutils command line tool to encrypt this configuration file as described in the ""The TDI Encryption utility"" on page 96" section.
3. In order to run this encrypted configuration file you must start the TDI server in secure mode as described in the "Server Security Modes" section.
4. In order to edit this encrypted configuration file you can use one of two options described in the ""Editing an encrypted TDI configuration file"" on page 96" section.

Editing an encrypted TDI configuration file

You can first decrypt the encrypted configuration file using the `cryptoutils` command line tool as described in the "The TDI Encryption utility" section. Then you can edit the decrypted configuration using the Configuration Editor and finally you can encrypt back the modified configuration file using the `cryptoutils` tool.

Standard TDI encryption of `global.properties` or `solution.properties`

The `global.properties` and `solution.properties` files store a number of properties, some of which can represent sensitive data such as passwords. In order to protect this sensitive data Tivoli Directory Integrator 7.0 is capable of encrypting this data.

All properties whose names are prefixed with `{protect}-` are PKI encrypted by the Server using the Server's public key. The Server's key is specified by the `com.ibm.di.server.encryption.key.alias` property from the keystore specified by the `api.keystore` property. For example, if you want to encrypt a property `com.ibm.di.server.encryption.keystore` you can add the following line in the `global.properties` or `solution.properties` file:

```
{protect}-com.ibm.di.any.property=some_value
```

The next time the Server runs it detects that this property has to be encrypted and it immediately overwrites the file, writing the plain text value "some_value" in encrypted form.

Note: On some operating systems (z/OS, Linux/UNIX systems if so configured) the file might not be accessible for writing. In this case the server outputs a warning message that the file has not been written/encrypted.

Protecting the properties in `global.properties` or `solution.properties` is also accessible from the "Global-Properties" and "Solution-Properties" Property Stores accessible from the "Properties" folder in the Configuration Editor.

Encryption of properties in external property files

Properties stored in external property files can be protected by encryption in just the same way as properties in the `global.properties` or `solution.properties` can. For more information on encrypting properties stored in these files, see the "Standard TDI encryption of `global.properties` or `solution.properties`" section. The syntax of properties in an external property file is as follows:

```
[{protect}-]keyword <colon | equals> [{encr}][{java}]value
```

- The optional `{protect}-` prefix signals that the value either is or should be encrypted. When the value starts with the character sequence `{encr}` it means that the value is already encrypted.
- The optional `{java}` value prefix signals that the value is a serialized java object. The value must be b64-encoded. For example:

```
{protect}-api.truststore.pass={encr}J8AKimpEutu3Bb1OVg55F/5d5v02kXWcNUWnCq3vINUc6K0719z9dEk3H430t2iTT1dZTI6FSSVin9KsCy  
BLmgv+n84w7He1K13ro2dFmZbTYKMxuxGoqN9nL2V0vZoptNqzoWvs6IN/p3VkiIBt1ao/9mEPEKuIwRnKtkQ89Bg=
```

The TDI Encryption utility

In the `TDI_install_dir/serverapi` directory you find a utility (`cryptoutils`) which enable you to decrypt and re-encrypt files, (for example, the Identity Registry file) such that you can edit the file manually.

The tool recognizes the following command-line parameters:

`inputFile`

{required} Specifies the file to be encrypted or decrypted.

`outputFile`

{required} Specifies the new file that is created with the resulting data after the encryption or decryption is done. If the file exists, it is overwritten.

mode {required} Specifies the mode in which the tool operate; it can be one of the following modes:

- *encrypt*: encrypt user registry
- *decrypt*: decrypt user registry
- *encrypt_config*: encrypt a TDI configuration file or a JavaScript file
- *decrypt_config*: decrypt a TDI configuration file or a JavaScript file
- *encrypt_props*: encrypt the values of all protected properties in a TDI properties file
- *decrypt_props*: decrypt the values of all protected properties in a TDI properties file

Note: User Registry files are encrypted differently from configuration and JavaScript files.

keyStore

{required} Specifies the keystore file which contains the key for encryption/decryption.

storepass

{required} Specifies the password of the keystore file.

alias {required} Specifies the alias of the encryption/decryption key in the keystore

keypass

{optional} Specifies the password of the encryption/decryption key; by default, the keystore password is used to access the key

transformation

{optional} Specifies the name of the cryptography transformation used for encryption/decryption; can be RSA or any secret key transformation (for example, AES/CBC/PKCS5Padding); the default is RSA.

storetype

{optional} Specifies the type of the keystore file (for example, JKS); this parameter is case-insensitive (JCEKS and jceks are equivalent); if this parameter is missing, the default keystore type of the JRE (configured by the "keystore.type" security property in the java.security file of the JRE) is used.

cryptoproviderclass

{optional} Specifies the Java security provider which is used for encryption/decryption (but not for keystore access); by default the providers from the security provider list of the JRE (configured in java.security JRE file) is used.

Examples:

Encrypt the User Registry

A TDI Server running in secure mode requires that the User Registry is encrypted with the Server key.

You can encrypt a plaintext User Registry file like this:

```
cryptoutils -input registry.txt -output registry.enc -mode encrypt -keystore ../testserver.jks -storepass server -alias server
```

Decrypt a TDI configuration

```
cryptoutils -input myconfig.enc.xml -output myconfig.xml -mode decrypt_config -keystore ../testserver.jks  
-storepass server -alias server
```

This command decrypts the "myconfig.enc.xml" configuration file (possibly created by a TDI Server, which runs in secure mode). Now the decrypted configuration "myconfig.xml" can be easily modified using the Configuration Editor. After modifying the configuration, it can be encrypted again, so that a TDI Server in secure mode can read and use it.

Encrypt a TDI configuration using a symmetric cipher (rather than the default "RSA")

```
cryptoutils -input myconfig.xml -output myconfig.enc.xml -mode encrypt_config -keystore ../server.jck  
-storepass server -alias server -transformation AES/CBC/PKCS5Padding -storetype jceks
```


The above command assumes that the keystore "server.jck" exists. That keystore is supposed to contain an AES secret key under alias "server".

Decrypt the global.properties file

The Tivoli Directory Integrator Server automatically encrypts the values of protected properties when reading the global.properties or solution.properties file.

You can decrypt all encrypted values in the global.properties file like this:

```
cryptoutils -input ../etc/global.properties -output ../etc/global.properties -mode decrypt_props
-keystore ../testserver.jks -storepass server -alias server
```

Note: When the cryptoutils tool is used to encrypt and decrypt the User Registry, configuration files (see the "Server Security Modes" on page 93" section for details how the server treats encrypted configurations) or "Encryption of TDI Server Hooks" on page 104, it encrypts and decrypts a file as a whole.

On the other hand, the encryption/decryption mode for property files encrypts/decrypt only the values of the protected properties and not the whole file. Thus after encrypting a .properties file using *encrypt_props* mode, the property keys and the comments in the file are still readable by humans. For more information on protected properties see sections "Standard TDI encryption of global.properties or solution.properties" on page 96 and "Encryption of properties in external property files" on page 96.

The TDI secret key tool

The standard key management utilities which come with IBM JRE 5.0 – keytool and Ikeyman cannot work with secret keys. That is why TDI offers its own tool for handling secret keys. The tool is named *skeytool* and is located in the "bin" folder of the TDI installation.

skeytool can perform the following operations:

Generate a secret key and put it in a keystore

To use this operation, specify the *-gensecretkey* command-line switch. The operation accepts the following parameters:

alias [required] the alias under which the generated key is stored; if there is already a key under that alias, it is overwritten.

keyalg [required] the algorithm for which the key is generated.

keysize [required] the size of the generated key.

keystore [required] the keystore file in which the generated key is stored; if the file does not exist it is created.

storepass [required] the password for the keystore file.

storetype [required] the type of the keystore file (that is, "JCEKS"); the parameter is case-insensitive ("jceks" is the same as "JCEKS"); note that JKS keystores are not suitable for secret keys.

keypass [optional] a password to protect the generated key in the keystore; by default, the keystore password is also used for the key.

keygenproviderclass [optional] the class name of a Java security provider, which is used for key generation

(but not for keystore access); by default the providers from the security provider list of the JRE (configured in `java.security` JRE file) is used.

Example:

```
skeytool -genseckey -alias myseckey -keyalg AES -keysize 128 -keystore ks.jck -storepass secret -storetype jceks
```

This command create an AES key of size 128 and put in the `ks.jck` keystore file of type JCEKS under alias `myseckey`. The `ks.jck` file is created if it does not exist. The key have the same password as the keystore.

Delete a secret key from an existing keystore

To use this operation, specify the `-delete` command-line switch. The operation accepts the following parameters:

`alias` [required] the alias under which the key to be deleted is stored.

`keystore`
[required] the keystore file in which the key to be deleted is stored.

`storepass`
[required] the password for the keystore file.

`storetype`
[required] the type of the keystore file (that is, "JCEKS"); the parameter is case-insensitive ("jceks" is the same as "JCEKS"); note that JKS keystores are not suitable for secret keys.

Example:

```
skeytool -delete -alias myseckey -keystore ks.jck -storepass secret -storetype jceks
```

This command deletes the key entry under alias `myseckey` from the existing keystore file `ks.jck` of type JCEKS.

Import an existing key from one keystore (source) into another (destination)

To use this operation, specify the `-importkey` command-line switch. The operation accepts the following parameters:

`srcalias`
[required] the alias, under which the key is located in the source keystore.

`srckeystore`
[required] the source keystore file.

`srcstorepass`
[required] the password for the source keystore file.

`srcstoretype`
[required] the type of the source keystore file (that is, "JCEKS"); the parameter is case-insensitive ("jceks" is the same as "JCEKS"); note that JKS keystores are not suitable for secret keys.

`srckeypass`
[optional] the password of the key in the source keystore; by default, the source keystore password is also used for the key.

`destalias`
[required] the alias, under which the key is stored in the destination keystore; if there is already a key under that alias, it is overwritten.

`destkeystore`
[required] the destination keystore file; if the file does not exist it is created.

`deststorepass`
[required] the password for the destination keystore file.

deststoretype

[required] the type of the destination keystore file (that is, "JCEKS"); the parameter is case-insensitive ("jceks" is the same as "JCEKS"); note that JKS keystores are not suitable for secret keys.

destkeypass

[optional] the password for the key in the destination keystore; by default, the destination keystore password also be used for the key

Example:

```
skeytool -importkey -srcalias myoriginalkey -srckeystore ks1.jck -srcstorepass firstpass  
-srcstoretype jceks -destalias mycopiedkey -destkeystore ks2.jck -deststorepass secondpass  
-deststoretype jceks
```

This command import the key under alias *myoriginalkey* from keystore *ks1.jck* into keystore *ks2.jck* under alias *mycopiedkey*.

List the contents of a keystore

This operation prints the aliases of all entries in the keystore and their type (trusted certificate, private key or secret key).

To use this operation, specify the *-list* command-line switch. The operation accepts the following parameters:

keystore

[required] the keystore file in which the key to be deleted is stored.

storepass

[required] the password for the keystore file.

storetype

[required] the type of the keystore file (that is, "JCEKS"); the parameter is case-insensitive ("jceks" is the same as "JCEKS"); note that JKS keystores are not suitable for secret keys.

Example:

```
skeytool -list -keystore ks.jck -storepass mypass -storetype jceks
```

A sample output from the command could look like this:

```
mycert : trusted certificate entry  
myrsakeypair : private key entry  
mysecretkey : secret key entry
```

Print information about a secret key

This operation prints the algorithm name of the key, the key size and the key itself as raw data.

To use this operation, specify the *-printsecretkey* command-line switch. The operation accepts the following parameters:

alias [required] the alias, under which the generated key is stored; if there is already a key under that alias, it is overwritten.

keystore

[required] the keystore file, where the generated key is stored; if the file does not exist it is created.

storepass

[required] the password for the keystore file.

storetype

[required] the type of the keystore file (that is, "JCEKS"); the parameter is case-insensitive ("jceks" is the same as "JCEKS"); note that JKS keystores are not suitable for secret keys.

keypass

[optional] a password to protect the generated key in the keystore; by default, the keystore password is also used for the key.

Example:

```
skeytool -printsecretkey -alias mysecretkey -keystore ks.jck -storepass secret -storetype jceks
```

A sample output from the command could look like this:

Key algorithm: DES

Key size: 56

Raw key: AE:01:E4:28:5D:6A:12

TDI System Store Security

The Tivoli Directory Integrator System Store is the database or persistent layer where all the information which is required by a Tivoli Directory Integrator Server is persisted. Traditionally, this layer did not have any security around itself. Any user was able to access the System Store. However from Tivoli Directory Integrator 7.0, there is configurable security provided around the System Store.

In Tivoli Directory Integrator 7.0, the System Store by default is used in Network Mode. This way, a number of Tivoli Directory Integrator instances and other applications is able to access the System Store concurrently. In view of the System Store being available over the Network there is a need to have some security built around it in order to protect the data which is maintained by the Tivoli Directory Integrator Server.

Derby (previously known as Cloudscape) provides several ways to define the repository of users and passwords. To specify which of these services to use with your Derby system, set the property `derby.authentication.provider` to the appropriate value as discussed in the appropriate section listed below.

External Directory Service

A directory service stores names and attributes of those names. Derby uses the Java naming and directory interface (JNDI) to interact with external directory services that can provide authentication of users' names and passwords.

You can allow Derby to authenticate users against an existing LDAP directory service within your enterprise. LDAP (lightweight directory access protocol) provides an open directory access protocol running over TCP/IP. An LDAP directory service can quickly authenticate a user's name and password.

On configuring a set of properties defined by Derby you can start using the External Directory Service as a repository for user names and passwords.

User-defined class

The user defined class approach enables you to hook Derby to any other external authentication service other than LDAP.

Set `derby.authentication.provider` to the full name of a class that implements the public interface `org.apache.derby.authentication.UserAuthenticator`. By writing your own class that fulfills some minimal requirements, you can hook Derby up to an external authentication service.

The class that provides the external authentication service must implement the public interface `org.apache.derby.authentication.UserAuthenticator` and throw exceptions of type `java.sql.SQLException` where appropriate.

Built-in Derby Users

Derby provides a simple repository for storing the user names and passwords. For using this built-in repository the property `derby.authentication.provider=BUILTIN` should be set..

The Tivoli Directory Integrator System Store is using the Built-in repository for storing the user name and password. Since Tivoli Directory Integrator have only one user for accessing the System Store this is the most viable provider that can be used.

User Authentication:

The user authentication details deal with the authentication of users. The user authentication mechanism only authenticates if the user name is present in the mentioned repository (it can any one of the repositories which are mentioned above) and if the password is correct for the specified user. However if you want to have more control over the access rights, you can use the User Authorization mechanism provided by Derby.

The master switch for requiring that users be authenticated against provided parameters is the property `derby.connection.requireAuthentication` - the default is *TRUE*.

The access modes can be set using the property `derby.database.defaultAccessMode=fullaccess`. This property sets the default access mode for all the users in the Derby repository. This property also defines the access level for the System Store user. The different access levels supported by Derby are *fullAccess*, *readOnly*, and *noAccess*. However if you want to have specific access modes for specific users, you can assign access using the properties mentioned below:

- `derby.database.fullAccessUsers=<usernames>` for allowing full access to users.
- `derby.database.readOnlyAccessUsers=<usernames>` for allowing read only access to users.
- `derby.database.noAccessUsers=<usernames>` for not allowing users to access the database.

The *usernames* should be a comma separated list of users for example

`derby.database.fullAccessUsers=sa, mary`

In the current version of Tivoli Directory Integrator we have only one user accessing the System Store. This user is required to perform all the operations on the System Store hence we have set the access mode to *fullAccess*.

Miscellaneous Config File features

The "password" configuration parameter type

The configuration parameters of a Tivoli Directory Integrator component in a Config can be "string", "number", "boolean", and so on. One of the available types is "password". If a configuration parameter is of type password, then the Configuration Editor shows its value in the component configuration window as a sequence of '*' characters – both when typing in a new password, and when opening an existing configuration for editing or running.

Component Password Protection

Tivoli Directory Integrator saves configuration information in an XML file which contains clear text for all configuration values. This includes sensitive information like passwords. Tivoli Directory Integrator supports encryption of the entire configuration file but does not encrypt or protect sensitive information when the configuration file is saved in clear text.

Tivoli Directory Integrator provides a way to better protect passwords that are needed for its various components; it hides the passwords in a clear text configuration and provides default security for passwords that are stored. In order to do this component passwords are defined (stored and retrieved) in a default property store, instead of in the configuration file. In Tivoli Directory Integrator 7.0, a user defined property store can be any system for which there is a connector and the default property store most likely be an external properties file. All component passwords will by default go to this default

property store, instead of in the configuration file (as it is in older versions of the product). Thus, passwords can be isolated from the configuration file unless explicitly overridden by the user (may be appropriate for initial development).

Saving passwords to configured Properties

The password protection mechanism is directly related to the configuration windows offered to the user. The configuration windows, or forms, contain descriptions of each parameter and its syntax. One type of syntax is *password* which causes the Configuration Editor to use a password text field for editing. Whenever the value for a password syntax component parameter is changed, the value of the password is saved in an external repository, called the *Password Store*. This external repository for passwords is configured in the *Properties* page in the configuration editor (*Password-Store*) and is specified in the configuration file for the current TDI solution. If no such property store is configured the password is saved in clear text in the configuration file.

If a default password store is configured, a unique property name is generated the first time a protected/password parameter is saved. This key is used as the key in the password store. The same property name is written to the configuration file as a standard property reference. When the value is later retrieved, standard property resolution takes place to retrieve the actual value from the password store.

If a Password Store is specified, a unique key is generated for the password and the password is saved encrypted in the Password Store under that key. In the configuration file, the password is referenced only by that key.

If no Password Store is specified, the password appears in plain text in the configuration file.

For example:

1. Create a new project from the Configuration Editor
2. Right-click on the "Properties" folder in the navigation view and select "New Property Store" called "MyProps".
3. From the "Connector" tab of the newly created Property Store, type in "MyProps.properties" in the "Collection Path/URL" field.
4. Specify that the new Property Store is used as the Password Property Store (right-click on the new properties store in the navigation view and select **Password Property Store**).
5. Add a new assembly line with a FTP Client Connector.
6. Enter a password in the "Login Password" field of the FTP Client Connector.
7. Save the solution and close the Configuration Editor.

After the above procedure, in the configuration file of the created solution will contain lines that resemble the following text:

```
<parameter name="ftpPass">@SUBSTITUTE{property.MyProps:ftpPass-38ae53e8779cfd65}</parameter>
.....
<PasswordStore>MyProps</PasswordStore>
```

...and in the "MyProperties.properties" file there is a line like the following text:

```
{protect}-ftpPass-38ae53e8779cfd65={encr}GVJC01A7VUIW=
```

This means that the FTP password configuration in the solution file references an encrypted property from the current Password Store - "MyProps". The property key used is "ftpPass-38ae53e8779cfd65".

Protecting attributes from being printed in clear text during tracing

Tivoli Directory Integrator solution builders need a way to protect sensitive data, such as passwords, from being printed in clear text when tracing on the solution is needed. Therefore in Tivoli Directory

Integrator 7.0 some of the methods dealing with the Attribute class have been enhanced to say whether an attribute is protected or not. If the attribute is marked as protected and tracing is on, a fixed number of stars '*' is output instead of the actual value.

When connection parameters are found in the TaskCallBlock (TCB), the values never be logged directly by Tivoli Directory Integrator. The fact that parameters were given is logged, but not the values themselves. If the solution needs to be debugged, those values can be dumped manually, for example using scripting.

Encryption of TDI Server Hooks

Server Hook scripts are defined and made available by creating files in the "serverhooks" subdirectory of the solution directory. Scripts that contain sensitive information should be encrypted with the Server API before adding it to the directory. Scripts can be encrypted by using `cryptoutils` (see "The TDI Encryption utility" on page 96). Note that the TDI server only decrypts script files that have the ".jse" filename extension. The ".jse" extension indicates to the TDI server that the script file is encrypted. That is why, after you encrypt a Server Hook script file, make sure to change its filename extension to ".jse".

Remote Configuration Editor and SSL

The Configuration Editor used to edit remote Config Files (that is, Config files on a remote system) is called the Remote Configuration Editor (Remote CE). The Tivoli Directory Integrator 7.0 Remote CE is capable of starting AssemblyLines in configurations opened for editing. The Remote Configuration Editor is a client of the Server API of the remote TDI Server. Consequently the Remote Configuration Editor is authenticated and authorized as a client of the Server API. In order for this to work when SSL is used:

1. The server to which the Remote Configuration Editor connects must be configured to require SSL client authentication. This is a configuration of the Server API – for details see "SSL-based authentication" on page 78.
2. The Remote Configuration Editor TDI instance must be configured to supply SSL client authentication. This is configured in a "SSL client authentication" on page 65.

This SSL client authentication is needed because the Remote Configuration Editor uses listener objects so that it can be notified when an AssemblyLine has terminated and for this to work with SSL both the client must trust the server identity and server must trust the client identity.

Using the Remote Configuration Editor

Using a Remote Configuration Editor is a little different from using a local CE. When running a remote Configuration Editor to manage a Config on a remote system, you must be mindful of restrictions that apply to the CE in remote mode. Notable restrictions include:

- When editing Config files locally, it is sufficient to have appropriate file system access (read and write) to the Config file. However, when editing a remote Config, you must have Admin privileges on the remote Config Instance.
- When connecting to a data source (using **Connect** buttons in mapping windows), these connections are evaluated locally.

For example: `ldap://localhost:389` results in the Configuration Editor (CE) attempting to connect to the local LDAP server, rather than to the LDAP server on the remote config file computer (that is, on computers running the following operating systems: See "Using the Remote Configuration Editor" on page 88 and "Using the Remote Configuration Editor on z/OS" on page 215.

- HP-UX Integrity
- HP PA-RISC
- Solaris Opteron
- z/OS
- i5/OS
- Linux PPC

- Linux 390
- When generating WebServices-related connectors results in function components that generate the WSDL file, .jar files (using Complex Type Generator), and so on, you are generating them locally. These components are not generated on the remote system to which the CE is connected and must be uploaded manually to the remote system for deployment.
- The remote CE only allows editing and viewing of those Configs that are present in the folder specified by the `api.config.folder` property.
- When working with **System Store** operations (such as deleting the Iterator state key, and so on) that are available in the CE, work with the local system store and not with the remote Tivoli Directory Integrator computer's System Store. Only when the AssemblyLine (AL) is executed does the AL connect to the remote System Store, because at AL execution time, the AL is running inside the remote JVM.
- When you use the **Parameter Substitution** editor (available with Cntl+E), the editor shows only the local properties, and not the properties set on the remote system. Similarly, creating and saving a new property store (file type) stores the property store (file) locally.
- When using the Configuration Editor to edit remote Config Files, you are subject to Server API authentication and authorization, because the CE is acting as a client application. Therefore, in order to use the CE in this way, you must have *admin* access on the Remote Server.
- When using the Remote Server, the Remote Server itself must have sufficient access to the local file system where the Config Files are stored. If the Config Files are stored on a read-only file system or a file storage location where the user ID under which the Remote Server is running does not have write access, you cannot edit remote Configs.

Summary of configuration files and properties dealing with security

Table 13. The table of configuration files that were discussed above and what is contained in each.

Configuration file	Location	Description
global.properties	<i>TDI_home/etc</i>	This file is the primary configuration file for the server.
solution.properties	Solution folder	This file (<i>solution.properties</i>) is initially a copy of <i>global.properties</i> used by the current solution. After you make changes, values in this file override corresponding values in <i>global.properties</i> .
registry.txt	<i>TDI_home/serverapi</i>	This file is the User registry for the Server API, defined by the "api.user.registry" property in <i>global.properties</i>
build.properties	<i>TDI_home/etc</i>	This file contains the Tivoli Directory Integrator build information, build date, version, and so on; it is a text file, and by default the file is in the platform-native encoding.
tdisrvctl-log-4j.properties	<i>TDI_home/etc</i>	This file controls the logging strategy for the <i>tdisrvctl</i> command line utility.
Log4J.properties	<i>TDI_home/etc</i>	This file controls the logging strategy for the server (<i>ibmdisrv</i>) when started from the command line.
jlog.properties	<i>TDI_home/etc</i>	This file controls the tracing and First Failure Data Capture (FFDC) strategy
ibmdi.ico	<i>TDI_home/etc</i>	This file lists the icons for Tivoli Directory Integrator.
idisrv.sth	<i>TDI_home</i>	This file contains the Tivoli Directory Integrator server stash; it is a binary file that contains the encrypted password for the sample server keystore file (<i>testserver.jks</i>).

Table 13. The table of configuration files that were discussed above and what is contained in each. (continued)

Configuration file	Location	Description
derby.properties	<i>TDI_home</i> /etc	This file contains the default configuration for the Derby System Store shipped with Tivoli Directory Integrator.
reconnect.rules	<i>TDI_home</i> /etc	This file contains text that defines reconnect rules for how Tivoli Directory Integrator should handle reconnect exceptions.
global.properties.v611	<i>TDI_home</i> /etc	This file serves as a sample place holder and is useful during migration.
TDI0700.SYS2	<i>TDI_home</i> /etc	This is the product signature (license) file used by the ITLM agent to recognize Tivoli Directory Integrator.
pkcs11.cfg	<i>TDI_home</i> /etc	This file is used for initializing IBM pkcs11 implementation provider. For details refer to section PKCS11 Configuration File.
testadmin.der	<i>TDI_home</i> /serverapi	This file is the exported certificate from testadmin.jks.
testadmin.jks	<i>TDI_home</i> /serverapi	This file contains an example keystore and truststore for a Server API remote client.
cryptoutils.bat(sh)	<i>TDI_home</i> /serverapi	This file is a command line utility (shell script) used for encrypting and decrypting Tivoli Directory Integrator configurations and the user registry file.
testserver.jks	<i>TDI_home</i>	This file is a sample server keystore and truststore, referenced as an example.
testserver.der	<i>TDI_home</i>	This file is an exported sample server certificate, ready to be imported in a truststore.
am_config.properties	<i>TDI_home</i> /ActionManager	This file configures the Action Manager.
am_logging.properties	<i>TDI_home</i> /ActionManager	This file configures Action Manager logging.
ibmdiservice.props	<i>TDI_home</i> /wind32_service	This file configures the Windows service.
mqeconfig.props	<i>TDI_home</i> /jars/plugins/	This file allows configuration of the MQe service. In Tivoli Directory Integrator, you can access MQe using authentication for the MQe Mini-Certificate Server to issue certificates; the certificates are then used for authentication. When authenticating, additional properties available in Tivoli Directory Integrator that must be added to the mqeconfig.props properties file.

Note: The file registry.txt can be encrypted and decrypted using the “The TDI Encryption utility” on page 96. The cryptoutil tool should not be applied on global.properties or solution.properties. You can encrypt individual property values but not the whole properties file.

Table 14. The table of properties that are referenced above, characteristics about them, what they do, what their value can be, what they are used for.

Name	Possible values	Description
com.ibm.di.server.securemode	true/false	On or off switch for secure mode.
api.keystore	file name	Server keystore used for SSL certificates. Previously com.ibm.di.server.keystore.

Table 14. The table of properties that are referenced above, characteristics about them, what they do, what their value can be, what they are used for. (continued)

Name	Possible values	Description
api.key.alias	Key alias	Key alias from keystore for SSL certificates. Previously com.ibm.di.server.key.alias.
{protect}-api.keystore.password	SSL keystore password	Keystore password for SSL. Added in TDI 7.0.
{protect}-api.key.password	SSL key password	Key password for SSL. Added in TDI 7.0.
com.ibm.di.server.encryption.keystore	file name	Data encryption for the keystore that hosts the key used by the Server. Added in TDI 7.0.
com.ibm.di.server.encryption.key.alias	Key alias	Encryption keystore key alias. Added in TDI 7.0
com.ibm.di.server.encryption.keystoretype	Keystore type, that is, "JKS", "JCEKS", and so on.	Keystore type that hosts the encryption key of the Server. Added in TDI 7.0.
com.ibm.di.server.encryption.transformation	"RSA" or some secret key transformation	Server transformation used for encryption. Can be set to either "RSA" (public key encryption) or to some secret key transformation (94 of the Tivoli Directory IntegratorServer Security section). Added in TDI 7.0.
api.on	true/false	On or off Server API switch.
api.user.registry	file name	Server API users registry file
api.user.registry.encryption.on	true/false	User registry switch for encrypted or not encrypted.
api.remote.on	true/false	On or off switch for remote Server API. The default setting is true.
api.remote.ssl.on	true/false	On or off switch requiring, or not requiring, SSL for the remote Server API.
api.remote.ssl.client.auth.on	true/false	On or off switch requiring, or not requiring, SSL client authentication for the remote Server API
api.truststore	file name	Server truststore.
api.truststore.pass	*	Trustore password.
api.remote.nonssl.hosts		Non-SSL addresses for accepting non-SSL IP connections.
api.custom.method.invoke.on	true/false	Server API methods for custom method invocation =true when allowed to be used, and =false when disallowed.
api.custom.method.invoke.allowed.classes		Server API classes that can be directly invoked by the Server API methods for custommethod invocation.

Table 14. The table of properties that are referenced above, characteristics about them, what they do, what their value can be, what they are used for. (continued)

Name	Possible values	Description
api.custom.authentication	Script file name or "[ldap]/[jaas]" for built in LDAP or JAASAuthentication	Custom authentication method.
api.custom.authentication.ldap.*		LDAP authentication configuration set of properties.
javax.net.ssl.*		Standard JSSE set of properties for keystore, truststore and their passwords
com.ibm.di.server.pkcs11	false	pkcs11 compliant crypto devices for SSL, required or not required. Added in TDI 7.0
{protect}-com.ibm.di.server.pkcs11.pass	administrator	Access password for pkcs11 compliant crypto device. Added in TDI 7.0
com.ibm.di.server.pkcs11.accl	false	Hardware cryptographic devices to be used for encryption when this property is set to true.

Note: All properties listed in the above table can be set in the configuration file `global.properties`, and can be protected by encryption using the `{protect}-` prefix (see section "Standard TDI encryption of `global.properties` or `solution.properties`" on page 96" for details).

Web Admin Console Security

See "AMC and Action Manager security" on page 195.

Miscellaneous security aspects

Ports and files used by TDI

HTTP Basic Authentication

Some Tivoli Directory Integrator components give you the opportunity to use HTTP Basic Authentication as authentication mechanism. As the name says it is basic (simple) authentication. HTTP Basic Authentication should not be considered secure for any particularly rigorous definition of secure, because the credentials are base64 encoded and they can be easily decoded by someone. You should use more complex schemes to protect their data (for example a combination of turned on SSL and HTTP Basic Authentication). If the component supports HTTP Basic Authentication, then you see the following parameter:

authenticationMethod

Specifies the type of HTTP authentication. If the type of HTTP authentication is set to Anonymous, then no authentication is performed. If HTTP basic authentication is specified, HTTP basic authentication is used with user name and password as specified by the username and password parameters.

Lotus Domino SSL specifics

The Domino APIs for SSL do not use JSSE, and are instead Domino-specific. This means that the Tivoli Directory Integrator truststore and keystore (see section "Client SSL configuration of TDI components" on page 65) do not play any part in SSL configuration for the Domino Change Detection connector. For SSL

configuration of the Domino Change Detection connector, a `TrustedCerts.class` file is used. This file is generated every time the DIIOP process starts (in the Domino Server) and must be in the classpath of TDI (that is, the `ibmdisrv` or `ibmditk` shell scripts which start the TDI server and TDI Configuration Editor respectively). The user must copy the `TrustedCerts.class` to a local path included in the `CLASSPATH` or have the `Lotus\Domino\Data\Domino\Java` of your Domino installation in the classpath. Whether the TDI truststore or keystore are set or not in `global.properties` (or `solution.properties`) is of no consequence to this connector.

Note: Note: The above is related to the configuration of SSL for the Notes Connector and the Domino Change Detection Connector since they use SSL over IIOP.

Certificates for the TDI Web service Suite

The `cn=` portion of the distinguished name (dn) of a certificate to be used with the TDI Web services Server Connectors must match the DNS name or IP address of the host computer on which TDI is running. Otherwise an Exception is thrown, because the client not be able to establish an SSL connection to the TDI Web services Server Connector. An example of the `cn=` portion of the distinguished name of a certificate follows: `cn=www.myserver.com`. (This constraint about the distinguished name in the server's certificate comes from the HTTPS protocol – see rfc2818 "HTTP over TLS.")

Note: If TDI needs to use both a client and a server certificate only the default certificate configured in `global.properties` or `solution.properties` is used, then this must be the same certificate. An alternative would be to write a custom implementation of the `SSLSocket` or the `SSLServerSocket` Java class and make it use a certificate different from the default.

Example Server certificate creation

The following command line creates a self-signed server certificate in the keystore named "MyServerKeyStore.jks".

```
keytool -alias MyServerCertAlias -keyalg RSA -genkey -dname cn=<server_ip_address>
-validity 365 -keystore MyServerKeyStore.jks -storepass mystorepass -keypass mykeypass
```

The alias of the created certificate is "MyServerCertAlias". The RSA algorithm is used to create the key pair. The distinguished name of the certificate is the IP of the server. The certificate is valid for 365 days (one year). The password of the keystore is "mystorepass". The password of the created private key is "mykeypass". The created certificate can then be configured for use by setting the following properties in the `global.properties` or `solution.properties` file:

```
api.key.alias=MyServerCertAlias
api.keystore=MyServerKeyStore.jks
```

MQe authentication with mini-certificates

Tivoli Directory Integrator MQe components can be deployed to take advantage of MQe Mini-Certificate authenticated access. To use these MQe features, it is necessary to download and installation Websphere MQ Everyplace version 2.0.1.7 and WebSphere MQ Everyplace Server Support ES06. Use of certificate authenticated access prevents an anonymous MQe client Queue Manager or application submitting a change password request to the MQe Password Store Connector.

For more information on configuring MQe authentication with Mini-Certificates, see the "Authenticated MQe Access" section in Chapter 8 "MQ Everyplace Password Store" of the *IBM Tivoli Directory Integrator V7.0 Password Synchronization Plug-ins Guide*.

Chapter 7. Reconnect Rule Engine

Introduction

The Tivoli Directory Integrator Server supports reconnect rules that apply to certain error situations during the life of a Connector. The server takes measures, laid out in rules, based on conditions occurring when communicating with target systems.

The AssemblyLine polls the Reconnect Rule Engine every time a Connector raises an exception and the engine recommends a course of action for the current situation. The AssemblyLine code then acts in the proposed way. The possible actions to attempt are:

- to reconnect
- to leave the exception unhandled and let further error mechanisms like error hooks process it.

The *reconnect action* leads to a reconnect attempt only if reconnect is enabled by means of the options available in the Connector's "Connection Errors" tab in the CE. If reconnect is not enabled in this configuration, reconnect is not attempted in case of error regardless of the decision of the Reconnect Rule Engine.

Reconnecting basically involves automatic restart of the Connector and bringing it to its previous position (if so configured). This is done by executing terminate for the Connector, then executing initialize for the Connector and in case of Iterator Connectors, optionally skipping entries until the position before the reconnect is reached. On each reconnect attempt the corresponding reconnect hook is invoked. The script in the hook may eventually change the configuration so that a subsequent reconnect would be successful.

The *error action* implies that no automatic reconnect is attempted and that the corresponding error hooks are invoked. The hooks can eventually perform some custom recovery or error reporting.

Reconnect Rules

The Reconnect Rule Engine makes decisions based on configured rules. Each rule describes what should be done when a given kind of error situation ensues. The engine uses two types of rules:

- **Built-in rules**, which are stored in the `tdi.xml` files of each connector file and are packaged in the connector's jar file; as a result these rules are always specific to the particular connector class and match all connector names; this list of rules is the default list of the Reconnect Rule Engine when working on an error situation for a given Connector; if you have programmed your own Connectors in Java, then for information about how to construct your own built-in rules see section "Connector Reconnect Rules definition" in the "Implementing your own Components in Java" appendix in *IBM Tivoli Directory Integrator V7.0 Reference Guide*.
- For backward compatibility with previous releases of TDI, when the Reconnect Rule Engine is set up it implicitly adds to the built-in rules, a set of rules that prescribe to attempt reconnect on all `IOException`-s and all `CommunicationException`-s (`java.io.IOException` and `javax.naming.CommunicationException`);
- **User-defined rules**, which are loaded from an external text file named `etc/reconnect.rules`; this list of rules overrides the built-in rules. See "User-defined rules configuration" on page 112.

Each rule applies to certain connectors and certain error situations.

A rule has the following parts:

- **Connector Class**: the Java Class of the connectors to which the rule applies

- **Connector Name:** the name of the connector component as it is specified in the configuration file of the currently executed solution
- **Exception Class:** the base class of the exceptions to which the rule applies
- **Regular Expression:** a regular expression that matches the messages of the exceptions to which the rule applies
- **Action:** the action, prescribed by the rule. Can be *error* or *reconnect*.

An error situation is described by the following parts:

- **Connector Class:** the class of the connector that raised the exception
- **Connector Name:** the name of the connector that raised the exception
- **Exception:** the exception raised by the connector – a subclass of `java.lang.Throwable`.

A rule applies to an error situation if all of the following conditions are fulfilled at the same time:

- the rule applies to the connector in the error situation (subclasses of the connector class, described in the rule are also matched)
- the rule applies to the name of the connector which caused the error situation
- the exception is an instance of the exception class, to which the rule applies
- the rule does not have a regular expression to match the exception message or the regular expression matches the message of the exception.

When a given error situation occurs, the reconnect rule engine finds the most specific rule that matches the error situation. First the engine searches through the user-defined rules and if no matching rule is found, it searches the built-in rules. If still no matching rule is found, the engine prescribes the default action, which is "error". If a matching rule is found in the user-defined rules, then the built-in rules are not searched, even if there exists a more specific rule in the built-in rules.

Note: If two or more rules match an error situation, the most specific rule is selected; if there are several most-specific rules and none of them is more specific than the rest, then the first rule in the list is selected. That is why the order of the rules in the rule lists matters. For example: suppose the following rules exist (this is pseudo-syntax used for clarity only):

```
...exceptionClass = "java.io.IOException", exceptionMessageRegExp = ".*", action = "error"...
...exceptionClass = "java.io.IOException", exceptionMessageRegExp = "\w*", action = "reconnect"...
```

If an exception of type `java.io.Exception` with message "problem" is raised, then the first rule is selected, although both rules match the error and no rule is more specific than the other (the outcome of the regular expression match is not considered for weighting purposes.)

Nested Exceptions:

Some exceptions are nested inside other exceptions. When the reconnect rule engine searches through a list of rules (for example the built-in rules), the engine searches for a rule that matches the top-level exception first. If no matching rule is found, then the engine searches again the same list of rules but this time it searches for a match for the nested exception (if the top-level exception has no nested exception this search is skipped). Note that only the first-level nested exception is attempted to be matched by the reconnect rule engine; if there are more levels of nested exceptions they are ignored.

User-defined rules configuration

The list of user-defined rules is configured in a text file named `reconnect.rules` in the "etc" subfolder of the TDI solution folder (or the TDI installation folder, if no solution folder has been defined). Each rule is placed on a single line. The format of a rule is as follows:

```
<connector_class>:<connector_name>:<exception_class>:<action>:<regular_expression>
```

where

- <connector_class> is the fully qualified name of the Java class of the Connector
- <connector_name> is the name of the Connector as inserted in the AssemblyLine
- <exception_class> is the fully qualified name of the Java class of the exception
- <action> can be either 'error' or 'reconnect'
- <regular_expression> is a Java regular expression as described in the JavaDoc of the `java.util.regex.Pattern` class at <http://java.sun.com/j2se/1.5.0/docs/api/java/util/regex/Pattern.html>.

Notes:

1. Each part except the action can be empty. If a part is empty that means "match-all".
2. Each part is mandatory – even if it is empty the surrounding colons must be present. (Consequently on each line there must be at least 4 colons – each colon separating two adjacent parts of the rule. At least 4, because the regular expression may contain colons too. These colons do not interfere with the rule parsing because the regular expression comes last in a rule.)
3. No redundant white space is allowed.
4. The regular expression starts just after the fourth colon and spans until the end of the line.
5. The user-defined rules file is not a Java properties file. The main reason is that a key for a rule must include all rule parts, except the reconnect action, in order to be unique. So the only value from using the Java properties mechanism would be the separation of the action from the other rule parts. However, it would come at the price of escaping white-spaces, colons and equal signs (requirements for a valid property key). Even if the Java property framework was used, custom parsing of the property key would still be required in order to extract the rule parts from it.
6. The regular expression (not the reconnect action) comes last on each line. This pattern is chosen such that it is unnecessary to escape colons (which are considered rule part delimiters) in the regular expression.
7. The regular expression must match the entire message text: Suppose the message text you want to match contains the words "Some Error" somewhere in the message text. A suitable regular expression might then be:

```
.*Some Error.*
```

The character "." matches any character except new line, and the * modifier specifies 0 or more. Now suppose the message ends with a new line. If that is the case, the previous regular expression does not match. You can try a regular expression like this instead:

```
.*Some Error.*\r?\n?
```

"\r" and "\n" specify return and new line characters, and the ? modifier specifies 0 or 1 occurrence.

8. You must still configure reconnect in the Connector's configuration; see "General reconnect configuration" on page 114.

Examples

An example, consisting of two rules:

```
com.ibm.di.connector.ReconnectTestConnector:myconnname:java.io.IOException:error:.*\Wfatal\W.*  
::java.io.IOException:reconnect:
```

Reconnect with the JDBC Connector:

Tivoli Directory Integrator's JDBC connector is configured in Iterator mode to iterate a table from DB2 and is enabled for the reconnect feature. However, at the time of running the solution, DB2 instance is not started yet. In order to have reconnect working, the following exception details need to be mentioned in the `reconnect.rules` file:

```
com.ibm.di.connector.JDBCCConnector::com.ibm.db2.jdbc.DB2Exception:reconnect:
```

Reconnect with the RAC Connector:

Tivoli Directory Integrator's RAC connector is configured in Iterator mode and is enabled for the reconnect feature. In case the Agent Controller server is down, in order for the RAC connector to try to reattempt (reconnect), the following exception details need to be mentioned in the `reconnect.rules` file:

```
com.ibm.di.connector.RACConnector::org.eclipse.tptp.platform.execution.exceptions.AgentControllerUnavailableException:reconnect:
```

Exception considerations

Every environment and solution created for a particular environment using Tivoli Directory Integrator is typically unique. User-defined rules are custom-built and the functionality is made available so solutions can automatically attempt to reconnect based on the exceptions specific to the environment or solution. Refer to the Tivoli Directory Integrator Java API documentation for information about specific exceptions that are returned by the Tivoli Directory Integrator APIs for each Connector.

Additionally, some Tivoli Directory Integrator components rely on underlying libraries and the APIs of these libraries throw exceptions for specific situations. Below we list a few core TDI components where you can look for additional information on exceptions and what may be the cause of the exceptions. This information is helpful when deciding if you want to attempt to create custom reconnect rules for specific exceptions that may be encountered:

- **LDAP Connector** - The LDAP Connector depends on the JNDI libraries shipped with the JRE. For more information on the JNDI interface, its APIs, and the exceptions it may throw, see <http://java.sun.com/j2se/1.5.0/docs/api/javax/naming/package-summary.html>.
- **JDBC Connector** - The JDBC Connector depends on the configured JDBC Driver. The Java API documentation or reference material for the configured JDBC driver should be consulted for more information on the possible exceptions that may be thrown. The "Understanding JDBC Drivers" subsection in the JDBC Connector chapter in *IBM Tivoli Directory Integrator V7.0 Reference Guide* contains links to the JDBC Driver documentation for a set of commonly used JDBC drivers.

General reconnect configuration

Specifying a reconnect rule is necessary for a reconnect to be attempted. However it is not sufficient by itself. The other requirement is enabling reconnect in the general reconnect configuration. This can be done under the **Connection Errors** tab in the Configuration Editor. If reconnect is not enabled in this configuration, reconnect is not attempted in case of error regardless of the decision of the Reconnect Rule Engine. Here is a list of Configuration options:

Number Of Retries

The number of times a reconnect attempt is made when a problem occurs, before giving up. If a new problem occurs later on, the same number of attempts is made.

Delay Between Retries

The number of seconds to wait between each reconnect attempt, and before the first reconnect attempt.

Auto Retry to Connect on Initialize

If this flag is set, and a connection cannot be established when the connector is being initialized, a "reconnect" attempt is made. Not really reconnect, since a connection was not established in the first place, but generally the same mechanism.

Auto Reconnect on Connection Loss

If this flag is set, and the connection is lost after the connector is initialized, a reconnect attempt is made.

Auto Skip Forward

After a reconnect, automatically skip forward as many times as the number of successful reads.

Note: For both the **Auto Retry to Connect on Initialize** and **Auto Reconnect on Connection Loss** flags, the reconnect engine is determine if the exception leads to a reconnect attempt, or is a more general error.

Chapter 8. System Queue

The System Queue is a TDI JMS messaging subsystem similar to the TDI System Store. It facilitates the storing and forwarding of messages between TDI Servers and AssemblyLines. The System Queue simplifies the development of TDI solutions in which asynchronous communication is required to share work amongst multiple AssemblyLines. The System Queue can use either the IBM WebSphere MQ or IBM WebSphere® MQ Everyplace® as its underlying JMS messaging system, as well as any other JMS system provided the JMS Script Driver can properly address this JMS system.

Note: The System Queue Connector (see *IBM Tivoli Directory Integrator V7.0 Reference Guide*) does not talk directly to the System Queue, but rather uses the Server API as an intermediary.

In IBM Tivoli Directory Integrator 7.0, the System Queue is enabled by default by the install process. The MQe Queue Manager is set up in the Solution Directory which was configured during installation. If a new solution directory needs to be created, the MQe Configuration Utility needs to be used to setup the MQ Queue Manager in the new solution directory.

The MQe Queue Manager set up during the installation process has two predefined queues:

- **_default** – serves as a general purposes queue
- **passwords** – this queue is used by the JMS Password Store components for storage of password changes.

This makes the System Queue usable right after the installation of Tivoli Directory Integrator. The sections below discuss the configuration details and the means of changing them, should you want to alter the default configuration.

System Queue Configuration

The System Queue is configured using the following driver-specific Java properties specified in the Tivoli Directory Integrator `global.properties` or `solution.properties` file:

systemqueue.on

This parameter specifies whether the System Queue is to be started and initialized on Tivoli Directory Integrator Server startup. The valid values are `true` and `false`. The default value is `true`.

systemqueue.jmsdriver.name

This parameter specifies the fully qualified name of the Java class to be used as a JMS Driver for the System Queue. This value can be the name of a user-provided class or one of the four standard Tivoli Directory Integrator 7.0 JMS Driver implementations:

- `com.ibm.di.systemqueue.driver.IBMMQe` (“WebSphere MQe parameters” on page 118)
- `com.ibm.di.systemqueue.driver.IBMMQ` (“WebSphere MQ parameters” on page 118)
- `com.ibm.di.systemqueue.driver.IBMMB` (“Microbroker parameters” on page 118)
- `com.ibm.di.systemqueue.driver.JMSScriptDriver` (other JMS system by way of the “JMSScript Driver parameters” on page 119)

The default value is `com.ibm.di.systemqueue.driver.IBMMQe`.

Depending on the `com.ibm.di.systemqueue.driver.IBMMQe` parameter, one of the following sections is applicable:

WebSphere MQe parameters

In order to be able to use MQe as the JMS provider for the System Queue an MQe Queue Manager needs to be created. This can be done using the “MQe Configuration Utility” on page 121 bundled with TDI

systemqueue.jmsdriver.param.mqe.file.ini

This is an MQe-specific parameter that specifies the relative file system file name of the MQe initialization file. This property is required and takes effect only if the MQe JMS driver is specified in the `systemqueue.jmsdriver.name` property. The default value is `MQePWStore/pwstore_server.ini`. This is the default location for the MQe initialization file created by the “MQe Configuration Utility” on page 121.

The system queue is turned on by default, except on z/OS. If you want to use MQe as a system queue you then an abridged enabling procedure is as follows:

1. Set the `systemqueue.on` property in the `global.properties` or `solution.properties` file to `true`.
2. Configure MQe by invoking:

```
cd solution_dir (if using the installation directory, use cd TDI_install_dir)
TDI_install_dir/jars/plugins/mqeconfig.sh
TDI_install_dir/jars/plugins/mqeconfig.props create server (one line)
```

WebSphere MQ parameters

These are WebSphere MQ-specific parameters; for more information about these parameters, see the MQ JMS driver initialization properties in the “System Queue Configuration Example” on page 120 section.

systemqueue.jmsdriver.param.jms.broker

(IP address and TCP port number)

systemqueue.jmsdriver.param.jms.serverChannel

(server channel defined for the MQ server instance)

systemqueue.jmsdriver.param.jms.qManager

(name of the Queue Manager defined for the MQ server instance)

systemqueue.jmsdriver.param.jms.sslCipher

(cipher suite name corresponding to the cipher selected when configuring the MQ server channel, for example, `SSL_RSA_WITH_RC4128_MD5`)

systemqueue.jmsdriver.param.jms.sslUseFlag

(true for SSL connection requested, false if not)

Microbroker parameters

In order to use Microbroker (MB) as the JMS provider for the System Queue, the `systemqueue.jmsdriver.name` property in `global.properties` or `solution.properties` must be set to `com.ibm.di.systemqueue.driver.IBMMB`.

The Microbroker driver has the following parameters (listed here without the “`systemqueue.jmsdriver.param.`” prefix):

jms.broker

the MB server address (IP address and TCP port number); an example value would be “9.126.6.120:1883”

jms.clientID

the client ID; it is required.

Note: In order to be able to use Microbroker as the JMS provider for the System Queue, some Microbroker jars are needed. A sample list of the required jars is available in section External System Configuration, Microbroker of the JMS Connector in *IBM Tivoli Directory Integrator V7.0 Reference Guide*.

JMSScript Driver parameters

The JMS Driver allows you to provide connectivity to any JMS provider through scripting in JavaScript, without writing and building Java code. The JMS Driver acts as a bridge between the System Queue and a user-specified piece of JavaScript, residing on the local file system, which is responsible for creating a `javax.jms.QueueConnectionFactory` object or a `javax.jms.TopicConnectionFactory` object. These objects are obtained in a provider-specific way.

systemqueue.jmsdriver.param.js.jsfile

This is a JMS Script Driver specific parameter (that is, taken into account when the **systemqueue.jmsdriver.name** is set to `com.ibm.di.systemqueue.driver.JMSScriptDriver`) that specifies the name of the file that contains the user-supplied JavaScript code to handle your JMS system of choice. For more information about this parameter, see the JMS driver settings in the “System Queue Configuration Example” on page 120 section. Note that the names of the Java properties do not have the `systemqueue.jmsdriver.param.` prefix.

systemqueue.jmsdriver.param.js.jsscript

The script body which contains JavaScript code for interfacing with the corresponding JMS provider. If this parameter is not provided, then the **systemqueue.jmsdriver.param.js.jsfile** parameter is used for loading the JavaScript to execute.

systemqueue.jmsdriver.param.user.xxxx

These are user-defined properties which are passed by the System Queue to the configured JMS Driver implementation. For example if the following property is set:

```
systemqueue.jmsdriver.param.user.my.prop1=myvalue1
```

the configured JMS Driver get a property with a name of `user.my.prop1` and a value of `myvalue1`.

systemqueue.auth.username

This is the user name used by the System Queue for authentication to the configured JMS system. If this parameter has not been set then the System Queue does not use authentication to the configured JMS system.

systemqueue.auth.password

This is the password used by the System Queue for authentication to the configured JMS system. This parameter is only used when the `systemqueue.auth.username` parameter has been specified.

The env JavaScript object

The piece of JavaScript executed by the JMS Driver needs access to a JavaScript object named *env*. This is an object of type `java.util.Hashtable`, which contains provider-specific parameters for connecting to the JMS provider. These parameters are intended to be used by the JavaScript code in order to access the specific JMS system server instance.

These parameters can be specified in `global.properties` or `solution.properties` using the `systemqueue.jmsdriver.param` prefix. For example, if a URL param is needed for some JMS system, then the following property can be set in `global.properties` or `solution.properties`:

```
systemqueue.jmsdriver.param.myjmssystem.url=myjmsserver.mydomain.com:12345
```

This definition would cause the System Queue to pass it to the JavaScript code as an entry in the *env* `Hashtable`, whose key would be `"myjmssystem.url"` (the System Queue removes the prefix) and whose value would be `"myjmsserver.mydomain.com:12345"`.

The ret JavaScript object

The piece of JavaScript executed by the JMS Driver has access to a JavaScript object named *ret*. This is an object of type `com.ibm.di.systemqueue.driver.JMSScriptDriver.Ret`. It is an instance of the *Ret* inner class of the JMS Script driver class. This *ret* object is to be used to return to the JMS Script driver and eventually to the System Queue the provider-specific objects which the JavaScript code obtains from the JMS system. The *ret* object can also be used to return any error information to the JMS Driver and the System Queue.

This *ret* object has the following members which can be set from JavaScript:

- `queueConnectionFactory` - an object of type `javax.jms.QueueConnectionFactory`. This is the place where the `javax.jms.QueueConnectionFactory` object obtained from the specific JMS system should be stored.
- `topicConnectionFactory` - an object of type `javax.jms.TopicConnectionFactory`. This is the place where the `javax.jms.TopicConnectionFactory` object obtained from the specific JMS system should be stored.
- `errorcode` - an object of type `java.lang.Object`. This is the place where any error information object should be stored. An example of such an object would be a `java.lang.Exception` object.
- `errordescr` - an object of type `java.lang.String`. This is the place where any textual error description should be stored.

JavaScript example for Fiorano MQ

An example configuration and JavaScript code to use the third-party Fiorano MQ system is provided in the *TDI_install_dir/examples* folder, and reproduced below:

```
var ctx = new Packages.java.util.Hashtable();
ctx.put("jms.username", "anonymous");
ctx.put("jms.password", "anonymous");
ctx.put("jms.broker", "http://192.168.113.220:1856");
ctx.put("jms.qManager", "fiorano.jms.runtime.naming.FioranoInitialContextFactory");

var ic = new javax.naming.InitialContext(ctx);

var queueFactory = ic.lookup("primaryQCF");
var topicFactory = ic.lookup("primaryTCF");

ret.queueConnectionFactory = queueFactory;
main.logmsg("driverFiorano.js : QueueConnectionFactory : " + queueFactory);

ret.topicConnectionFactory = topicFactory;
main.logmsg("driverFiorano.js : TopicConnectionFactory : " + topicFactory);
```

Note: This piece of JavaScript demonstrates how the parameters can be hard-coded in the JavaScript code. An alternative is to use the *env* JavaScript object to get any user-supplied parameters from `global.properties` or `solution.properties`. Using the *env* object for parameter retrieval would make changing the configuration easier, because only properties in `global.properties` or `solution.properties` would have to be changed, and no JavaScript code editing would be necessary. This means that users without JavaScript skills would be able to change the configuration.

System Queue Configuration Example

```
##-----
## System Queue settings
##-----
## If set to "true" the System Queue is initialized on startup and can be used;
## otherwise the System Queue is not initialized and cannot be used.
systemqueue.on=true

## Specifies the fully qualified name of the class that is used as a JMS Driver.
# systemqueue.jmsdriver.name=com.ibm.di.systemqueue.driver.JMSScriptDriver
# systemqueue.jmsdriver.name=com.ibm.di.systemqueue.driver.IBMMQe
systemqueue.jmsdriver.name=com.ibm.di.systemqueue.driver.IBMMQ
```

```

### MQe JMS driver initialization properties
## Specifies the location of the MQe initialization file.
## This file is used to initialize MQe on TDI server startup.
# systemqueue.jmsdriver.param.mqe.file.ini=MQePWStore/pwstore_server.ini

### MQ JMS driver initialization properties
systemqueue.jmsdriver.param.jms.broker=192.168.113.54:1414
systemqueue.jmsdriver.param.jms.serverChannel=S_s04win
systemqueue.jmsdriver.param.jms.qManager=QM_s04win
systemqueue.jmsdriver.param.jms.sslCipher=SSL_RSA_WITH_RC4128_MD5
systemqueue.jmsdriver.param.jms.sslUseFlag=true

### JMS Javascript driver initialization properties
## Specifies the location of the script file
# systemqueue.jmsdriver.param.js.jsfile=driver.js

## This is the place to put any JMS provider specific properties needed by a JMS Driver,
## which connects to a 3rd party JMS system.
## All JMS Driver properties should begin with the 'systemqueue.jmsdriver.param.' prefix.
## All properties having this prefix are passes to the JMS Driver on initialization after
## removing the 'systemqueue.jmsdriver.param.' prefix from the property name.
# systemqueue.jmsdriver.param.user.param1=value1
# systemqueue.jmsdriver.param.user.param2=value2
# ...

## Credentials used for authenticating to the target JMS system
# {protect}-systemqueue.auth.username=<username>
# {protect}-systemqueue.auth.password=<password>

```

Security and Authentication

Encryption

Of the standard JMS Drivers, only the driver for MQ supports SSL. The MQe JMS Driver works only with a local Queue Manager – this is mandated by the MQe architecture. The JMS Script Driver is a generic driver which supports whatever the corresponding user-provided JavaScript supports.

Authentication

Some JMS systems, such as WebSphere MQ, can use or even require the use of user name and password authentication. The System Queue provides two standard properties in `global.properties` or `solution.properties` which can be used to configure and supply a user name and password to the System Queue. These properties are `systemqueue.auth.username` and `systemqueue.auth.password`. These two properties are protected by the standard TDI server encrypting of properties which are marked as `{protect}-`. In this way after these properties are set and the TDI server is started the properties' values get encrypted. For more information about these two properties, see the “System Queue Configuration” on page 117 section.

MQe Configuration Utility

The Tivoli Directory Integrator 7.0 MQe Configuration Utility (a command line utility) creates a default MQe Queue when initially setting up the MQe Queue Manager. This default MQe Queue is named “**_default**”. This default Queue is created for convenience only – so that a user can use the MQe Configuration Utility to set up MQe (using the appropriate MQe Configuration Utility command) and then start using the System Queue and the System Queue Connector right away.

IBM Tivoli Directory Integrator 7.0 provides a default configuration file, in `<TDI_install_folder>/MQePWStore/pwstore_server.ini`, the configured parameters of which is used unless you modify them by using this utility.

Additionally the Tivoli Directory Integrator 7.0 MQe Configuration Utility can be used to create and delete user MQe Queues to be used by the System Queue and the System Queue Connector.

Creating an MQe Queue using the MQe Configuration Utility

Typing the following command line creates an MQe Queue named "queue_name" using the mqeconfig.props configuration file:

```
mqeconfig mqeconfig.props create queue queue_name
```

Deleting an MQe Queue using the MQe Configuration Utility

Typing the following command line delete the MQe Queue named "queue_name" using the mqeconfig.props configuration file:

```
mqeconfig mqeconfig.props delete queue queue_name
```

If your solution needs any special configuration, then you can use the MQe Explorer to fine tune your MQe configuration. The MQe Explorer is not bundled with Tivoli Directory Integrator, but can be downloaded as part of the MQe Server Support ES06 pack at http://www-1.ibm.com/support/docview.wss?rs=0&dc=D400&q1=MQe&q2=MQ+Everyplace&uid=swg24007943&loc=en_US&cs=utf-8&ccc=us&lang=en.

Authentication of MQe messages to provide MQe Queue Security

In Tivoli Directory Integrator 7.0 access to MQe can be secured by means of authentication using the MQe Mini-Certificate Server to issue certificates to be used for authentication. For that purpose several additional properties available in Tivoli Directory Integrator 7.0 must be added to the mqeconfig.props properties file, which contains the configuration properties of the MQe Configuration Utility.

The certificates issued by the MQe Mini-Certificate server have a configurable validity period. The default validity period is 12 months. The MQe documentation states that issued certificates should be renewed before the period expires. To enable this, the MQe configuration utility include an option to renew certificates. Typing the following command renews the certificates:

```
mqeconfig mqeconfig.props renewcert {client | server}
```

1. When the last command option is "client", the following values must be set in the mqeconfig.props file:

- **clientRootFolder** - The directory where MQe configuration instance is located.
- **certServerReqPin** - This value is used as a one time authentication PIN for the given authenticatable entity when requesting certificate renewal from the MQe Mini-Certificate server.
- **certServerIPAndPort** - This value is used as the destination address for MQe Mini-Certificate server requests. The format of the value is "FastNetwork:<host>:<port>", where host must be the computer name or TCP IP address or hostname where the MQe Mini-Certificate server is running.
- **certRenewalEntityName** - The MQe authenticatable entity name requiring certificate renewal. Typical entity names include those below, however, any entity name configured in the MQe Mini-Certificate may be used assuming the entity does indeed exist in the queue manager registry referred to by the value of "clientRootFolder":
 - PWStoreClient – client side MQe queue manager
 - PWStoreServer+passwords – remote queue proxy on the client side.

2. When the last command option is "server", the following values must be set in the mqeconfig.props file:

- **serverRootFolder** - The directory where MQe configuration instance is located.
- **certServerReqPin** - This value is used as a one time authentication PIN for the given authenticatable entity when requesting certificate renewal from the MQe Mini-Certificate server.
- **certServerIPAndPort** - This value is used as the destination address for MQe Mini-Certificate server requests. The format of the value is "FastNetwork:<host>:<port>", where host must be the computer name or TCP IP address or hostname where the MQe Mini-Certificate server is running.

- **certRenewalEntityName** - The MQe authenticatable entity name requiring certificate renewal. Typical entity names include those below, however, any entity name configured in the MQe Mini-Certificate may be used assuming the entity does indeed exist in the queue manager registry referred to by the value of "serverRootFolder":
 - PWStoreServer – server side MQe queue manager
 - PWStoreServer+passwords – real queue on the server side.

Support for DNS names in the configuration of the MQe Queue

There is no additional coding required to support this feature. It should be noted that DNS support is really an MQe feature, since the TDI component implementations simply pass the configuration properties from mqeconfig.props through to the MQe APIs. Themqeconfig.props properties which can accept DNS name or IP address values are:

- serverIP
- certServerIPAndPort

Configuration of High Availability for MQe transport of password changes

To support high availability deployments, you have the possibility to deploy and configure multiple instances of the TDI MQe components. In some deployments, it may be necessary to configure multiple TDI MQe Password Store components. For example, if password change plug-ins have been configured for multiple Windows Domain Controllers—in this case, it is likely that there separate instances of MQe client side Queue Managers with the name "PWStoreClient". Additionally, for each of the client Queue Managers, there is a remote queue proxy connection to the MQe server side Queue Manager queue used by the TDI MQe Password Connector. The remote queue proxy name is "PWStoreServer+passwords". When you use this type of deployment scenario, the authentication certificates associated with these two MQe entities (that is, "PWStoreClient", "PWStoreServer+passwords") is requested and issued multiple times. This happens each time the mqeconfig utility is executed. Before executing the second and each subsequent instances of the mqeconfig utility, it necessary to re-enable certificate issue for each of the MQe entities mentioned above.

For some deployments, you may prefer to configure the TDI MQe Password Connector such that it supports a particular high availability requirement. You may expect that an implementation supporting this type of requirement would employ multiple instances of the TDI MQe Password Connector, each with its own associated MQe Queue Manager configuration. In this case you would deploy multiple identical MQe server side configurations, allowing a network load balancer to route requests from the TDI MQe Password Store client to an available server instance. Each MQe Queue Manager on the server side is configured using the mqeconfig utility. When this utility executes it automatically request authentication certificates from the MQe Mini-Certificate server for the entities named "PWStoreServer" and "PWStoreServer+passwords". These represent the Queue Manager and Queue names respectively. Before executing the second and each subsequent instance of the mqeconfig utility, it necessary to re-enable certificate issue for the two MQe entities mentioned above.

Providing remote configuration capabilities in the MQe Configuration Utility

Creating a remote MQe Queue using the MQe Configuration Utility

Typing the following command line create a remote MQe Queue named "queue_name" using the mqeconfig.props configuration file:

```
mqeconfig mqeconfig.props create remotequeue queue_name targetQMname [QM_ip_or_hostname comm_port]
```

In the above command line QM_ip_or_hostname and comm_port parameters are optional; if they are missing only a remote queue definition is created. If you provide these two parameters, a Connection definition also be created before creating the remote queue definition.

Note: A remote queue is not usable without a Connection definition. In addition several remote queues can be defined to share a single Connection. The targetQMname parameter specifies the name of the remote MQe Queue Manager.

Deleting a remote MQe Queue using the MQe Configuration Utility

Typing the following command line delete a remote MQe Queue named "queue_name" using the mqeconfig.props configuration file:

```
mqeconfig mqeconfig.props delete remotequeue queue_name targetQMname
```

In the above command line the targetQMname parameter specifies the name of the remote MQe Queue Manager.

Chapter 9. Encryption and FIPS mode

To provide confidentiality of data, IBM Tivoli Directory Integrator 7.0 can encrypt:

- configuration files
- property values in properties files
- server user registry files
- JavaScript files

Encryption is the process of selecting some humanly readable text, called *plaintext*, and hiding its content and meaning to make the data in the plaintext format more secure. Plaintext is written in lowercase letters. Encrypted text is called *ciphertext*. Ciphertext is written in capital letters.

Note: In Config files, if the `{protect}-` prefix precedes the name of a property, then the property value is, or should be, encrypted. The prefix, `{protect}-` is optional. The values that are already encrypted values start with `{encr}`.

See “Working with encrypted TDI configuration files” on page 94 and “Encryption of properties in external property files” on page 96.

For example:

```
[{protect}-]keyword <colon | equals> [{encr}][{java}]value
```

The `{java}` value must be b64-encoded. For example:

```
{protect}-api.truststore.pass={encr}J8AKimpEutu3Bb10Vg55F/5d5v02kXWcNUWnCq3vINUc6K0719z9dEk3H430t2iTT1dZTI6FSSV  
in9KsCyBLmgv+n84w7He1K13ro2dFmZbTYKMxux6oqN9nL2V0vZoptNqzoWvs6IN/p3VkiIBt1ao/9mEPEKuiWrNkTkQ89Bg=
```

Configuring Tivoli Directory Integrator to run FIPS mode

The Federal Information Processing Standard (FIPS) Publication 140-2, FIPS PUB 140-2, is a U.S. government computer security standard used to accredit cryptographic modules.

When the Tivoli Directory Integrator server is configured to run in FIPS mode, that Server is using the FIPS 140-2 certified cryptographic modules. Tivoli Directory Integrator does not generate cryptographic keys – keys are created using external utilities such as *keytool* (see “Creating a self-signed certificate using keytool” on page 68) and *Ikeyman* (see “Creating a self-signed certificate using Ikeyman” on page 69). For information on Tivoli Directory Integrator use of encryption, see Chapter 6, “Security and TDI,” on page 63. In order to create, edit, export and overall manage keystores and truststores the Ikeyman GUI utility or the keytool command line utility can be used. The executable file, keytool.exe is found in `root_directory/jvm/jre/bin`, or `root_directory/jvm/bin`, depending on your platform.

Symmetric cipher support

Symmetric cipher support in Tivoli Directory Integrator is one of the requirements for achieving a FIPS 140-2 compliance.

The reason for encrypting a message is to change the message into a meaningless form of text called cipher text that is meaningless to whoever intercepts the message. There are many different encryption algorithms called ciphers. One of the most widely known ciphers is the *symmetric* cipher. The symmetric cipher has a key that both the sender and the receiver can keep. The sender uses that key to encrypt the message. The receiver uses the same key to decrypt the message.

An optional configuration is provided to use a symmetric cipher (specifically, the Advanced Encryption Standard, or AES). The symmetric cipher encoded using AES allows customers that need FIPS-compliant solutions to use a supported cipher around encryption.

The following property defines the cipher:

```
com.ibm.di.securityTransformation=DES/ECB/NoPadding
```

This property defines a cipher for the password-based encryption or decryption of Tivoli Directory Integrator configurations.

FIPS encryption

You can run Tivoli Directory Integrator and the Tivoli Directory Integrator server in a secure way using FIPS. You can also configure additional properties when you want to operate Tivoli Directory Integrator in a specific mode, for example, FIPS mode.

Connectors, Function Components, Parsers: FIPS 140-2 is concerned only with cryptographic functionality such as SSL, digital signing, encryption, cryptographic hashing and random number generation.

SSL:

FIPS 140-2 requires TLS to be the protocol for SSL communication. SSLv3 and its predecessors are not allowed. When FIPS mode is turned on, Tivoli Directory Integrator components that use SSL will fail to communicate with external systems that do not support TLS.

JDBC and the System Store:

The DB2 Type 4 JDBC driver (`com.ibm.db2.jcc.DB2Driver`) that is shipped with Tivoli Directory Integrator, supports SSL in a FIPS conformant way.

The Apache Derby drivers, network and embedded, do not support SSL in version 10.2 (which is the one bundled with Tivoli Directory Integrator).

However, the Apache Derby 10.2 database engine can perform database encryption. By default Tivoli Directory Integrator uses Derby for its System Store. If you use database encryption functionality of Derby in FIPS mode, be sure to specify the IBM certified cryptographic provider `IBMJCEFIPS` as the provider used for encryption and also choose a FIPS approved encryption cipher. Here is an example of how to configure the System Store to use Derby with FIPS compliant database encryption:

```
com.ibm.di.store.database=jdbc:derby://localhost:1527/C:\TDI\TDISysStoreEnc;create=true;
  dataEncryption=true;encryptionKey=c566bab9ee8b62a5ddb4d9229224c678;encryptionAlgorithm=AES/CBC/NoPadding;
  encryptionProvider=com.ibm.crypto.fips.provider.IBMJCEFIPS
com.ibm.di.store.jdbc.driver=org.apache.derby.jdbc.ClientDriver
com.ibm.di.store.jdbc.urlprefix=jdbc:derby:
com.ibm.di.store.jdbc.user=APP
```

JMS and the System Queue:

MQe Mini-Certificates involve cryptography that is not FIPS compliant, so this security feature of MQe should not be used in FIPS mode.

The WebSphere MQ 5.3 JMS provider is not capable of running SSL in a FIPS compliant mode. In FIPS mode SSL should not be used with that provider.

The WebSphere MQ 6.0 JMS provider can use SSL in a FIPS compliant mode. To take advantage of it, however, you need MQ 6.0.1.0 or higher, because in earlier versions of MQ 6.0 the FIPS conformant mode does not work properly with Java 5.0 that Tivoli Directory Integrator uses. To use FIPS compliant SSL communications between Tivoli Directory Integrator and WebSphere MQ:

1. Ensure that the WebSphere MQ installation is of version 6.0.1.0 or higher.
2. Ensure that the corresponding Queue Manager on the MQ side requires FIPS compliant SSL communications.
3. Ensure that the corresponding SSL channel of the Queue Manager uses a FIPS compliant SSL Cipher Spec.
4. Turn on FIPS mode for Tivoli Directory Integrator. When FIPS mode is enabled for Tivoli Directory Integrator, it automatically enables FIPS mode on all JMS SSL connections to WebSphere MQ.
5. Copy the JMS client jars from the WebSphere MQ installation to TDI; refer to the JMS Connector documentation in *IBM Tivoli Directory Integrator V7.0 Reference Guide* for a list of necessary client libraries for MQ 6.0 and how to deploy them in TDI.
6. On the Tivoli Directory Integrator side, configure a FIPS compliant SSL Cipher Suite that is compatible with the SSL Cipher Spec configured on the SSL channel of the MQ Queue Manager. You can do this using the **jms.sslCipher** parameter of the JMS Connector and the `systemqueue.jmsdriver.param.jms.sslCipher` system property of the MQ driver for the System Queue. You can find a SSL Cipher Spec to Cipher Suite mapping and their FIPS compliance here: http://publib.boulder.ibm.com/infocenter/wmqv6/v6r0/index.jsp?topic=/com.ibm.mq.csqzaw.doc/uj34740_.htm.

The TDI server and FIPS: When run in this mode the Tivoli Directory Integrator Server is forced to use FIPS 140-2 cryptographic modules

Note: If the Server is running with FIPS and SSL enabled, then do not use clients with SSL for secure sockets communication. In this case the Server uses TLS and a connection will not succeed. Instead of using SSL make sure you are using TLS like the Server does for secure sockets communication.

Running the Tivoli Directory Integrator Server in FIPS mode has the following implications:

- Only FIPS compliant crypto algorithms are allowed for encryption and decryption of configurations, properties, and so forth.
- Auxiliary tools which use encryption/decryption should be used in FIPS compliant way - Ikeyman, createtash, cryptoutils, skeytol, and so forth.
- Components will not be able to communicate with external systems that do not use TLS for socket communication.
- Some components should not be used when the Server is in FIPS mode because they will break the FIPS compliancy. Refer to Table 15 on page 128 for a list detailing component compliance.

Enabling FIPS mode: When using FIPS, many Tivoli Directory Integrator configuration options are changed, so you must keep in mind several rules in order to maintain FIPS compliancy. Some of the rules are mentioned in this document and others can be found in <https://w3.webahead.ibm.com/w3ki/download/attachments/370821/FIPS+140+Guidelines.pdf?version=1> and <http://www-128.ibm.com/developerworks/java/jdk/security/50/FIPShowto.html>.

Enabling FIPS mode in TDI:

1. Set the `com.ibm.di.server.fipsmode.on` property to **true** in `global.properties` or `solution.properties`.
2. Make sure the `com.ibm.di.securityTransformation` property value is in an algorithm which is FIPS compliant, for example, AES/ECB/NoPadding. This algorithm is used when you attempt to open an encrypted configuration.
3. Hardware cryptography can not be used along with SSL in FIPS mode.

The underlying SSL module – IBMJSSE2 does not support hardware cryptography in FIPS mode as stated here: <http://www-128.ibm.com/developerworks/java/jdk/security/50/secguides/jsse2Docs/JSSE2RefGuide.html#runfips>. You cannot use hardware-based SSL keys for the Server API in FIPS mode; the `com.ibm.di.server.pkcs11` property must be absent or set to false in `global.properties` and `solution.properties`.

4. Make sure Server encryption uses a transformation that is FIPS 140-2 compliant.

By default the Server uses public key encryption with the RSA algorithm. However, the RSA encryption option is not compliant with FIPS 140-2. That is why you must manually configure another cryptographic transformation that is FIPS allowed. Below are sample steps that setup Tivoli Directory Integrator to use the AES cipher for encryption:

- Generate an AES secret key and put it in a keystore. This can be done using the “The TDI secret key tool” on page 98 utility located in the bin folder of a Tivoli Directory Integrator installation like this:

```
skeytool -genseckey -alias server -keyalg AES -keysize 128 -keystore server.jck -storepass mypass  
-storetype jceks -keypass mykeypass -keygenproviderclass com.ibm.crypto.fips.provider.IBMJCEFIPS
```

This command creates a new keystore file `server.jck` of type JCEKS (JKS keystores cannot host secret keys) with an AES key of size 128 under alias `server`. The password for the created keystore is `mypass`. Pay special attention to the `keygenproviderclass` parameter – it is absolutely necessary to specify the FIPS certified provider if you strive for FIPS 140-2 compliance. Note that this is just an example, you can use whatever file names, passwords and aliases you wish.

- Change the Tivoli Directory Integrator settings to use secret key encryption with the newly generated key. For example in `global.properties` or `solution.properties` file, set the following properties:

```
com.ibm.di.server.encryption.keystore=server.jck  
com.ibm.di.server.encryption.keystoretype=jceks  
com.ibm.di.server.encryption.key.alias=server  
com.ibm.di.server.encryption.transformation=AES/CBC/PKCS5Padding
```

- Migrate all existing files that have been encrypted with the old key:

All encrypted files that existed prior to the introduction of the new key, need to be migrated. Migration involves decryption with the old key and (optionally) re-encryption with the new one (see “Maintaining encryption artifacts – keys, certificates, keystores, encrypted files” on page 137). For example you can migrate `global.properties` as follows:

```
cryptoutils -input ../etc/global.properties -output ../etc/global.properties  
-mode decrypt_props -keystore ../testserver.jks -storepass server -alias server  
-transformation RSA -storetype jks -keypass server
```

```
cryptoutils -input ../etc/global.properties -output ../etc/global.properties  
-mode encrypt_props -keystore ../server.jck -storepass mypass -alias server  
-transformation AES/CBC/PKCS5Padding -storetype jceks -keypass mykeypass
```

- Regenerate the Tivoli Directory Integrator Server “Stash File” on page 93 to reflect the new passwords of the encryption keystore and the encryption key. This is done using the `createstash` utility found in the bin folder of a Tivoli Directory Integrator installation. For example:

```
createstash mypass mykeypass com.ibm.crypto.fips.provider.IBMJCEFIPS
```
- Use only FIPS compatible Tivoli Directory Integrator components in your solutions, as listed in the table below:

Table 15. FIPS compatible components

Directory Integrator Component	Allowed in FIPS mode?	Remarks
Connectors		
ACT Connector	yes	Operates as a Server API client
Active Directory Change Detection Connector	yes	Uses default JSSE factories for SSL
AssemblyLine Connector	yes	Operates as a Server API client
Axis Easy Web Service Server Connector	yes	Uses default JSSE factories for SSL
Command line Connector	yes	Provides no cryptography features

Table 15. FIPS compatible components (continued)

Directory Integrator Component	Allowed in FIPS mode?	Remarks
Domino/Lotus Notes Connectors	no	Domino/Notes 7 cryptographic capabilities are not FIPS conformant. (Some FIPS enablement may be included in Notes 8.0.1.)
ITIM DSMLv2 Connector	yes	Uses default JSSE factories for SSL
DSMLv2 SOAP Connector	yes	Uses default JSSE factories for SSL
DSMLv2 SOAP Server Connector	yes	Uses default JSSE factories for SSL
Exchange Changelog Connector	yes	Uses default JSSE factories for SSL
File system Connector	yes	Provides no cryptography features
FTP Client Connector	yes	Provides no cryptography features
GLA Connector	yes	Provides no cryptography features
HTTP Client Connector	yes	Uses default JSSE factories for SSL
Old HTTP Client Connector	yes	Uses default JSSE factories for SSL
HTTP Server Connector	yes	Uses default JSSE factories for SSL
Old HTTP Server Connector	yes	Provides no cryptography features
IBM Directory Server Changelog Connector	yes	Uses default JSSE factories for SSL
ITIM Agent Connector	yes	Provides no cryptography features
JDBC Connector	depends	If no cryptography features are used (SSL, encryption), the Connector is FIPS conformant. Otherwise FIPS conformance depends on the FIPS conformance of the cryptographic functionality of the JDBC driver that is used. See “Connectors, Function Components, Parsers” on page 126 for a discussion on the FIPS conformance of JDBC drivers.
JMS Connector	depends	If no cryptography features are used (SSL, encryption), the Connector is FIPS conformant. Otherwise FIPS conformance depends on the FIPS conformance of the cryptographic functionality of the JDBC driver that is used. See “Connectors, Function Components, Parsers” on page 126 for a discussion on the FIPS conformance of JMS providers.
JMX Connector	yes	Provides no cryptography features
JNDI Connector	yes	Uses default JSSE factories for SSL
LDAP Connector	yes	Uses default JSSE factories for SSL
LDAP Server Connector	yes	Uses default JSSE factories for SSL
Mailbox Connector	yes	Uses default JSSE factories for SSL
Memory Queue Connector	depends	Depends on the FIPS compliance of the JDBC driver used for the System Store. (The Memory Queue uses the System Store for persistence.) See “Connectors, Function Components, Parsers” on page 126 for a discussion on the FIPS conformance of JDBC drivers.
Memory Stream Connector	yes	Provides no cryptography features

Table 15. FIPS compatible components (continued)

Directory Integrator Component	Allowed in FIPS mode?	Remarks
MQe Password Store Connector	depends	Only PKCS#7 is allowed in FIPS mode for message protection. The RSA encryption option must not be used. The MQe Mini Certificates are not FIPS compliant, so they must not be used in FIPS mode.
Netscape/iPlanet/Sun Directory Changelog Connector	yes	Uses default JSSE factories for SSL
Properties Connector	depends	If encryption is turned off, the Connector is FIPS conformant. Otherwise FIPS conformance depends on the cipher used for encryption. An example of a FIPS 140-2 approved cipher is AES. Other approved ciphers can be found at: http://csrc.nist.gov/publications/fips/fips140-2/fips1402annexa.pdf The Server encryption option will always be FIPS conformant as long as TDI is configured correctly for FIPS mode. (See “Enabling FIPS mode” on page 127.)
Server Notifications Connector	yes	Operates as a Server API client
System Queue Connector	depends	If no cryptography features are used by the System Queue (SSL, encryption), the Connector is FIPS conformant. Otherwise FIPS conformance depends on the FIPS conformance of the JMS provider that is used by the System Queue. See “Connectors, Function Components, Parsers” on page 126 for a discussion on the FIPS conformance of JMS providers.
Windows Users and Groups Connector	yes	Provides no cryptography features
System Store Connector	depends	Depends on the FIPS compliance of the JDBC driver used by the System Store.
RAC Connector	yes	Provides no cryptography features
RDBMS Changelog Connector	depends	Same as the JDBC Connector
SNMP Connector	yes	Provides no cryptography features
SNMP Server Connector	yes	Provides no cryptography features
TAM Connector	yes	Tivoli Access Manager Runtime for Java is FIPS conformant
TCP Connector	yes	Uses default JSSE factories for SSL
TCP Server Connector	yes	Uses default JSSE factories for SSL
Timer Connector	yes	Provides no cryptography features
URL Connector	yes	Provides no cryptography features
Web Service Receiver Server Connector	yes	Uses default JSSE factories for SSL
z/OS Changelog Connector	yes	Uses default JSSE factories for SSL
Function Components		
Castor Java to XML FC	yes	Provides no cryptography features

Table 15. FIPS compatible components (continued)

Directory Integrator Component	Allowed in FIPS mode?	Remarks
Castor XML to Java FC	yes	Provides no cryptography features
EMF XMLToSDO	yes	Provides no cryptography features
EMF SDOToXML	yes	Provides no cryptography features
AssemblyLine FC	yes	Operates as a Server API client
Java Class Function Component	depends	Depends on the FIPS compliance of the Java class, whose method will be invoked by the Function Component. If the Java class does not use cryptography (SSL, encryption, signing, cryptographic hash functions, and so forth) it can be safely used in FIPS mode.
Parser FC	depends	Depends on the FIPS compliance of the Parser that is configured for the Function Component
CBE Generator Function Component	yes	Provides no cryptography features
SendEmail Function Component	yes	Uses default JSSE factories for SSL
Memory Queue FC	depends	Depends on the FIPS compliance of the JDBC driver used by the System Store. (The Memory Queue uses the System Store for persistence.) See “Connectors, Function Components, Parsers” on page 126 for a discussion on the FIPS conformance of JDBC drivers.
Axis Java To Soap FC	yes	Provides no cryptography features
WrapSoap FC	yes	Provides no cryptography features
Invoke Soap WS FC	yes	Uses default JSSE factories for SSL
Axis Soap To Java FC	yes	Provides no cryptography features
Axis EasyInvoke Soap WS FC	yes	Uses default JSSE factories for SSL
Complex Types Generator Function Component	yes	Provides no cryptography features
Remote Command Line Function Component	depends	The cryptographic capabilities of the RXA toolkit are not FIPS compliant. If no cryptography features are used, the component can be used in FIPS mode.
z/OS TSO/E Command Line FC	depends	Depends on the FIPS compliance of the cryptography involved in the TSO command that is invoked by the Function Component
SAP ABAP Application Server Component Suite	no	The SAP cryptographic module has not been FIPS 140-2 certified. If no cryptography features are used, the components can be used in FIPS mode.
Parsers	yes	None of the Tivoli Directory Integrator Parser components use cryptography so all of them can be used in FIPS mode.

Setting com.ibm.di.server.fipsmode.on:

To enable FIPS mode in Tivoli Directory Integrator you must specify it in a property in `global.properties` or `solution.properties`. The property is named `com.ibm.di.server.fipsmode.on` and

can be set to either **true** or **false**. When this property is set to **true**, the TDI Server runs in FIPS mode. In this mode, the IBM FIPS security provider is set in the Tivoli Directory Integrator JVM before the IBM JCE security provider in the providers list. When the TDI FIPS enabling property is true, it also enables FIPS mode in the IBM JSSE2 provider and sets the default JSSE SSL socket factories to be the ones from the IBM JSSE2 provider. By default FIPS mode is not enabled in TDI, that is, the `com.ibm.di.server.fipsmode.on` property is set to **false**.

Using crypto algorithms in FIPS mode:

Only FIPS-compliant crypto algorithms can be used. This means that you must use only FIPS-compliant algorithms in order to stay in FIPS-compliant mode. Using other algorithms violates FIPS compliancy.

Setting `com.ibm.di.securityTransformation`:

When opening an encrypted configuration, TDI uses the `com.ibm.di.securityTransformation` property to get the algorithm that decrypts the configuration. If this property is set to an algorithm which is not FIPS-compliant, and the Tivoli Directory Integrator Server FIPS mode is turned on, then an Exception is thrown. The Exception message which is shown would look something like this:

```
CTGDIC012E Could not load file<FILE_PATH>. No such algorithm: <ALGORITHM_NAME>.
```

In order to avoid this Exception, always set FIPS-compliant algorithms for this property when running in FIPS mode. By default the `com.ibm.di.securityTransformation` property is set to `DES/ECB/Nopadding` which is **not** a FIPS-compliant algorithm. This property also defines a cipher for the password-based encryption and decryption of TDI configurations.

Setting properties automatically when running in FIPS mode:

- Tivoli Directory Integrator sets a relevant System property which is not present in the `global.properties` file by default. This property is called `com.ibm.di.cryptoProvider` and is set to the `IBMJCEFIPS` security provider when run in FIPS mode. You should note that if this property is set in `global.properties` then that particular value is used; if this property is set to a non-FIPS compliant provider, then even if TDI is run in FIPS mode, TDI is **not** FIPS-compliant.
- When in FIPS mode, specific JSSE Socket Factories are used. These are the `IBMJSSE2` Socket Factories. This is done automatically by the TDI Server which sets the `ssl.SocketFactory.provider` and the `ssl.ServerSocketFactory.provider` properties to the JSSE implementation classes in `IBMJSSE2` provider.

Using the create stash file command line tool in FIPS mode:

To create a stash file that conforms to FIPS 140-2 standards you must provide the `IBMJCEFIPS` provider class as the third parameter when using the `createstash` file tool. For example:

```
TDI_install_dir\bin\createstash Password Password com.ibm.crypto.fips.provider.IBMJCEFIPS
```

Using alternatives to RSA encryption in FIPS mode:

In FIPS mode, configure Tivoli Directory Integrator to use the Advanced Encryption Standard (AES) instead of the RSA encryption algorithm. A secret key cipher that is FIPS 140-2 compliant is required. As an acronym, RSA stands for Rivest, Shamir, and Adelman, the inventors of the algorithm. The RSA algorithm is a strong encryption algorithm used for sending data over the internet. The RSA cipher is allowed only to encrypt and decrypt keys for transport (SSL, TLS) to stay within the boundaries of the Approved Mode of FIPS 140-2 Level 1, as stated at: <http://www-128.ibm.com/developerworks/java/jdk/security/50/FIPShowto.html>.

Running auxiliary tools in FIPS mode: This section provides command line syntax for identifying the appropriate crypto provider, and when generating a secret key.

createstash:

Pass the FIPS 140-2 certified crypto provider IBMJCEFIPS as an explicit provider parameter on the command-line:

```
createstash mypass mykeypass com.ibm.crypto.fips.provider.IBMJCEFIPS
```

cryptoutils:

Pass the FIPS 140-2 certified crypto provider IBMJCEFIPS as an explicit provider on the command-line using the *cryptoproviderclass* option like this:

```
cryptoutils -input registry.txt -output registry.enc -mode encrypt -keystore ../testserver.jks -storepass server  
-alias server -cryptoproviderclass com.ibm.crypto.fips.provider.IBMJCEFIPS
```

skeytool:

Special care is necessary only when generating a secret key. In that case you should pass the FIPS 140-2 certified crypto provider IBMJCEFIPS as an explicit provider on the command-line using the *cryptoproviderclass* option like this:

```
skeytool -gensecretkey -alias mysecretkey -keyalg AES -keysize 128 -keystore ks.jck -storepass secret  
-storetype jceks -keygenproviderclass com.ibm.crypto.fips.provider.IBMJCEFIPS
```

Configuring FIPS properties for TDI: **Running keytool/Ikeyman in FIPS mode:**

To use the keytool and Ikeyman utilities in FIPS mode, edit the *java.security* file in *TDI_install_dir/jvm/jre/lib/security*. In the first two lines in the *java.security* file, set the IBMJCEFIPS provider first and the IBMJCE security provider second. For example:

```
security.provider.1=com.ibm.crypto.fips.provider.IBMJCEFIPS  
security.provider.2=com.ibm.crypto.provider.IBMJCE
```

However, on Solaris and HP-UX, the SUN provider should always be the first provider in the security providers list.

Configuring SSL and PKI certificates

Tivoli Directory Integrator uses both Secure Socket Layer (SSL) and Public Key Infrastructure (PKI) encryption methods. SSL and PKI provides an important foundation for many of the TDI and TDI server features. SSL provides for encryption and authentication of network traffic between two remote communicating parties. Similarly, PKI (public key infrastructure) enables users of unsecured networks to securely and privately exchange data by using a public and a private cryptographic key pair that is obtained and shared through a trusted authority. See “Working with certificates” on page 134.

SSL certificate:

An SSL certificate resides on a secure server and is used to encrypt the data that identifies the server. The SSL certificate helps to prove the site belongs to the entity who claims it and contains information about the certificate holder, the domain that the certificate was issued to, the name of the Certificate Authority who issued the certificate, and the root and the country it was issued in.

PKI certificate:

A PKI certificate enables users of an unsecured network to add security and privacy to data exchanges. PKI uses a cryptographic key pair that it gets and shares through a trusted authority called a Certificate Authority (CA). Using PKI, you can obtain a certificate that can identify an individual or an organization and directory services that can store the certificates. The CA can also revoke the certificates when necessary. The most common use of a digital certificate is to verify that a user sending a message is who the sender claims to be, and to provide the receiver with the encryption of the reply.

Follow these steps to provide separate configuration options for certificates to be used for PKI Encryption and SSL:

1. Add the following properties:

```
com.ibm.di.server.encryption.keystore  
com.ibm.di.server.encryption.key.alias  
api.keystore.password  
api.key.password
```

2. Rename the following properties as shown:

```
com.ibm.di.server.keystore -----> api.keystore  
com.ibm.di.server.key.alias -----> api.key.alias
```

Note: The `disrv.sth` file now holds the password only for the encryption file.

Encrypting and decrypting using CryptoUtils

Using Tivoli Directory Integrator, you can PKI encrypt sensitive properties in the `global.properties` or in the `solution.properties` file. One method of decrypting PKI-encrypted properties is to use the Configuration Editor (CE) properties editor. The `CryptoUtils` command-line utility is another method of decrypting PKI-encrypted properties files. Decryption requires you to give your PKI credentials so that unauthorized users cannot access sensitive information. You can properties files that contain PKI encrypted properties using `CryptoUtils`. see “The TDI Encryption utility” on page 96.

Working with certificates

Someone who wishes to send an encrypted message applies for a digital certificate from a CA. The CA issues an encrypted digital certificate that contains the applicant's public key and a other identification data. The CA makes reveals its own public key through printed media or perhaps on the Internet. The recipient of an encrypted message uses the CA's public key to decode the digital certificate attached to the message, verifies the certificate as issued by the CA, and then gets the sender's public key and identification data from the certificate. With this information, the recipient can send an encrypted reply.

Digital certificates are of two types:

- CA-signed certificates
- Self-signed certificates

CA-signed certificates are signed by a Certificate Authority such as VeriSign and thawte. A self-signed certificate is an identity certificate that is signed by its own creator. See “Managing a CA-signed certificate using keytool” on page 68 and “Creating a self-signed certificate using keytool” on page 68.

Comparing CA-signed and Self-signed certificates

Certificate authority signed certificates:

Certificate authorities such as VeriSign require a procedure whereby applicants can prove their identities and obtain certificates that authenticate both the identity of the certificate applicants and its own identity as a signer of a certificate.

Self-signed certificates:

Typically there is a local certification authority (CA), that is, the certificates do not come from any of the well known CAs like VeriSign, and so on. The local CA itself should have a root certificate issued by a well-known CA, but even this is not always true. If the local CA's root certificate is self-signed, you must import it into the truststore of each server or client that is using SSL.

In this case, each server for an SSL connection, and each client doing PKI authentication, generates its own self-signed certificate. It is then necessary to export the certificate to a file and to import it into various truststores. If a client **C** connects to a server **S**, **C** must have **S**'s self-signed certificate in its

truststore. If a client C does PKI authentication (symmetric SSL) to a server S, S must have C's self-signed certificate in its truststore. Note: Self-signed certificates can be used for either a client or a server certificate. See the "Keystore and truststore management" on page 68 on information how to do this. Each server for an SSL connection and each client doing PKI authentication must then issue a request for a certificate to the local CA, and must add the resulting certificate into its keystore.

Configuring certificates using PKI and SSL

Tivoli Directory Integrator provides separate configuration options for certificates to be used for public key infrastructure (PKI) encryption and Secure Socket Layer (SSL) connection. Independent configuration of PKI and SSL certificates allows you to migrate your encrypted properties separately from the process of upgrading your SSL certificates.

Under PKI, a Certificate Authority (CA) binds public keys to user identities. The user identity must be unique for each CA. Public Key certificates collect each user, user identity, public key, their binding, validity conditions, and other attributes that are made unforgetable in public key certificates issued by the CA.

The certificates used for SSL may expire, or for security reasons, SSL certificates may have to be refreshed frequently. Certificates used for PKI encryption can be persisted longer than it is appropriate to persist SSL certificates. PKI certificates should be maintained in case there is data that has been encrypted using the public key certificate. As a result, Tivoli Directory Integrator allows you to configure PKI and SSL certificates separately. Each server for an SSL connection and each client performing PKI authentication must issue a request for a certificate to the local CA, and must add the resulting certificate into its keystore.

These properties are added to the `global.properties` file:

```
com.ibm.di.server.encryption.keystore  
com.ibm.di.server.encryption.key.alias
```

These properties variables are set to the same values as the ones already in `global.properties`:

```
api.keystore=truststore  
api.key.alias=server
```

Using cryptographic keys located on hardware devices

The RSA signing and encryption algorithm (developed by Ron Rivest, Adi Shamir, and Leonard Adleman) is a well-known public key cipher. RSA Laboratories (Part of EMC Corp.) have published the PKCS#11 standard, which defines a platform-independent API to hardware cryptographic tokens, such as Hardware Security Modules and smart cards. The PKCS#11 API defines most commonly used cryptographic object types, including:

- RSA keys
- X.509 Certificates
- Data Encryption Standard (DES)DES/Triple DES keys
- All the functions required for using, creating or generating, modifying, and deleting the above objects

Public-Key Cryptography Standards (PKCS) PKCS#11 is a standard that provides a common application interface to cryptographic services on various platforms using various hardware cryptographic devices. Hardware Cryptographic key storage devices allow keys to be stored on hardware devices. IBM Tivoli Directory Integrator supports private keys and certificates on crypto devices that are PKCS#11 compliant. Support is provided on all hardware devices supported by the IBM Java PKCS libraries shipped with the IBM Java Runtime Environment (JRE). PKCS standards are a set of common protocols that allow secure information exchange over networks using a public key infrastructure (PKI). IBM Tivoli Directory Integrator can store Secure Socket Layer (SSL) keys on the hardware devices. For the requirement to store keys on hardware devices, the following new properties are available in the `global.properties` file:

```

##PKCS11 options
##Set the value of following properties to use PKCS11 enabled devices to store TDI servers private key /
##certificate.
com.ibm.di.pkcs11cfg=etc\pkcs11.cfg
com.ibm.di.server.pkcs11=false
com.ibm.di.server.pkcs11.library=
com.ibm.di.server.pkcs11.slot=
{protect}-com.ibm.di.server.pkcs11.password=PASSWORD

```

The default value of the property `com.ibm.di.server.pkcs11` is `false`. The value corresponding to the property `com.ibm.di.server.pkcs11.password` is encrypted.

Using IBMPCKS11 to access devices and to store SSL keys and certificates

IBM Tivoli Directory Integrator uses IBMPCKS11 to access crypto hardware devices that store the SSL keys and certificates. Support is provided for all hardware devices supported by the IBM Java PKCS libraries and shipped with the IBM JRE.

Table 16. SSL supported properties

Property	Default value	Description
<code>com.ibm.di.pkcs11.cfg</code>	<code>etc\pkcs11.cfg</code>	Use CFG file to point to the path of the configuration file required to initialize the IBM PKCS11 implementation provider.
<code>com.ibm.di.server.pkcs11</code>	<code>false</code>	Use PKCS#11 compliant crypto devices for ssl.
<code>com.ibm.di.server.pkcs11.library</code>		Use this property to specify the path to the PKCS11client library.
<code>com.ibm.di.server.pkcs11.slot</code>		Specify the slot number of the device.
<code>{protect}-com.ibm.di.server.pkcs11.pass</code>		Use this password to access the pkcs11 compliant crypto device.
<code>com.ibm.di.server.pkcs11.accl</code>	<code>false</code>	Use <code>=true</code> to set hardware cryptographic devices for cryptographic operations.

Enabling or disabling padding

Padding means adding extra bits to a transmission so that the transmission is the exact, required, size. Some encryption and decryption algorithms require their input to be an exact multiple of the block size. If the plaintext to be encrypted is not an exact multiple, you must *pad* before encrypting by adding a padding string. When decrypting let the receiving party know how to remove the padding.

Note: All properties listed in the `global.properties` file can be set in the configuration file by the same name; it is recommended that you edit your `solution.properties` file instead if you have one. These properties can be protected by encryption using the `{protect}-` prefix (see section “Standard TDI encryption of `global.properties` or `solution.properties`” on page 96 for details).

When setting the property for padding, the default value is `DES/ECB/NoPadding`. The padding property defines an algorithm or cipher for password-based encryption and decryption of Tivoli Directory Integrator configurations. The property is: `com.ibm.di.securityTransformation`.

Maintaining encryption artifacts – keys, certificates, keystores, encrypted files

Changed encryption key

Any change of the key that the Server uses for encryption leads to a need for migration of existing encrypted files. To migrate an encrypted file, you should decrypt it with the old encryption key and encrypt it with the new one. Encryption and decryption can be done using the `cryptoutils` tool.

Files which are often encrypted or contain encrypted parts are: configurations, the User Registry and properties files (Tivoli Directory Integrator properties files can contain encrypted properties, although the files are usually not encrypted as a whole).

Note: By default all sensitive properties (such as passwords) inside `global.properties` or `solution.properties` are encrypted. As a rule of thumb you should always migrate `global.properties` and `solution.properties` files when you change the Server encryption key.

Changed password for encryption key or keystore

The Server reads the password for the keystore that holds the encryption key and the password for the encryption key itself from the Server stash file. Thus if any of those passwords is changed, the stash file must be updated. This can be done using the `createstash` tool.

Expired encryption certificate

If the Server uses public-key encryption, the certificate associated with the encryption key-pair can potentially expire at some point in time. If this happens, the certificate can be renewed using the procedure described in section “Extending the validity of a certificate using `keytool`” on page 69. That procedure preserves the underlying keys, so no migration of existing encrypted files is necessary.

Chapter 10. Configuring the TDI Server API

The IBM Tivoli Directory Integrator 7.0 Server API provides a set of programming calls that can be used to develop Tivoli Directory Integrator solutions and interact with the server locally and remotely. It also includes a management layer that exposes the Server API calls through the Java Management Extensions (JMX) interface. This section provides information about properties you can use to configure the Server API.

- For information on using the Server API, see "Appendix C. Server API" in the *IBM Tivoli Directory Integrator V7.0 Reference Guide*.
- For additional information on configuring the server API, see "Configuring the Server API" on page 73.
- For information on Tivoli Directory Integrator server security, see "Remote Server API" on page 72.

Server ID

Remote clients can identify the server they are talking to if the server can be distinguished using a unique ID. Tivoli Directory Integrator allows you to specify unique server IDs that allow remote clients, such as the Administration and Monitoring Console (AMC), to connect to different Tivoli Directory Integrator servers at different times using the same IP and port. In order to connect using the same ID and port but at different times, Tivoli Directory Integrator client applications must be able to register these client applications as different Tivoli Directory Integrator servers that you can associate with different data and databases.

Users are to assign the unique IDs manually, ensuring that any remote client (such as AMC) can connect to a Tivoli Directory Integrator server based on the IP address, the port, and the unique ID of the Tivoli Directory Integrator server. AMC registers every server with a unique ID, so that a Tivoli Directory Integrator server cannot be registered more than once by mistake or intentionally. When assigning IDs manually, users must ensure that different Tivoli Directory Integrator servers have distinguishable IDs.

You can configure the unique server ID property using `com.ibm.di.server.id`. To give a server a unique server id for a given server, provide a unique ID string for this property in the `global.properties` or `solution.properties` file on the server you are identifying. The default value for `com.ibm.di.server.id` is blank.

Exception for password protected Configs

The server API throws an exception if you use a password protected Config without using a password. The Tivoli Directory Integrator server API can detect and handle server problems when the server is faced with clients trying to access password protected configurations without supplying a password. A message displays, notifying the user about the problem. The error message is invoked when no password is supplied or the when the password entered is wrong. See "AMC and encrypted configs" on page 198.

Server RMI

Due to increasing needs for remote access to each Tivoli Directory Integrator server, Remote Method Invocation (RMI) is enabled by default. To ensure adequate security, default remote access requires Secure Socket Layer (SSL) for client authentication. The SSL access is facilitated with the sample keystore and truststore that are deployed with Tivoli Directory Integrator. See "SSL client authentication" on page 65 and "Summary of Server API Authentication options" on page 83.

Config load time-out interval

If a config instance does not completely load the configuration file when the server API makes a call, the server API returns a null object. You can add a time-out interval by adding the following property: `api.config.timeout` in the `global.properties` file. The interval for loading the configuration file is set to two minutes by default. If the config file does not load within the time interval, an exception is thrown.

Chapter 11. Properties

You can use Properties to configure TDI components and the TDI Server. Properties are simple keyword:value pairs of parameters kept outside your configuration files (configs), stored in External Properties files. This enables you to keep confidential information like passwords outside of your Config files. The `global.properties` file is the main configuration file for TDI. Properties are defined in the `global.properties` file or in the `solution.properties` file. The `solutions.property` file is a writable copy of the `global.properties` file and is used when the server is started from the solution directory. If a solution directory different from the installation directory is specified during installation, then a copy of the `global.properties` file named `solution.properties` is created in the Tivoli Directory Integrator solution directory. Both files are text files, and are written so that they can be understood by the operating system that is running on the platform.

Properties are single-valued data containers that hold parameter information, for example, *true* or *5000*. You can access properties from script using entry functions like `getProperty()` and `setProperty()`. Get and set methods work directly with property values. Entry objects can also contain properties. Like Attributes, properties are data containers. Attributes are used to store data content, but properties hold parametric information. Property values and Attributes can be any type of Java Object. Properties do not show up in:

- Attribute map selection.
- Work entry lists.

Working with properties

This section introduces the primary concepts you need when working with properties. Properties set in any properties files form a baseline for the entire IBM Tivoli Directory Integrator installation for all users on that computer. However, if your Solution Directory is different from the installation directory, you can have a set of text files in your Solutions Directory that mirror their counterparts in the installation directory. A property listed in any of those files override anything set in any of the global installation property files, including the `global.properties` and `solution.properties` files. Furthermore, a Java property set inside a Config file takes the highest precedence, and overrides anything in a global property file or the property files in the Solution Directory.

You can specify the Solution Directory in multiple ways:

- Set the environment variable `TDI_SOLDIR` before starting the Configuration Editor or the Server.
- Specify the `-s` parameter to the `ibmditk` script to start the Configuration Editor or the `ibmdisrv` script to start the TDI Server. This takes precedence over setting `TDI_SOLDIR`. If `TDI_SOLDIR` equals the installation directory, all property files are read from there, and the remarks about property files in the Solutions Directory do not apply.

In any other case, the first time you run the TDI Server, it makes a copy of all the property files into your Solutions Directory (it does not overwrite these files if they already exist). You can now tailor these files to your particular needs, without affecting the property files in the installation directory. The files remaining in the installation directory continue to form a baseline configuration for other instances of TDI.

Note: The file `global.properties` is copied to a file called `solutions.properties` in your Solutions Directory. Other files, like `Log4J.properties` and the files in the `amc` and `serverapi` folders are copied under their own names.

For documentation purposes, the original `global.properties` file from the installation directory is copied to the `<Solution directory>/etc` folder; this file is **not** used for any other purpose.

Migrating using properties and the tdimiggb1 tool

The tdimiggb1 tool helps you to migrate your global.properties file from one version of Tivoli Directory Integrator to a higher version. See the following chapter Chapter 5, “Migrating,” on page 41.

Global properties

Global properties are used to configure the Tivoli Directory Integrator Server settings that are kept in a file called global.properties in the etc folder of your installation directory. All properties included in the global.properties file are listed with their default values and explained in this chapter. A reference to more detailed documentation is provided, where possible, in the beginning of the groups of properties. The Configuration Editor (CE) (ibmditk) and the Tivoli Directory Integrator server (ibmdisrv) read the global.properties file on startup. This file is read and applied before a file called solution.properties from your Solution Directory is read.

Table 17. Some important Tivoli Directory Integrator global properties

Use of the property	Property	Default value	Description
Add your own .jar or .zip files	com.ibm.di.loader.userjars	c:\myjars	Specifies directories or jar files, separated by the Java Property "path.separator", which is ":" on Linux and ";" on Windows. Directories are searched recursively by the TDILoader for jar files containing classes and resources. Only files with a .zip or .jar extension are searched.
Define cipher	com.ibm.di.securityTransformation	DES/ECB/NoPadding	Defines a cipher for the password-based encryption or decryption of Tivoli Directory Integrator configurations. Changed in Tivoli Directory Integrator 7.0.
Enable config autoload	com.ibm.di.server.autoload	autoload.tdi	Looks for *.xml files in the directory specified by the "ibmdisrv -d" command. Executes each *.xml file found in the directory defined by -d.

Solution properties

Solution properties typically override global properties and are found in a file in your solution directory called solution.properties. The solution.properties file is by default a copy of the global.properties file, and you should edit the solution.properties file when configuring TDI, because it is read last out of all the properties files. If you want to, you can delete properties in your solution.properties file and add property configuration statements that you specifically want to override the global.properties defaults.

Java properties

Java properties are variables and settings of the Java Virtual Machine (JVM). Java log (Jlog) file properties are shown in “Useful JLOG parameters” on page 181.

Note: A Java property set inside a Config file takes the highest precedence, and overrides anything in a global property file or the property files in the Solution Directory.

Table 18. Java properties

Property	Default value	Description
<code>javax.net.debug</code>	none	Sets debug mode for the JSSE provider.
<code>com.ibm.di.javacmd</code>	none	Overrides the Java interpreter.
<code>com.ibm.di.installdir</code>	none	Uses this path to the Java executable file when running AssemblyLines from the Configuration Editor.
<code>com.ibm.di.jvmdir</code>	<code>TDI_root/jvm</code>	Defines the directory path where the JRE that Tivoli Directory Integrator uses is installed.
<code>com.ibm.di.server.maxThreadsRunning</code>	500	Sets this number of threads Tivoli Directory Integrator. Must be set higher than 3 to have any effect.
<code>com.ibm.di.server.securemode</code>	false	Sets the mode in which Tivoli Directory Integrator is running. (standard or secure)
<code>com.ibm.di.server.keystore</code>	<code>testserver.jks</code>	Names the keystore of the Server's SSL certificate. Renamed in Tivoli Directory Integrator 7.0.
<code>com.ibm.di.server.key.alias</code>	server	Names the key alias of the Server's SSL certificate. Renamed in Tivoli Directory Integrator 7.0.
<code>{protect}-api.keystore.password</code>	server (encrypted by default)	Provides the password for the server API keystore. Added in TDI 7.0.
<code>{protect}-api.key.password</code>		Provides the key password. If not specified, uses server keystore password. Added in Tivoli Directory Integrator 7.0.
<code>com.ibm.di.server.encryption.keystore</code>	<code>testserver.jks</code>	Names the keystore of the server encryption key. Added in Tivoli Directory Integrator 7.0.
<code>com.ibm.di.server.encryption.key.alias</code>	server	Provides the key alias of the server encryption key. Added in Tivoli Directory Integrator 7.0.
<code>com.ibm.di.server.encryption.keystoretype</code>	jks	Provides the type of the keystore that hosts the key used by the server for encryption. Added in Tivoli Directory Integrator 7.0.
<code>com.ibm.di.server.encryption.transformation</code>	RSA	Names the cryptographic transformation used by the server for encryption. Can be set to either "RSA" (public key encryption) or to some secret key transformation. Added in Tivoli Directory Integrator 7.0.

Table 18. Java properties (continued)

Property	Default value	Description
<code>com.ibm.di.server.fipsmode.on</code>	false	Enables or disables FIPS standards in TDI. If this property is set to true, TDI runs in FIPS-compliant mode. For more information on FIPS mode, see Added in Tivoli Directory Integrator 7.0.

System properties

System properties are stored in the System Store instead of being stored in an external properties file such as `solution.properties`. Certain system properties and Java properties are read-only. These system properties are shown in the respective Property Stores (for example, System Store). Attempting to modify these read-only properties has no effect. See also Chapter 12, “System Store,” on page 145.

Chapter 12. System Store

IBM Tivoli Directory Integrator supports persistent storage (that is, storage of objects that survive across JVM restarts), by means of a tightly-coupled relational database, the System Store.

The product deployed by default to implement the system store is a relational database implemented fully in Java, known as Apache Derby, and previously known as Cloudscape.

The System Store supports the following objects:

- Delta Tables
- User Property Store
- Password Store

Property stores

Password store and User property stores are types of system stores.

Password Store

The *Password Store* is an external repository that stores a value which results from changing the value for a password syntax component. The password protection mechanism is directly related to the configuration windows offered to the user. The configuration windows, or forms, contain descriptions of each parameter and its syntax. One type of syntax is *password* which causes the Configuration Editor to use a password text field for editing. This external repository for passwords is configured in the *Properties* page in the configuration editor (*Password-Store*) and is specified in the configuration file for the current TDI solution. If no such property store is configured the password is saved in clear text in the configuration file.

If a default password store is configured, a unique property name is generated the first time a protected/password parameter is saved. This key is used as the key in the password store. The same property name is written to the configuration file as a standard property reference. When the value is later retrieved, standard property resolution takes place to retrieve the actual value from the password store.

If a Password Store is specified, a unique key is generated for the password and the password is saved encrypted in the Password Store under that key. In the configuration file, the password is referenced only by that key.

User property stores

The User Property Store is a System Store table used for maintaining serialized Java objects associated with a key value. This is where persistent component parameters and properties (such as the **Iterator State Store**) are maintained, as well as any data you store. The System Store implements User property stores as one of its three types of persistent stores for IBM Tivoli Directory Integrator components. For information on TDI user interfaces that allow you to select properties from a property store, see “Add a Solution View” on page 206.

Third-party RDBMS as System Store

The System Store can also be configured to use other multi-user RDBMS systems, as opposed to using the bundled database, Apache Derby. This is done by specifying appropriate SQL Data Definition Language (DDL) statements and driver parameters as system properties in `global.properties` or `solution.properties`. Example statements, commented out, are present in the distribution version of `global.properties` in the `TDI_install_dir/etc` directory, for the supported configurations of IBM DB2, Oracle and MS SQL*Server.

It is also possible to take advantage of suitable templates built in to the Configuration Editor, by going to the appropriate Tivoli Directory Integrator Server document. Right-click on the Server in the Servers pane, and select **Edit system store settings. The Server System Store** header in the window is a context-sensitive menu; it has selections for Derby Embedded, Derby Networked, Oracle and DB2.

JDBC Driver parameters provide a path to the database; additional properties are used to specify tailored SQL for certain operations Tivoli Directory Integrator must be able to perform in the System Store. Multiple SQL statements can be specified per property. Each separate statement should be terminated with a semicolon. An example property could be (note that for display purposes, the statements in this document are broken up in multiple lines; however, in your property file all statements for a given property should be on one line):

```
com.ibm.di.store.create.delta.systable=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
SEQUENCEID int, VERSION int);
ALTER TABLE {0} ADD CONSTRAINT IDI_DS_{UNIQUE} PRIMARY KEY (ID)
```

where {0} => is replaced by the Table name; and

{UNIQUE} => is a special variable which can be used to generate a unique name based on the current system time.

The following section lists example connection parameters and statements for each of the supported RDBMS systems.

Oracle

Usage of Oracle requires that you drop the JDBC driver client library, `ojdbc14.jar`, in the `TDI_install_dir/jars` directory.

JDBC connection parameters

```
com.ibm.di.store.database=jdbc:oracle:thin:@itdidev.in.ibm.com:1521:itimdb
com.ibm.di.store.jdbc.driver=oracle.jdbc.OracleDriver
com.ibm.di.store.jdbc.urlprefix=jdbc:oracle:thin:
com.ibm.di.store.jdbc.user=SYSTEM
{protect}-com.ibm.di.store.jdbc.password=password
```

Where *itimdb* is the SID of the database to be used as System Store.

Create table statements

```
com.ibm.di.store.create.delta.systable=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
SEQUENCEID int, VERSION int);ALTER TABLE {0} ADD CONSTRAINT IDI_CS_{UNIQUE} PRIMARY KEY (ID)
com.ibm.di.store.create.delta.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
SEQUENCEID int, ENTRY BLOB );ALTER TABLE {0} ADD CONSTRAINT IDI_DS_{UNIQUE} Primary Key (ID)
com.ibm.di.store.create.property.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
ENTRY BLOB );ALTER TABLE {0} ADD CONSTRAINT IDI_PS_{UNIQUE} Primary Key (ID)
com.ibm.di.store.create.sandbox.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
ENTRY BLOB )
com.ibm.di.store.create.recal.conops=CREATE TABLE {0} (METHOD varchar(VARCHAR_LENGTH), RESULT BLOB,
ERROR BLOB)
```


MS SQL Server

Use of MS SQL Server requires that you install a number of Microsoft client libraries in the *TDI_install_dir/jars* directory.

JDBC connection parameters

```
com.ibm.di.store.database=jdbc:Microsoft:sqlserver://localhost:1433;DatabaseName=master;selectMethod=cursor;
com.ibm.di.store.jdbc.driver=com.microsoft.jdbc.sqlserver.SQLServerDriver
com.ibm.di.store.jdbc.user=sa
com.ibm.di.store.jdbc.password=passw0rd
```

The above connection parameters are used with these Microsoft JDBC jars:

1. Msutil.jar
2. MsBase.jar
3. MSsqlserver.jar

Note: For Microsoft SQL Server 2005, the driver jar file to be placed in the *TDI_install_dir/jars* directory is *sqljdbc.jar* (only one file is required) and it can be obtained from your SQL Server 2005 installation at <Microsoft SQL Server 2005-Install-Dir>/sqljdbc_<version>/<language>/sqljdbc.jar; the JDBC connection parameters need to be specified as follows:

```
com.ibm.di.store.database=jdbc:sqlserver://localhost:1433;DatabaseName=name;selectMethod=cursor;
com.ibm.di.store.jdbc.driver=com.microsoft.sqlserver.jdbc.SQLServerDriver
com.ibm.di.store.jdbc.user=sa
com.ibm.di.store.jdbc.password=passw0rd
```

The *selectMethod* property is optional to the *jdbc* URL. When this property is set to "cursor", a database cursor is created. This is useful when reading very large result sets that cannot be contained in the clients memory.

The default behavior of *selectMethod* is not "cursor", but "direct", which keeps result sets in clients memory, thus providing much faster performance. So unless memory is a problem, it is better to use the default "direct" behavior. For more information: [http://msdn.microsoft.com/en-us/library/ms378988\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms378988(SQL.90).aspx).

JDBC connection parameters (for JSQLConnect driver)

```
com.ibm.di.store.database= jdbc:JSQLConnect://itdidserver/database=reqpro
com.ibm.di.store.jdbc.driver= com.jnetdirect.jsql.JSQLDriver
com.ibm.di.store.jdbc.urlprefix= jdbc:JSQLConnect:
com.ibm.di.store.jdbc.user=administrator
{protect}-com.ibm.di.store.jdbc.password=password
```

These connection parameters are used with JSQLConnect drivers. You must download the JSQLConnect.jar file and copy it into the *TDI_install_dir/jars* directory.

Create table statements

The DATA TYPE for MS SQL is IMAGE.

```
com.ibm.di.store.create.delta.systable=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
SEQUENCEID int, VERSION int);
ALTER TABLE {0} ADD CONSTRAINT IDI_MYCONSTRAINT_{UNIQUE} PRIMARY KEY (ID)
com.ibm.di.store.create.delta.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
SEQUENCEID int, ENTRY IMAGE );
ALTER TABLE {0} ADD CONSTRAINT IDI_DS_{UNIQUE} Primary Key (ID)
com.ibm.di.store.create.property.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
ENTRY IMAGE );
ALTER TABLE {0} ADD CONSTRAINT IDI_PS_{UNIQUE} Primary Key (ID)
com.ibm.di.store.create.sandbox.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
ENTRY IMAGE)
com.ibm.di.store.create.recal.conops=CREATE TABLE {0} (METHOD varchar(VARCHAR_LENGTH),
RESULT IMAGE, ERROR IMAGE)
```

IBM DB2 for z/OS

JDBC connection parameters

```
com.ibm.di.store.database=jdbc:db2:net://localhost:50000/idi db
com.ibm.di.store.jdbc.driver=com.ibm.db2.jcc.DB2Driver
com.ibm.di.store.jdbc.urlprefix= jdbc:db2:net:
com.ibm.di.store.jdbc.user=db2admin
{protect}-com.ibm.di.store.jdbc.password=db2admin
```

Where *idi db* in the database URL is the DSN for a DB2 instance.

The above connection parameters are used with the db2jcc_license_cisuz.jar license jar file.

Create table statements

Tablespace and Indexes must be unique.

```
com.ibm.di.store.create.delta.systable=CREATE TABLESPACE TS1DSYS LOCKSIZE ROW BUFFERPOOL BP32K;
CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SEQUENCEID int, VERSION int) IN TS1DSYS;
CREATE UNIQUE INDEX DSTIX1 ON {0} (ID ASC);
ALTER TABLE {0} ADD CONSTRAINT IDI_DT_{UNIQUE} PRIMARY KEY (ID);
com.ibm.di.store.create.delta.store=CREATE TABLESPACE TS1DST LOCKSIZE ROW BUFFERPOOL BP32K;
CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SEQUENCEID int, ENTRY BLOB) IN TS1DST;
CREATE UNIQUE INDEX DSIX1 ON {0} (ID ASC);
ALTER TABLE {0} ADD CONSTRAINT IDI_DS_{UNIQUE} Primary Key (ID);
CREATE LOB TABLESPACE DSENT11 BUFFERPOOL BP32K LOCKSIZE LOB;
CREATE AUX TABLE TBDSEN1 IN DSENT11 STORES {0} COLUMN ENTRY;
CREATE INDEX IXEN1 ON TBDSEN1
com.ibm.di.store.create.property.store=CREATE TABLESPACE PS3DST LOCKSIZE ROW BUFFERPOOL BP32K;
CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, ENTRY BLOB) IN PS3DST;
CREATE UNIQUE INDEX PSIX3 ON {0} (ID ASC);
ALTER TABLE {0} ADD CONSTRAINT IDI_PS_{UNIQUE} Primary Key (ID);
CREATE LOB TABLESPACE PSENT31 BUFFERPOOL BP32K LOCKSIZE LOB;
CREATE AUX TABLE TBPSEN3 IN PSENT31 STORES {0} COLUMN ENTRY;
CREATE INDEX PSIXEN3 ON TBPSEN3
com.ibm.di.store.create.sandbox.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, ENTRY BLOB)
com.ibm.di.store.create.recal.conops=CREATE TABLESPACE IM{UNIQUE} LOCKSIZE ROW BUFFERPOOL BP32K;
CREATE TABLE {0} (METHOD VARCHAR(VARCHAR_LENGTH), RESULT BLOB, ERROR BLOB) IN IM{UNIQUE};
CREATE LOB TABLESPACE LB{UNIQUE} BUFFERPOOL BP32K LOCKSIZE LOB;
CREATE AUX TABLE AT{UNIQUE} IN LB{UNIQUE} STORES {0} COLUMN RESULT;
CREATE INDEX IX{UNIQUE} ON AT{UNIQUE};
CREATE LOB TABLESPACE LS{UNIQUE} BUFFERPOOL BP32K LOCKSIZE LOB;
CREATE AUX TABLE AE{UNIQUE} IN LS{UNIQUE} STORES {0} COLUMN ERROR;
CREATE INDEX IN{UNIQUE} ON AE{UNIQUE}
```

DB2 for other OS

JDBC connection parameters

```
com.ibm.di.store.database=jdbc:db2:net://localhost:50000/idi db
com.ibm.di.store.jdbc.driver=com.ibm.db2.jcc.DB2Driver
com.ibm.di.store.jdbc.urlprefix= jdbc:db2:net:
com.ibm.di.store.jdbc.user=db2admin
{protect}-com.ibm.di.store.jdbc.password=db2admin
```

Where *idi db* in the database URL is the DSN for a DB2 instance.

Create table statements

```
com.ibm.di.store.create.delta.systable=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
SEQUENCEID int, VERSION int);
ALTER TABLE {0} ADD CONSTRAINT IDI_MYCONSTRAINT_{UNIQUE} PRIMARY KEY (ID)
com.ibm.di.store.create.delta.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
SEQUENCEID int, ENTRY BLOB );
ALTER TABLE {0} ADD CONSTRAINT IDI_DS_{UNIQUE} Primary Key (ID)
com.ibm.di.store.create.property.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
ENTRY BLOB );ALTER TABLE {0} ADD CONSTRAINT IDI_PS_{UNIQUE} Primary Key (ID)
com.ibm.di.store.create.sandbox.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL,
ENTRY BLOB )
```

Using Derby to hold your System Store

The remainder of this chapter discusses the operational aspects of using Derby, in particular in conjunction with using Derby to hold your System Store.

Note: With regards to third party RDBMSs, in order to hold encrypted password values you may have to size the fields that hold them quite large. A typical small password might use as much as 178 characters. It depends on both your server's key, and the length of the unencrypted data you try to store (in bytes). Since this is a blocked encoding a larger password might use the same space, or double or triple that amount. Also, the size of the block depends on the server's key. One way to find the size you need, is to store the password (protected) to a file first, and then look at that file to see how many characters were used.

Derby can run in either of two modes: *embedded* and *networked*. By default, as specified in the `global.properties` file, Derby runs in *networked* mode.

The System Store used by Tivoli Directory Integrator releases before V7.0 was Derby (then called Cloudscape) in *embedded* mode. There are drawbacks to the way Derby runs in embedded mode. In embedded mode, Derby runs as a separate thread within the JVM when required. Startup and shutdown of Derby is automatic in embedded mode. However, when run this way, this Derby thread claims exclusive access to the database files. This can become problematic when different JVMs, each with its own Derby thread, try to access the same System Store.

In embedded mode, these actions cause a new, independent JVM to be started, triggering an access conflict when more than one JVM is active at any given time:

- A command line invocation of the IBM Tivoli Directory Integrator Server with a config file, causing one or more AssemblyLines to run.
- Startup of the Configuration Editor (GUI)
- Startup of an AssemblyLine from within the Configuration Editor

None of these actions by themselves causes the Derby thread to start. However, the Derby thread does start if access to any of the objects in the System Store is required (for example, Objects supported by the System Store such as Delta Tables and the User Property Store).

The solution to the access conflicts as outlined previously is to run Derby in *networked* mode, which enables concurrent access to the System Store. Also enable user authentication in derby to avoid security concerns in networked mode. To provide security at the database level, TDI uses the BUILTIN security provider for Derby. BUILTIN ensures that only valid users are able to access the Derby database. When you have Derby configured in networked mode, you can work with multiple instances of Derby databases booted as System Stores. You can also configure a Derby instance to work with a specific Configuration file instance.

Note: Depending on how Derby was started, instances of Derby may be left running in networked mode, even after all other Tivoli Directory Integrator processes have terminated.

When you set the property `derby.drda.startNetworkServer` to true (by default, this is the case, in `global.properties`), the Network Server automatically starts when you start Derby (in this context, Derby starts when the embedded driver is loaded). You may have to terminate Derby manually, if desired.

Configuring Derby Instances

To configure and manage multiple Derby instances and to provide facilities to start, stop and restart Derby servers in networked mode a menu option called **System Store** is provided in the Tivoli Directory Integrator Configuration Editor, as part of **Log & Settings** configuration of a project. Many of the configuration options listed take default values from the `global.properties` file, which was the configuration base for previous versions of Tivoli Directory Integrator; now.

The **System Store** menu option also provides ways to configure the System Store to use other databases like IBM DB2 as the backend RDBMS. For more information, refer to "System Store settings" under **The Configuration Editor -> Log & Settings** in *IBM Tivoli Directory Integrator V7.0 Users Guide*.

Starting Derby in networked mode

If the `com.ibm.di.store.hostname` property is set to `localhost` then remote connections are not allowed. If the `com.ibm.di.store.hostname` property is set to the IP address of the local computer running Tivoli Directory Integrator, then remote clients can access this Derby instance by using the IP address. You can only start the network server for the local computer.

Table 19. Starting Derby in networked mode

Property	Default value	Description
<code>com.ibm.di.store.start.mode</code>	automatic	The mode for starting up the Derby server process when required – set to automatic or manual.
<code>Com.ibm.di.store.hostname</code>	localhost	The URL of the Derby server.
<code>Com.ibm.di.store.port</code>	1527	The port for connecting to the Derby server.
<code>Com.ibm.di.store.sysibm</code>	true	The state for using the SYSIBM schema or not; values true or false.
<code>com.ibm.di.store.varchar.length</code>	512	The <code>varchar(length)</code> for the ID columns used in system store and System Store (PES) connector tables.

Enabling user authentication in System Store

Add these properties to the `global.properties` file after the System Store network mode properties.

Table 20. Enable user authentication in System Store

Property	Default value	Description
<code>derby.connection.requireAuthentication</code>	true	Enables user authentication for the System Store.
<code>Derby.authentication.provider</code>	BUILTIN	Sets the user authentication provider to BUILTIN. This is the most basic and simple authentication provider that Derby has.
<code>Derby.database.defaultAccessMode</code>	fullAccess	Defines the access level to the System Store user. The different access levels supported by Derby are "fullAccess", "readOnly" and "noAccess".

Create statements for System Store tables

You can configure create table SQL statements for

- Delta systable
- Delta table
- Property table
- Sandbox tables

- Record AssemblyLine table
- Tombstone manager table
- ibmsnap_commitseq column name

Table 21. Create statements for System Store

Property	Default value	Description
com.ibm.di.store.create.delta.systable	CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SEQUENCEID int, VERSION int);ALTER TABLE {0} ADD CONSTRAINT IDI_CS_{UNIQUE} PRIMARY KEY (ID)	Create table SQL statements for the delta systable.
com.ibm.di.store.create.delta.store	CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SEQUENCEID int, ENTRY BLOB);ALTER TABLE {0} ADD CONSTRAINT IDI_DS_{UNIQUE} Primary Key (ID)	Create table SQL statements for the delta table.
com.ibm.di.store.create.property.store	CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, ENTRY BLOB);ALTER TABLE {0} ADD CONSTRAINT IDI_PS_{UNIQUE} Primary Key (ID)	Create table SQL statements for the property table.
com.ibm.di.store.create.sandbox.store	CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, ENTRY BLOB)	Create table SQL statements for the Sandbox tables.
com.ibm.di.store.create.recal.conops	CREATE TABLE {0} (METHOD varchar(VARCHAR_LENGTH), RESULT BLOB, ERROR BLOB)	Create table SQL statements for Record AL.
com.ibm.di.store.create.tombstones	CREATE TABLE IDI_TOMBSTONE (ID INT GENERATED ALWAYS AS IDENTITY, COMPONENT_TYPE_ID INT, EVENT_TYPE_ID INT, START_TIME TIMESTAMP, CREATED_ON TIMESTAMP, COMPONENT_NAME VARCHAR(1024), CONFIGURATION VARCHAR(1024), EXIT_CODE INT, ERROR_DESCR VARCHAR(1024), STATS LONG VARCHAR FOR BIT DATA, GUID VARCHAR(1024) NOT NULL, USER_MESSAGE VARCHAR(1024), UNIQUE (ID, GUID))	Specify the SQL statement for creating the Tombstone Manager table. Keep the same table names and field names.
com.ibm.di.conn.rdbmschlog.cdcolname	ibmsnap_commitseq	Provide the ibmsnap_commitseq column name to be used by the RDBMS changelog connector.

Backing up Derby databases

Another matter that needs to be given some thought is **backup** of the data contained in a Derby database. The recommended (and simplest) way of doing this is to

- Shutdown the Derby database (if running in embedded mode, shut down all Tivoli Directory Integrator instances and Configuration Editor instances)
- Copy the entire Derby directory in your Tivoli Directory Integrator home directory (or whatever Derby directory your `global.properties` file is pointing to) to a different location, and ensure that this data is safe
- Restart the Derby database (if running in networked mode).

To restore a database, reverse source and destination of the copy operation in the above list of steps.

Troubleshooting Derby issues

This section does not attempt to be a comprehensive Troubleshooting Guide for Derby, but there are a number of symptoms that are observed sometimes in the context of usage of Derby as the underlying database in Tivoli Directory Integrator. These are:

Schema 'SYSIBM' does not exist error

Question:

I'm trying to use Derby in networked mode and having issues. I've figured out how to start it up and I'm able to query it with `sysinfo` and `testconnection`, but when I run TDI and try to open the system store I get an error stating:

```
[com.ibm.db2.jcc.a.SQLException: Schema 'SYSIBM' does not exist]
```

How do I fix this?

Explanation:

The reason you get this error is because you are trying to boot a database that was created in embedded mode into a networked mode server without starting the server using the `-ld` flag. Note that for a networked mode Derby server to open an embedded mode database, the `SYSIBM` schema **MUST** be loaded. The `SYSIBM` schema is a special schema loaded by the Derby server. The `SYSIBM` contains stored prepared statements that return result sets to determine metadata information.

Corrective action:

To solve this problem start the Derby networked server with the `"-ld"` flag, like:

```
./dbserver start -p 1527 -ld
```

Another Instance of Derby may already be booted

You may get the following error sometimes, especially when using Derby in embedded mode:

```
[ERROR XSDB6: Another instance of Derby may have already booted the database D:\tdi60\Derby.]
```

Explanation:

Derby try to prevent two instances of Derby from booting the same database (in this case `D:\tdi60\Derby`). This can happen if you are running two `AssemblyLines` which are trying to update the same Derby database running in embedded mode. This error might also crop up if you have an unclosed connection to the database.

Corrective Action:

1. If you want two AssemblyLines to update the same Derby database, then the correct mode of Derby should be networked mode; this mode of operation does not have that limitation.
2. You can work around this by closing the database using the **Browse Server Stores** option and then clicking on the **Close** button. Even if the database is not open, just opening and closing again through the **Browse Server Stores** option help solve this problem.

Future versions of Tivoli Directory Integrator attempt to handle this situation automatically, and stop and start Derby as required.

Can I use DB2 as a system store?

In Tivoli Directory Integrator it is possible to use DB2 as a system store, instead of the bundled Derby database system. However, some modification of system properties files is required for this to function correctly. You must replace the section on Derby networked mode with a section similar to the following (insert the correct parameters for your installation).

If you look at the default `global.properties` file, there are some `CREATE_TABLE` statements for using and setting up the system store. If you use the right syntax, you can use non-Derby databases as system store. Here is the DB2 syntax:

```
## Location of the DB2 database (networked mode)
com.ibm.di.store.database=jdbc:db2://168.199.48.4:3700/tdidb
com.ibm.di.store.jdbc.driver=com.ibm.db2.jcc.DB2Driver
com.ibm.di.store.jdbc.urlprefix=jdbc:db2:
com.ibm.di.store.jdbc.user=db2inst1
com.ibm.di.store.jdbc.password=*****
com.ibm.di.store.start.mode=automatic
com.ibm.di.store.port=3700
com.ibm.di.store.sysibm=true

# the varchar(length) for the ID columns used in system store and PES Connector tables
com.ibm.di.store.varchar.length=512

# create statements for DB2 system store tables
com.ibm.di.store.create.delta.systable=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH)
NOT NULL, SEQUENCEID int, VERSION int)
com.ibm.di.store.create.delta.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH)
NOT NULL, SEQUENCEID int, ENTRY BLOB )
com.ibm.di.store.create.property.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH)
NOT NULL, ENTRY BLOB )
com.ibm.di.store.create.sandbox.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH)
NOT NULL, ENTRY BLOB )
```

Note: Each `com.ibm.di.store.create.xxx` statement must be specified on one line, even though they are broken up in this example for illustration purposes.

Why can't remote connections be made to my Derby network server?

This may be because the Derby Server has been started by passing "localhost" as the hostname. This disallows any remote connections to be made to Derby. Stop the Derby server and start it with hostname parameter specified as the computer's IP address. This can be done by going to the Configuration Editor's **Server System Store** Server Settings window (available from the context menu on a server in the Servers view).

For more details, check <http://publib.boulder.ibm.com/infocenter/cldscp10/topic/com.ibm.cloudscape.doc/hubprnt22.htm>

Pre-6.0 (properties file) configuration of Cloudscape

Previous versions of Tivoli Directory Integrator configured the System Store using Cloudscape by means of a set of properties in the `global.properties` file, and version 7.0, now using Derby, still derives its base configuration from there. You should migrate any non-standard installations of Cloudscape configuration to the methods described in the previous chapter, "Configuring Derby Instances" on page 150.

In the `global.properties` file in the installation directory, there are two sections that are concerned with the configuration of the System Store:

- The first section (commented out by default) deals with running Derby in embedded (dedicated), non-shared mode.
- The second section (enabled by default) deals with Derby in networked or shared mode. If you determine that you must use dedicated mode, uncomment the first section and comment the second.

Note that the section with configuration parameters for embedded mode specifies a relative path for the Derby database; this typically means the database is created in the Solutions directory. The set of parameters for Networked mode uses an absolute path: it points to the installation directory. If you switch between the different modes you must be aware of this, and possibly change the path to an appropriate value, for the mode you intend to use.

Best practice is to keep databases, keystores and all other user data out of the program installation directory.

When deploying networked or shared mode, startup of the Derby database thread is automatic; however, shutdown is not. You should shut down the instance when you are finished with your last AssemblyLine.

Cloudscape command-line utility:

To make working with the Derby database more convenient, consider creating a script ("dbserver") with the following line (this example is for Unix/Linux):

```
export DB_JAR_DIR=jars/3rdparty/IBM
export DB_CLASSPATH=$DB_JAR_DIR/derby.jar:$DB_JAR_DIR/derbyclient.jar:\
$DB_JAR_DIR/derbynet.jar:$DB_JAR_DIR/derbytools.jar
java -classpath $DB_CLASSPATH org.apache.derby.drda.NetworkServerControl "$@"
```

You may have to join the middle two lines together at the "\" point.

The equivalent dbserver.bat file for Windows would be:

```
set DB_JAR_DIR=jars/3rdparty/IBM
set DB_CLASSPATH=%DB_JAR_DIR%\derby.jar;%DB_JAR_DIR%\derbyclient.jar;\
%DB_JAR_DIR%\derbynet.jar;%DB_JAR_DIR%\derbytools.jar;
java -classpath %DB_CLASSPATH% org.apache.derby.drda.NetworkServerControl %*
```

Note: The script must be started from within the IBM Tivoli Directory Integrator installation path as the working directory, as the following classpath is relative to this directory.

The following is an example of usage of this utility script:

Show all available commands: `./dbserver`

Start DBServer `./dbserver start -p 1527`

Stop DBServer `./dbserver shutdown`

The full list of sub-commands that you can specify to the dbserver script, and which are sent to Derby is:

- `start [-h <host>] [-p <portnumber>]`: This starts the network server on the port/host specified or on localhost, port 1527 if no host/port is specified and no properties are set to override the defaults. By default Network Server will only listen for connections from the machine on which it is running. Use `-h 0.0.0.0` to listen on all interfaces or `-h <hostname>` to listen on a specific interface on a multiple IP machine.
- `shutdown [-h <host>] [-p <portnumber>]`: This shutdowns the network server on the host and port specified or on the local host and port 1527 (default) if no host or port is specified.
- `ping [-h <host>] [-p <portnumber>]`: This will test whether the Network Server is up.

- `sysinfo [-h <host>] [-p <portnumber>]`: This prints classpath and version information about the Network Server, the JVM and the Cloudscape server.
- `runtimeinfo [-h <host>] [-p <portnumber>]`: This prints extensive debugging information about sessions, threads, prepared statements, and memory usage for the running Network Server.
- `logconnections {on | off} [-h <host>] [-p <portnumber>]`: This turns logging of connections and disconnections on and off. Connections and disconnections are logged to `derby.log`. Default is off.
- `maxthreads <max> [-h <host>] [-p <portnumber>]`: This sets the maximum number of threads that can be used for connections. Default 0 (unlimited).
- `timeslice <milliseconds> [-h <host>] [-p <portnumber>]`: This sets the time each session can have using a connection thread before yielding to a waiting session. Default is 0 (no yield).
- `trace {on | off} [-s <session id>] [-h <host>] [-p <portnumber>]`: This turns drda tracing on or off for the specified session or if no session is specified for all sessions. Default is off .
- `tracedirectory <tracedirectory> [-h <host>] [-p <portnumber>]`: This changes where new trace files will be placed. For sessions with tracing already turned on, trace files remain in the previous location. Default is `cloudscape.system.home` .

When running in networked mode, the Derby database is of course reachable over the network, not only by IBM Tivoli Directory Integrator instances but also by other applications using the appropriate drivers. The credentials required for such access are defined in the `global.properties` file, and might have to be tailored for your particular site needs. Pay particular attention to the username and password parameters as these govern integrity and security of the data.

If you often alternate between running Derby in dedicated mode and in networked mode, consider having two different "prototype" `global.properties` files on your file system, one each with the correct set of parameters for each of the two modes. Just before starting a server instance, copy in place the appropriate `global.properties` file, according to your needs. Alternatively, use separate Solution Directories; in a Solution Directory you can have a file called `solution.properties`, which property values defined in there override the ones defined system-wide in `global.properties`.

See also

The official Derby home at <http://db.apache.org/derby>, documentation at <http://db.apache.org/derby/manuals/index.html>.

For further information on the command options and the specification on each command, see <http://db.apache.org/derby/docs/10.0/publishedapi/org/apache/derby/drda/NetworkServerControl.html>.

Chapter 13. Command-line options

Command-line options must have their value followed immediately after the option. Do not use a space between the option and the value. There are options for:

- "Configuration Editor"
- "Server"
- "Command Line Interface – tdisrvctl utility" on page 160

Configuration Editor

The CE is launched using the `ibmditk` wrapper script. This script invokes the Eclipse launcher for Tivoli Directory Integrator (`ce/eclipsece/miadmin`) with the proper settings for the Java VM and the Tivoli Directory Integrator install location property, both of which are required to run the current CE.

The Eclipse launcher (`ce/eclipsece/miadmin`) is a standard Eclipse launcher that takes command line parameters of its own. See Eclipse Command Line Options for a complete description of the Eclipse command line options.

```
"%TDI_HOME_DIR%\ce\eclipsece\miadmin" -vm "%TDI_JAVA_BIN_DIR%\javaw" -vmargs -Dcom.ibm.di.loader.IDILoader.path="%TDI_HOME_DIR%" %*
```

The above is a fragment of the `ibmditk` script showing the two required parameters (eclipse command line parameters) that the CE needs.

Of notable interest is the `-data` command line option that specifies the location of the workspace to use. If you are going to run multiple instances of the CE, you have to specify a different workspace for each instance of the CE since the workspace is locked by each instance. For example:

```
ibmditk -data c:/instance1_workspace
```

The above command launches the CE using `c:/instance1_workspace` as its workspace location.

Shutdown Servers option:

This is a command line option that attempts to stop all running servers that uses the same installation directory as the CE. When this option is given on the command line like this:

```
ibmditk -tdishutdown
```

then the CE will start and look at every defined server in the Tivoli Directory Integrator servers project, filtering out those that do not use the same installation directory as the CE and attempt to stop it. When this is done the CE will exit the Java VM with an exit code of zero. There is no guarantee that the servers the CE tried to stop actually did stop. Some servers may linger beyond the time it takes the CE to complete this command and some servers may simply refuse to stop for various reasons.

Server

The following command-line options are for the IBM Tivoli Directory Integrator server (`ibmdisrv [options]`):

Example:

```
ibmdisrv -c"C:\demos\rs.xml" -r"Access2LDAP" -l"c:\metamerge\mydemo.log"
```

Notes:

1. There is no space between the option letter and the value. Use quotes to save against possible spaces or commas in the values.

2. The Windows Shell executive allows a maximum of nine (9) arguments, from the list below. There aren't any limitations on other platforms.

-s <dir>

Specifies the working directory where the solution is located. All relative file references in Tivoli Directory Integrator and in your Configs and so forth will be relative to this location. Must be the first parameter specified.

-c <file...>

Configuration file(s). If you don't specify this option, the items in the Autostart folder will be loaded and started (unless suppressed by specifying **-D**). Wildcards, as in *.xml, are allowed too.

Note: Submitting multiple configuration files is only allowed if the **-d** option is also specified.

-n <encoding>

Encoding to be used to write Config files. This must be a valid character set identifier valid in Java2; refer to the IANA Charset Registry (<http://www.iana.org/assignments/character-sets>) for the full list of values. Note that Java2 only supports a subset of those.

-r <al...>

List of AssemblyLine names to start. To start AssemblyLine **a** and **b**, use the command **-r a b**. Other syntaxes are supported as well: **-ra,b**; **-ra -rb**.

Note: If you use includes and namespaces, the AssemblyLine can be myNamespace:/AssemblyLines/alName (assuming namespace **myNamespace** and AssemblyLine name **alName**).

-T<name>

Enable JLOG-style tracing to file trace<name>.log, in directory <Tivoli_Common_Dir>/TDI/logs/. Default is trace to memory (from which it can be retrieved by the traceback routines of JFFDC in case of an unhandled exception.)

-D Flag to disable startup of items in the Autostart folder.

-w If **-r** (or **-t**) is specified then this flag causes IBM Tivoli Directory Integrator to wait for each AssemblyLine to complete before starting the next. If this flag is not specified then IBM Tivoli Directory Integrator starts all AssemblyLines specified by the **-r** parameter in parallel. When the last AssemblyLine has finished, the server stops.

-e Specifying this option causes the server to run in Secure mode. Using the master password specific to this server, it will decrypt and encrypt all Config files as well as the server API Registry.

-v Show version information and exit. This is logged in the log file only.

-P <password>

Password if configuration file(s) is/are encrypted.

-p Dump Java properties on startup. Note that you still must provide a configuration file, which is read before Java properties are dumped.

-d Start a "daemon", or *Config Instance* on this system.

If you start with **-d**, you will start one anonymous instance (the daemon), which will start one config instance for each config file specified on the command line; this allows you to start multiple config instances on the fly. You may specify 0 or more config files on the command line. It does not make sense to specify any AssemblyLines to run in this mode, since it is impossible to state which config file the AL will be in. You can autostart AssemblyLines, though, since those belong to the config instance that specifies the autostart.

If you start without **-d**, you will get one config instance that loads the config file specified on the command line. You must specify exactly one config file on the command line. (If you must use

multiple config files, they may be piped in on standard input.) In this mode, you can specify any number of AssemblyLines to run. This is the traditional way of running the server.

-q Takes 1 argument, mode. Mode=1 means run in record mode, mode=2 means run in playback mode.

-l <file>

Log file (default console output). Does very little as few messages go to the console. To change the log file for most of the logging, change `log4j.properties`.

-R Disables the Remote API, regardless of the setting in `global.properties`.

-W Start all Configs in the same thread; they do not terminate but wait forever.

-M Start AssemblyLines in simulation mode.

-S This option is for internal use for communication between the Configuration Editor and a server only; it is used to pass Config Files between them. Do not use this option yourself.

-f *extProp1=file1, extProp2=file2*

Where *extProp* is the name of the external Property Store. *file* specifies from where to read the properties. This option specifies a user-defined, external Property Store that can be entered when starting a TDI server. This optional command-line parameter **-f** can be used with the "ibmdisrv" server startup scripts. *extProp* is the name of the external Property Store. *file* specifies from where to read the properties. When the **-f** option is used to specify a properties file from the command line, the server changes the Property Store configuration in memory only, i.e. the server does not make this change permanent by changing the TDI Config file on disk – this change is valid for the current run of the TDI server.

If any property files are specified at the command line, they are valid only for the Config Instances specified with the **-c** command-line option (which are loaded on TDI server startup). The property files specified at the command line do not have any impact on Config Instances which have not been explicitly named with the **-c** command-line option (these can be Config Instances loaded by remote server API client for example).

If a Property Store whose name is specified with the **-f** command-line switch cannot be found in a Config Instance, an error message is logged in the server log (`ibmdi.log` in the Install-directory). When a Property Store name is specified more than once with the **-f** command-line switch then there are two effects: (1) a warning message is logged, and (2) the file specified last will take effect. This feature is implemented in the `com.ibm.di.server.RS` Java class (referenced by way of the main variable when scripting). After the `reload()` method is called, the `MetamergeConfig` object is loaded, and for each Property Store specified on the command line, the corresponding `PropertyStoreConfig` object is updated.

Note:

Although Copy/Paste of Config objects (ALs, Connectors, FCs, and so forth) are fully supported. You can easily copy ALs and components and then paste them into another Config. You can also exchange ALs and components using IM chats, e-mails and text files, because the copy-buffer is filled with the TDI Config XML definition of the selected item. This makes passing stuff around simple and easy, and is a great tool for support and online assistance (for example, ICT/NotesBuddy, forums, ...).

Note:

Make sure you select the entire `<MetamergeConfig>` node in your copy command, including the start and end tags.

-i This option specifies that the TDI server ignores any properties from the `global.properties` file, and reads only the `solution.properties` file. This option can be used when the `global.properties` file is unreadable - for example, when the encoding the TDI server is started with is different from the encoding of `global.properties`.

-? Prints a *usage* message, showing all options in brief.

When IBM Tivoli Directory Integrator ends it returns one of the following exit codes:

- 0 User started IBM Tivoli Directory Integrator with -v parameter (show info and exit)
- 1
 - Cannot open log file (-l parameter)
 - Cannot open Config file
 - Stopped by admin request
- 2 Exit after auto-run. When you start IBM Tivoli Directory Integrator specifying -w, the Server runs the AssemblyLines specified by the -r parameter and then exits.

Note: AssemblyLines run from the Configuration Editor are started in a different way and will not exit with status 2.
- 9 License expired or invalid (obsolete).

Command Line Interface – *tdisrvctl* utility

The Command Line Interface (CLI) to Tivoli Directory Integrator, called the *tdisrvctl* utility, is designed for remotely managing Configs, AssemblyLines, and so on. This utility connects to a remote Tivoli Directory Integrator server using the TDI Remote Server API, and performs the requested operations. As it is a client application interfacing to a Remote Server, it is subject to the same connection, authentication and authorization issues described in Chapter 6, “Security and TDI,” on page 63.

It exposes various command line options for the following functions:

- Start, stop, or reload Tivoli Directory Integrator Configs.
- Start or stop AssemblyLines in a particular config.
- Display a list of configs loaded on the server.
- Shutdown server.
- Display config report.
- Manage config properties through TDI-p, the Tivoli Directory Integrator properties framework
- Send custom notification events.
- View exposed AL Operations.
- View tombstones for terminated Configs and AssemblyLines
- View Tivoli Directory Integrator Server details.

Notes:

1. The command line utility is shipped in the *TDI_install_dir/bin* folder.
2. Remote Method Invocation (RMI) is enabled by default. Therefore, for any remote server API client (including the CLI), the property **api.remote.on** should be set to **true** and the client’s IP address must be mentioned in the property **api.remote.nonssl.hosts** in the *global.properties* (or *solution.properties*) file of the remote Tivoli Directory Integrator Server (if non-SSL mode is being used).
3. The remote Tivoli Directory Integrator server must be running.

Command Line Reference

The command has the following usage:

```
tdisrvctl [general_options] -op operation [operation_specific_options]
```

where **general_options** can be:

-h host	Enter the remote server IP address or hostname (default is localhost).
-K keystore	Enter the name of the SSL key database file.
-p port	Enter the port number (default is 1099).
-P key_pwd	Enter the key file password.
-s	Specify the working directory where the solution directory is located.
-T truststore	Enter the name of the SSL truststore database file.
-u userID	Enter the username (for custom authentication).
-v	Run in verbose mode.
-w user_pwd	Enter the user password (for custom authentication).
-W trust_pwd	Enter the trust file password.
-?	Display command usage.

And **operation** can be:

event	Send custom notification events
prop	Manage Config properties
queryop	Query for AssemblyLine (AL) operations
reload	Reload running Configs
report	Generate Config report or list Configs on remote server
shutdown	Shutdown the server
srvinfo	View TDI server information
status	View status of Configs or ALs
start	Start specific Config or ALs
stop	Stop specific Config or ALs
tombstone	View tombstone entries for specific Config or AL.
deletetombstone	delete a tombstone entry
debug	debug components of a running AssemblyLine

You can display help for any particular option like this:

tdisrvctl -op operation -?

Operations

event Use this option to send custom notification events to a particular server. All listeners registered for the particular event receive this notification. This allows TDI administrators to trigger listener applications based on planned custom events.

The usage for the event operation is:

tdisrvctl [general_options] -op event -e event_name [-s source] [-d data]

where:

-e event_name	The name of the event to send.
-s source	The name of the source invoking the event (default "tdisrvctl").
-d data	The data to be passed to an event listener (default is null).

Example:

To send an event "user.process.X.completed" from "admin".

tdisrvctl -h itditest -op event -e "process.X.completed" -s admin -d "Admin triggered event"

Note: All events sent from tdisrvctl using the **-e** option are prefixed by "user."

prop The "prop" option exposes the properties of a config via the TDI-p. It allows the user to get / set / view the properties of a particular config.

The usage for the prop operation is:

```
tdisrvctl [general_options] -op prop -c config_name
[ -l ] |
[-o property_store]
[-g key | all] |
[-s key=value] [-e] |
[-d key] ]
```

where:

-c config_name	Name the config to work with.
-l	List all the property stores configured.
-o property_store	Name the property store to work with.
-g key	Get the value of the specified key (or keyword 'all' implying get all keys).
-s key=value	Set the "key" to the specified "value."
-e	Encrypt the value when putting in the store (can be used with -s option only).
-d key	Delete the specified "key" from the store.

Notes:

1. The '-l', '-g', '-s', '-d' options are mutually exclusive, and cannot be used together.
2. The '-e' option can only be used with the '-s' option.
3. Managing properties stored in the **password store** is NOT supported.
4. While specifying the "-c" option specify the COMPLETE configuration file path on the remote server, or give a path relative to the "configs" folder. To see the relative paths use the "report" option of tdisrvctl:

```
tdisrvctl -op report -l
```

Examples:

To see a list of all the property stores for config C1.xml

```
tdisrvctl -op prop -c C1.xml -l
```

To get a list of all the properties for config C1.xml

```
tdisrvctl -op prop -c C1.xml -g all
```

To get a list of all the properties for config C1.xml from store MyStore

```
tdisrvctl -op prop -c C1.xml -o MyStore -g all
```

To set a property MY_PROP to value MY_VALUE for config C1.xml in store MyStore and mark it as protected:

```
tdisrvctl -op prop -c C1.xml -o MyStore -s MY_PROP=MY_VALUE -e
```

queryop

The queryop option returns the list of AL operations exposed in an AssemblyLine.

This option is useful in a scripting environment. A Tivoli Directory Integrator solution developer can develop a script to automatically query for exposed operations and then use the result to start an AssemblyLine with a specific operation using the start operation's **-r -alop** flag. The output of this operation is such that it can be grepped for or tokenized easily in a scripted environment.

The usage for the **queryop** operation is:

```
tdisrvctl [general_options] -op queryop -c <configFile> -r <ALname>
```

where

configFile	Config file name
ALName	Name of the AssemblyLine

Output:

```
ALOp:{attr_1;attr_2...attr_n;}
```

Note: While specifying the "-c" option specify the COMPLETE configuration file path on the remote server, or give a path relative to the "configs" folder. To see the relative paths use the "report" option of tdisrvctl:

```
tdisrvctl -op report -l
```

Examples:

To query for *operations* exposed in an AL:

```
tdisrvctl -h itditest -T trust.kdb  
-W secret -op queryop  
-c examples/ADCustomConnector.xml  
-r ADAssemblyLine
```

Example Output:

```
$initialize: {ldapurl;loginPasswd;loginUserName}
```

reload This option can be used to reload running Configs on a particular server.

The usage for reload operation is:

```
tdisrvctl [general_options] -op reload -c [config_list]
```

where:

config_list	Comma separated list of Configs to reload.
--------------------	--

Note: While specifying the "-c" option specify the COMPLETE configuration file path on the remote server, or give a path relative to the "configs" folder. To see the relative paths use the "report" option of tdisrvctl:

```
tdisrvctl -op report -l
```

Example:

To reload Configs C1.xml, C2.xml and C3.xml on remote host itditest:

```
tdisrvctl -h itditest -T trust.jks -W secret -op reload -c C1.xml,C2.xml,C3.xml
```

report This option can be used for generating a report for a particular config or for listing the configs available on the remote server's config folder.

The config report lists details of the particular config. The details are AssemblyLines, Connectors and Parsers in each AssemblyLine, Connector library, Parser library, Script library, Function Library. This option gives a one shot view of all the details of a particular config.

The config listing option helps the user in finding out the list of configs available on the remote server and what their exact names are. Of course, only those configs can be seen that are in the "config" folder of the remote server (see `global.properties` file for property **api.config.folder**). This command cannot obtain list of configs located "anywhere" on the system.

The usage for report operation is:

```
tdisrvctl [general_options] -op report [-c config | -l]
```

where:

-c config Name of the Config whose report is to be generated.
-l The Configs in the remote server's config folder.

The displayed details for each connector or function component part of an AssemblyLine look like this:

```
Name       : count
Mode       : Iterator
State      : Enabled
Debug      : Disabled
Template   : system:/Connectors/ibmdi.Timer
Parser     : [parent]
Comment    : None
```

Notes:

1. The specified config must be already loaded on the remote server.
2. Only one of the '-c' or '-l' option is allowed. Not both.
3. While specifying the "-c" option specify the COMPLETE configuration file path on the remote server, or give a path relative to the "configs" folder.
4. The argument to the -c option is case-sensitive, and must match the name of the config file exactly as known by the server instance, reported by for example "tdisrvctl -op status".

Examples:

To get a complete listing of the details of C1.xml on remote server:

```
tdisrvctl -h remoteserver -op report -c C1.xml
```

To get a list of the configs available in the "config" folder of the remote server:

```
tdisrvctl -h remoteserver -op report -l
```

shutdown

This option can be used to shutdown the Tivoli Directory Integrator server.

The format for this command is:

```
tdisrvctl [general_options] -op shutdown [-o return_code]
```

where:

-o return_code The return code with which the remote Tivoli Directory Integrator server should exit.

Examples:

To shutdown the local Tivoli Directory Integrator server:

```
tdisrvctl -op shutdown
```

To shutdown the server running on remote host itditest which is configured for SSL (server-auth only)

```
tdisrvctl -h itditest -T trust.kdb -W secret -op shutdown
```

srvinfo

This option is used to display the information of a Tivoli Directory Integrator server.

The usage of the command is:

```
tdisrvctl [general_options] -op srvinfo
```

Example:

To view the server information for a Tivoli Directory Integrator server running on localhost

```
tdisrvctl -h localhost -op srvInfo
```

status This option can be used to view status of AssemblyLines.

The usage for status operation is:

```
tdisrvctl [general_options] -op status -c [config_list | all]
-r [AL_list | all]
-listen
```

where:

config_list	Comma-separated list of Configs or keyword "all".
AL_list	Comma-separated list of ALs or keyword "all".
-listen	indicates to start receiving the logs of a running Config or AssemblyLine.

Notes:

1. At least one of the options ('-c', '-r' or '-t') must be specified. -
2. The keyword "all" indicates all configs or AssemblyLines.
3. The -listen option requires exactly one Config or AssemblyLine to be specified.
4. While specifying the "-c" option specify the COMPLETE configuration file path on the remote server, or give a path relative to the "configs" folder. To see the relative paths use the "report" option of tdisrvctl:

```
tdisrvctl -op report -l
```

Examples:

To see the status of all configs and ALs:

```
tdisrvctl [general_options] -op status -c all -r all
```

You could also write

```
tdisrvctl [general_options] -op status
```

To see the status of AL1, AL2:

```
tdisrvctl -h itditest -op status -c cl.xml -r AL1,AL2
```

Output:

```
(Component Type # Component Name # RUNNING / STOPPED # Statistics):
1 # AL1 # RUNNING # [get:571] [add:571] [del:3] [requests:2333]...
1 # AL2 # STOPPED #
```

The Component Types are:

- 0 for Config
- 1 for Assembly line

The Statistics contain the following details (valid for AssemblyLines only):

- Attribute "add" – total number of "add" operations performed
- Attribute "mod" – total number of "modify" operations performed
- Attribute "del" – total number of "delete" operations performed
- Attribute "get" – total number of "getNext" (Iterations) performed
- Attribute "request" – total number of requests accepted when there is a Server mode Connector in the AssemblyLine.
- Attribute "callReply" – total number of "callReply" operations performed
- Attribute "err" – total number of errors encountered
- Attribute "skip" – total number of 'skip' operations performed

- Attribute "lookup" – total number of "lookup" operations performed
- Attribute "ignore" – total number of "ignore" operations performed
- Attribute "reconnect" – total number of "reconnect" operations performed
- Attribute "exception" – the exception text if the component terminated with an exception

To see the details of Configs (running and stopped) on a particular server:

```
tdisrvctl -h itditest -op status -c all
```

5. To see the details of a running AssemblyLine on a particular server and start receiving its logs:

```
tdisrvctl -h itditest -op status -c rs.xml -r all -listen
```

start This option can be used to start a config or AssemblyLines.

The usage for the start operation is:

```
tdisrvctl [general_options] -op start -c [config]  

-e [password]  

-r [AL_list | all] -alop <alop_Name> [{requiredAttr_1; requiredAttr_2; ...  

requiredAttr_n}] | -f filename]  

-s [Simulate mode]  

-m [run name] -o [propStore1=filename1,propStore2=filename2...]  

-t [temp config instance]  

-listen
```

where

-c config	Name of config to start.
-e password	Password of config file if it is encrypted.
-r AL_list	Comma-separated list of ALs to start or keyword 'all'.
-o property file list	comma separated list of property store names and values
-alop operName	The specific AL operation and list of list required attributes for the specified operation.
-f filename	Name of the file where the input attributes and their values are configured for the operation.
-s Simulate mode	Run the specified AssemblyLines in simulate mode
-m multi-instance	Run multiple instances of same Config with different run names
-t temp config instance	Start temp config instance from the XML in the config file specified
-listen	receive the logs of the specified Config or AssemblyLine

Notes:

1. The '-c' option is mandatory. -
2. The keyword "all" indicates all AssemblyLines.
3. Required attributes list is mandatory with -alop option.
4. -alop option cannot be used with -r all option. It works only with a specific AL.
5. When running a temp config with solution or run name it is not possible to check if another config with the same name is already running on the server. If this happens an exception will occur. You could check the running config instances using the **status** command.
6. The -t option expects the Config specified in the -c option to be located on the client machine.
7. If the -t option is used and the config specified in the -c option is relative then it will be searched in the current folder.
8. The -listen option requires exactly one Config or AssemblyLine to be specified.

Examples:

1. To start assembly line AL1 and AL2 of config C1 on remote server itditest:

```
tdisrvctl -h itditest -T trust.kdb -W secret -op start -c C1.xml -r AL1,AL2
```

The `-r` option requires that `-c` option should also be specified. This is because the AssemblyLines mentioned in the command *must* belong to one of the Configs in the `-c` option.

2. To start assembly line AL1 on remote server itditest with AL operation:

```
tdisrvctl -h itditest -T trust.kdb -W secret -op start  
-c examples/ADCustomConnector.xml  
-r ADAssemblyLine  
-alop $initialize {ldapurl##ldap://9.182.190.149:390;loginPasswd##password;loginUsrname##cn=root}
```

3. To start AssemblyLine AL1 on remote server itditest with AL operation update:

```
tdisrvctl -h itditest -T trust.kdb -W secret -op start  
-c examples/ADCustomConnector.xml -r ADAssemblyLine  
-alop search { $init.ldapurl##ldap://9.182.190.149:390;$init.loginPasswd##password;$init.loginUsrname##cn=root;searchBase##o=ibm,c=us}
```

Note: All initialization attributes are to be prefixed with `$init`.

- 4.

```
tdisrvctl -h itditest -T trust.kdb -W secret -op start -c examples/ADCustomConnector.xml -r ADAssemblyLine -alop search -f inputFile
```

Input file format:

=====

Key1:value1

Key2:value2

5. Command to run an AssemblyLine AL1 in simulate mode:

```
tdisrvctl -h itditest -T trust.kdb -W secret -op start -c examples/ADCustomConnector.xml -r AL1 -s
```

6. Command to load multiple config instances:

```
tdisrvctl -op start -c C1.xml -m test -f PropertyStorename=TestProp.properties, PropStore2=propfile2 ... -r AL1,AL2
```

7. Command to run temp config instance:

```
tdisrvctl -op start -c C1.xml -t -r AL1
```

8. Command to start a config on a particular server and receive its logs:

```
tdisrvctl -h itditest -op start -c rs.xml -listen
```

9. Command to start an AssemblyLine on a particular server and receive its logs:

```
tdisrvctl -h itditest -op start -c rs.xml -r AL1 -listen
```

stop The usage for the stop operation is:

```
tdisrvctl [general_options] -op stop -c [config]  
-r [AL_list | all]
```

where:

`-c config`

Name of Config.

`-r AL_list`

Comma-separated list of ALs to stop or keyword "all."

Notes:

1. The `-c` option is mandatory.
2. While specifying the `-c` option specify the COMPLETE configuration file path on the remote server, or give a path relative to the "configs" folder. To see the relative paths use the "report" option of `tdisrvctl`:

```
tdisrvctl -op report -l
```
3. The keyword "all" indicates all AssemblyLines.
4. The `-r` option requires that `-c` option should also be specified. This is because the AssemblyLines mentioned in the command *must* belong to one of the Configs in the `-c` option.
5. The argument to the `-c` option is case-sensitive, and must match the name of the config file exactly as known by the server instance, reported by for example "tdisrvctl -op status".

Example:

To stop assembly line AL1 and AL2, of Config C1 on remote server itditest:

```
tdisrvctl -h itditest -T trust.jks -W secret -op stop -c C1.xml -r AL1,AL2
```

tombstone

This option can be used to view tombstone details of previously run Configs, AssemblyLines and EventHandlers (historical).

The usage for the tombstone operation is:

```
tdisrvctl [general_options] -op tombstone -c [config]  
          [-r [AL_name] ]  
          [-age n]  
          [[attribute_list] | all ]
```

where:

-age n	Tombstone record for the last 'n' days (default is 1 day).
-c config	Name of Config.
-r AL_name	Name of AssemblyLine.
all	Tombstone attributes: show all.

attribute_list:

-ct	Component type.
-cn	Component name.
-guid	Tombstone entry's guid
-et	Event type.
-ex	Exit code.
-stime	Component's start time.
-ctime	Tombstone create time.
-desc	Error description.
-um	User message.
-stat	Statistics (valid for ALs only).

Notes:

1. The '-c' option is mandatory.
2. While specifying the "-c" option specify the COMPLETE configuration file path on the remote server, or give a path relative to the "configs" folder. To see the relative paths use the "report" option of tdisrvctl:

```
tdisrvctl -op report -l
```
3. The argument to the -c option is case-sensitive, and must match the name of the config file exactly as known by the server instance, reported by for example "tdisrvctl -op status".

Examples:

1. To see the last 2 days tombstone entries (all attributes) for config C1.xml

```
tdisrvctl [general_options] -op tombstone -c C1.xml -age 2 all
```
2. To see tombstone entries for config C1 for the past 3 days:

```
tdisrvctl -h itdiserver -op tombstone -c C1 -age 3 all
```
3. To see tombstone entries for config C1 for the last 24 hours (specific attributes):

```
tdisrvctl -h itdiserver -op tombstone -c C1 -ct -ctime -cn -um
```
4. To see the tombstone entry for AL1 of "rs.xml"

```
tdisrvctl -h itdiserver -op tombstone -c C1 -r AL1
```

deletetombstone

This option can be used to delete tombstone entries for previously run AssemblyLines.

The usage for the delete tombstone operation is:

```
tdisrvctl [general_options] -op deletetombstone -guid <GUID number>
```

where

-guid "GUID number" is the unique identifier for the tombstone to be deleted. The GUID for a tombstone can be obtained by viewing the contents of the tombstone; see the entry about the tombstone option for details as to how to obtain the GUID.

debug This option can be used to set the Debug mode values of connectors and function components of a running AssemblyLine. When you set the Debug mode of a connector with specified parser, the Debug mode of the parser is also initialized with the same value.

The usage for the debug operation is as follows:

```
tdisrvctl [general_options] -op debug -c config  
          -r assembly_line  
          [-alc al_component]  
          -on/off
```

where:

-c config	Name of Config.
-r assembly_line	Name of AssemblyLine.
-alc al_component	name of the AssemblyLine component.
-on	flag to enable debug.
-off	flag to disable debug.

Notes:

1. The '-c' and '-r' options are mandatory and require exactly one Config/AssemblyLine to be specified.
2. While specifying the "-c" option specify the COMPLETE configuration file path on the remote server, or give a path relative to the "configs" folder. To see the relative paths use the "report" option of tdisrvctl:
tdisrvctl -op report -l
3. The argument to the -c option is case-sensitive, and must match the name of the config file exactly as known by the server instance, reported by for example "tdisrvctl -op status".
4. If the **-alc** option is not specified all components in the specified AssemblyLine will be affected.

Examples:

1. To view the Debug mode value of the components in the AssemblyLines of a specified Config:
tdisrvctl -op report -c C1
2. To enable the debug mode for all components in the running AssemblyLine al2:
tdisrvctl -op debug -c C1-r al2 -on C1-r al2 -on
3. To disable the debug mode for specified components in the running AssemblyLine al3:
tdisrvctl -op debug -c C1-r al3 -alc comp1,comp2 -off

Other points to note

- If the user specifies the **-T** option or the **-K** option, it means the command line utility must use SSL.

- If no **-h** (host) option is specified, the command line interface searches for the environment variable **TDI_RSRV**. If **TDI_RSRV** is not set or empty, then it uses "localhost" as default. This is also the case for the **-p** (port) option: if **-p** is not specified then it searches for **TDI_RPORT**, and if that is not specified either, it uses the default of "1099".
- The command returns "0" indicating that the command completed successfully without any errors. A "-1" is returned otherwise. For instance a command asking for starting 3 AssemblyLines returns 0 only if all 3 AssemblyLines started successfully, otherwise it returns -1.
- The **tdisrvctl** command line utility will use Log4J logging APIs for logging error messages. The Log4J configuration file is specified in the startup script (the .bat or .sh) file. The command uses a file called **tdisrvctl-log4j.properties** to set up the Log4J logging. If the solution directory is specified the command sets an environment variable for pointing to the log configuration file in the solution directory. If the solution directory is not specified then the command uses the log configuration file present in the install directory.
- The **tdisrvctl-log4j.properties** file has the complete path of location where the logs are to be created. The log files are created in the *TDI_install_dir/logs* directory by default. The location can be customized as needed.
- All reported error and warning messages are displayed with an error code prefix. This error code can be used to search the *IBM Tivoli Directory Integrator V7.0 Messages Guide* for an explanation of the error message and operator response.

Chapter 14. Logging and debugging

Tivoli Directory Integrator uses a logging class to record messages to a number of various log channels. All Tivoli Directory Integrator components use this logging class which in turn invokes an industry standard logging tool (Log4J). While Log4J provides a variety of output channels and formats, there are other logging utilities with overlapping and additional output channels that you as a Tivoli Directory Integrator user may need. Many of these are open source libraries that are not bundled with Tivoli Directory Integrator for legal reasons. To enable inclusion of these 3rd party logging utilities, the Tivoli Directory Integrator logging component is modeled to act as a proxy between Tivoli Directory Integrator and the actual logging implementations, called LogInterface implementations. Refer to the section "Creating additional Loggers" in *IBM Tivoli Directory Integrator V7.0 Reference Guide* for more information on how to create, configure and program your own LogInterface classes.

Note: Enable or disable logging in Tivoli Directory Integrator by configuring the `com.ibm.di.logging.enabled` property. To enable logging, use `com.ibm.di.logging.enabled=true` (default). To completely disable logging, use `com.ibm.di.logging.enabled=false`.

The remainder of this section describes how to use the logging class that is bundled with Tivoli Directory Integrator, called `com.ibm.di.log.TDILog4J`.

Logging and debugging by the system is mainly done through the Task object (the current AssemblyLine). Logging can either be done explicitly (in script) or done by the various components themselves.

The Log4J logging engine is a very flexible framework that lets you log to file, eventlog, syslog and more, and can be tuned so it suits most needs. It can be a great help when you want to troubleshoot or debug your solution. By means of the aforementioned logging class, Tivoli Directory Integrator has additional tracing facilities (discussed in Chapter 15, "Tracing and FFDC," on page 179), though in most cases, the logging functionality described here suffice.

Some Tivoli Directory Integrator components may have very specific troubleshooting guidelines; always check the particular component's section in the *IBM Tivoli Directory Integrator V7.0 Reference Guide* and the *IBM Tivoli Directory Integrator V7.0 Problem Determination Guide* for more information.

The log scheme for the server (ibmdisrv) is described by the file `Log4J.properties` and elements of the Config file, see "Log4J default parameters" on page 176.

Note: Any of the aforementioned properties files can be located in the Solution Directory, in which case the properties listed in these files override the values in the file in the installation directory.

You can create your own appenders to be used by the Log4J logging engine by defining them in the `Log4J.properties` file. You can use drivers built-in to Log4J like the default one, which is defined with the statement:

```
Log4J.appender.Default=org.apache.Log4J.FileAppender
```

The phrase `org.apache.Log4J.FileAppender` defines this appender to use the `FileAppender` class. Additional Log4J compliant drivers are available on the Internet, for example drivers that can log using JMS or JDBC. In order to use those, they need to be installed into the IBM Tivoli Directory Integrator installation jars directory after which appenders can be defined using those additional drivers in `Log4J.properties`. For more information, refer to <http://jakarta.apache.org/log4j/docs>.

In addition to the IBM Tivoli Directory Integrator built-in logging, you can log by adding script code in your AssemblyLine. This is described in much more detail in the *IBM Tivoli Directory Integrator V7.0 Users Guide*, in which you also find out how the interactive debugger works.

Script-based logging

You can issue messages to the AssemblyLine's configured loggers at any time using JavaScript, at any point where scripting is possible (hooks, script components, and so on.) The explicit `logmsg()` calls available to you (that is, `task.logmsg()` & `main.logmsg()`) can have an optional string parameter indicating the Log4J level at which the messages are to be logged. Default is INFO. If the log-level given by the user is invalid for Log4J, the message is logged at DEBUG level. Levels include DEBUG, INFO, WARN, ERROR, FATAL.

If you use

```
task.logmsg()
```

your messages will be logged along with the other messages from the AssemblyLine. If you are running your AssemblyLine from the Configuration Editor, that will be in the CE output window. If your AssemblyLine also uses other logging methods, the messages will be there too.

When you use

```
main.logmsg()
```

your message will be logged along with other messages from the Config Instance. This will be in the log file(s) or other loggers created by the Config Instance, which are typically not seen in the Configuration Editor.

Logging using the default Log4J class

Configuring the default logging of IBM Tivoli Directory Integrator, which uses Apache Log4J is done globally (using the file `Log4J.properties` which specifies global defaults for Server tasks) or specifically, using the Configuration Editor, for each AssemblyLine or Config File as a whole. To provide this level of flexibility and customization, the Java Log4J API is used.

Only the parameters that describe how messages are logged are described here.

All log configuration windows operate in the same way: For each one you can set up one or more log schemes. These are active at the same time, in addition to whatever defaults are set in the `Log4J.properties` file; see "Log4J default parameters" on page 176.

Many (but not all) loggers support a Character Encoding option, to control what character set the log files are written in. There are many different character sets; for an informal overview check <http://java.sun.com/j2se/1.5.0/docs/guide/intl/encoding.doc.html>.

The possible log schemes are as follows:

FileRollerAppender

Sometimes, you want to log to file but keep a limited number of files, as they can fill your disks. `FileRollerAppender` generates a new file for each run of the Server. The system saves only the specified number of previous logs. If your log is called `mylog.txt`, and you ask for 2 generations, then after 3 runs you have a `mylog.txt` (last run) as well as the files `mylog.txt.1` and `mylog.txt.2`, where `mylog.txt.2` is the oldest log. From this point, you do not get more files, only newer versions with the same name. Keep two generations of backup files.

`FileRollerAppender` has the following parameters:

File Path

The name of the file to log to. The path is relative to where you installed IBM Tivoli Directory Integrator. The special macro {0} used in filenames is replaced by the name of the Server. Similarly, {1} used in filenames is replaced by a unique identifier generated by the system for you. The {1} macro has no relevance for the special case where you use FileRollerAppender, but is important where you want unique file names.

Number of backup files

If your File Path was mylog.txt, and you select 2 backup-files, the two previous runs have their files renamed to mylog.txt.1 and mylog.txt.2 when you run a third time.

Layout

Determines the format of the log message. Options are:

- Pattern (used if you want to customize the way the messages are logged)
- Simple (format containing just the loglevel and the message)
- HTML (creates an HTML file containing some (relative) time info, thread info, loglevel, category, and message)
- XML (similar to HTML, but generates an XML file (using namespace-prefix Log4j))

Pattern

Only used when **Layout** is **Pattern**. See “Creating your own log strategies” on page 177.

Log level

Severity level of the log messages. Options are (from maximum to minimum information):

- DEBUG
- INFO
- WARN
- ERROR
- FATAL

Character Encoding

Character Encoding to be used; like Cp1252, ISO-8859-1, and so on.

Log Enabled

Click to enable the use of this Appender.

ConsoleAppender

Logs to the console (standard output). This is in the window where you started the server (ibmdisrv) or the execute task-window in the Configuration Editor (ibmditk). **Console** has the following parameters:

Layout

See **FileRollerAppender**, previous.

Pattern

See **FileRollerAppender**, previous.

Log level

See **FileRollerAppender**, previous.

Log Enabled

See **FileRollerAppender**, previous.

FileAppender

Logs to a file. **File** has the following parameters:

File Path

See **FileRollerAppender**, previous.

Append to file

Click to append log information to file. If option is not enabled, the file is overwritten.

Layout

See **FileRollerAppender**, previous.

Pattern

See **FileRollerAppender**, previous.

Log level

See **FileRollerAppender**, previous.

Character Encoding

Character Encoding to be used; like Cp1252, ISO-8859-1, and so on.

Log Enabled

See **FileRollerAppender**, previous.

This is the appender set up by default; see “Log4J default parameters” on page 176.

SyslogAppender

Enables IBM Tivoli Directory Integrator to log on UNIX Syslog. **Syslog** has the following parameters:

Host name/IP Address

Host to log on to.

Syslog Facility

Legal facilities found in the drop-down. Must be supported by the host you are logging to.

Print Facility String

If set, the printed message includes the facility name of the application.

Layout

See **FileRollerAppender**, previous.

Pattern

See **FileRollerAppender**, previous.

Log level

See **FileRollerAppender**, previous.

Log Enabled

See **FileRollerAppender**, previous.

NTEventLog

Enables applications to log to the Windows NT Event Log (on Windows platforms). **NTEventLog** has the following parameters:

Source

The "source" name appearing in the NT event log; usually the title of the application doing the logging.

Layout

See **FileRollerAppender**, previous.

Pattern

See **FileRollerAppender**, previous.

Log level

See **FileRollerAppender**, previous.

Log Enabled

See **FileRollerAppender**, previous.

DailyRollingFileAppender

The daily rolling file appender rotates the log file every day. When the output file is rolled it is given a name consisting of the base name plus a date pattern string; that is, filename.yyyy-mm-dd. It usually is used with the **Append to file** parameter set to **true**. **DailyRollingFile** has the following parameters:

File Path

See **FileRollerAppender**, previous.

Append to file

Create new file or append to existing file, depending on whether this is checked. You usually want this on when using the **DailyRollingFile**.

Date Pattern

How often the file is rotated. Use the drop-down to choose resolution from minutes to months. For example, if the File Path is set to example.log and the DatePattern set to ' . 'yyyy-MM-dd, on 2003-10-31 at midnight, the logging file example.log is copied to example.log.2003-10-31. Logging for 2003-11-01 continues in example.log until it rolls over the next day.

Layout

See **FileRollerAppender**, previous.

Pattern

See **FileRollerAppender**, previous.

Log level

See **FileRollerAppender**, previous.

Character Encoding

Character Encoding to be used; like Cp1252, ISO-8859-1, and so on.

Log Enabled

See **FileRollerAppender**, previous.

Also see the example under “Log4J default parameters” on page 176.

SystemLogAppender

This Appender creates log files in a catalog hierarchy under *TDI_install_dir/system_logs*. For each Config File, there is a corresponding directory with logfiles named *AL_xxx*, where xxx is the name of the AssemblyLine being run.

This Appender has the following parameters:

Pattern

Specifies the format of the log as defined by LOG4J. The default value is:

```
"%d{ISO8601} %-5p [%c] - %m%n"
```

Additional values available in the field are:

```
"%d{HH:mm:ss} %p [%t] - %m%n"
```

```
"%p [%t] %c %d{HH:mm:ss,SSS} - %m%n"
```

Log level

See **FileRollerAppender**, previous.

Character Encoding

Character Encoding to be used; like Cp1252, ISO-8859-1, and so on.

Log Enabled

See **FileRollerAppender**, previous.

Log Levels and Log Level control

Log levels can be

- ALL
- DEBUG
- INFO
- WARN
- ERROR
- FATAL
- OFF

ALL logs everything. DEBUG, INFO, WARN, ERROR and FATAL have increasing levels of message filtration. Nothing is logged on OFF.

You can issue log messages to the system or AssemblyLine logs by using the `logmsg()` method from JavaScript, wherever Tivoli Directory Integrator allows scripting. It can take one or two parameters. See the Java API documentation for the `logmsg()` declaration (package **com.ibm.di.server**, class *AssemblyLine* or class *RS*).

The interface for the `logmsg()` method (both main and task) with additional log level parameter is **logmsg (String logLevel, String msg)**. The legal values for `logLevel` are: "FATAL", "ERROR", "WARN", "INFO", "DEBUG", corresponding to the log levels available for log Appenders. Any unrecognized value is treated as "DEBUG".

Note that the IBM Tivoli Directory Integrator `logmsg()` JavaScript calls log on INFO level by default. This means that setting `loglevel` to WARN or lower silences your `logmsg` as well as all Detailed Log settings. However, with the level parameter to the `logmsg()` call you can override the log level for individual `logmsg()` calls.

Log4J default parameters

When Tivoli Directory Integrator is installed, a *FileAppender* is used for the default logger. If you want to change the default logger you must change the content of the `log4j.properties` file situated in the *TDI_installdir/etc* folder. The default configuration is as follows:

```
# This is the default logger, you will see that it logs to ibmdi.log
log4j.appender.Default=org.apache.log4j.FileAppender
log4j.appender.Default.file=logs/ibmdi.log
log4j.appender.Default.layout=org.apache.log4j.PatternLayout
log4j.appender.Default.layout.ConversionPattern=%d{ISO8601} %-5p [%c] - %m%n
log4j.appender.Default.append=false
```

The *FileAppender* logger truncates the content of the `ibmdi.log` file (situated in *TDI_installdir/logs*) each time the Tivoli Directory Integrator server is started. If want to change that behavior you must change the **log4j.appender.Default.append** property to true.

In the `log4j.properties` file you can find also two examples for changing the default logger to *RollingFileAppender* or *DailyRollingFileAppender*. If you want to use them just uncomment the preferred one and comment the *FileAppender* logger:

```
#####ROLLING FILE SIZE APPENDER
##RollingFileAppender rolls over log files when they reach a certain size specified by the
##MaxFileSize parameter

#log4j.appender.Default=org.apache.log4j.RollingFileAppender
#log4j.appender.Default.File=logs/ibmdi.log
#log4j.appender.Default.Append=true
#log4j.appender.Default.MaxFileSize=10MB
```

```
#log4j.appender.Default.MaxBackupIndex=10
#log4j.appender.Default.layout=org.apache.log4j.PatternLayout
#log4j.appender.Default.layout.ConversionPattern=%d{ISO8601} %-5p [%c] - %m%n

#####DAILY OUTPUT LOG4J SETTINGS
## With the DailyRollingFileAppender the underlying file is rolled over at a user chosen frequency.
##The rolling schedule is specified by the DatePattern option

#log4j.appender.Default=org.apache.log4j.DailyRollingFileAppender
#log4j.appender.Default.file=logs/ibmdi.log
#log4j.appender.Default.DatePattern='.'yyyy-MM-dd
#log4j.appender.Default.layout=org.apache.log4j.PatternLayout
#log4j.appender.Default.layout.ConversionPattern=%d{ISO8601} %-5p [%c] - %m%n
```

These are some of the parameters you find in the file Log4J.properties (for ibmdisrv and ibmditk).

Full documentation can be found at <http://jakarta.apache.org/log4j/docs>.

Log4J.rootCategory=DEBUG, Default

DEBUG is the loglevel for the named Appender (Log4J term called Default). If you set the loglevel to OFF or to the level above INFO, you do not get output from your script logmessages (see the following log terms):

Log4J.appender.Default

Defines what type of Appender the named appender Default is. It can be one of the following:

- FileRollerAppender (generates a new file for each run of the Server)
- ConsoleAppender (log to console)
- FileAppender (log to file)
- SyslogAppender (log to UNIX Syslog)
- NTEventLog (log to Windows NT EventLog)
- DailyRollingFileAppender (saves old files with a datestamp in their names)
- SystemLogAppender (In a folder structure under *root_directory/system_logs*)

Log4J.appender.Default.file

Default log file for FileAppender, relative to your installation directory (default ibmdi.log).

Log4J.logger.com.ibm.di.*

Log level of various IBM Tivoli Directory Integrator components. Note that, for example, ibmditk shows the log level of the IBM Tivoli Directory Integrator Configuration Editor itself (not the processes you are running inside it). Do not change these.

Creating your own log strategies

You can use this framework to differentiate how the different AssemblyLines log.

Note: This information is intended for users who want to continue using the `global.properties` file to customize logging output. You can customize logging output through the Configuration Editor (ibmditk).

The following section defines a log scheme called CONSOLE, which later can be used by specific AssemblyLines:

```
Log4J.appender.CONSOLE=org.apache.Log4J.ConsoleAppender
Log4J.appender.CONSOLE.layout=org.apache.Log4J.PatternLayout
Log4J.appender.CONSOLE.layout.ConversionPattern=%d [%t] %-5p - %m%n
```

Now in order to have the AssemblyLines myAL use this, you need the lines:

```
Log4J.logger.AssemblyLine.myAL=INFO, CONSOLE
```

Refer to the full Log4J (version 1.2) documentation for description of the ConversionPattern parameters. Here are some parameters:

%d Date/time depending on format.

%p Priority.

%c Category.

Note: this is typically in the form *Type.alName.xxx*. *Type* can be EventHandler or AssemblyLine, *alName* is the name of the AssemblyLine (or EventHandler as named by the creator), and *xxx* is a unique ID for the thread. **%c{2}** outputs *alName & unique ID*.

%m Message.

%n Newline.

%t Threadname.

Chapter 15. Tracing and FFDC

In addition to the user-configurable logging functionality described in Chapter 14, “Logging and debugging,” on page 171, IBM Tivoli Directory Integrator is instrumented throughout its code with tracing statements, using the JLOG framework. This is a logging library similar to Log4J, but which is used inside TDI specifically for tracing and First Failure Data Capture (FFDC). To which extent this becomes visible to you, the end user, depends on a number of configuration options in the global configuration file `jlog.properties`, and the Server command line option `-T`.

Note: Normally, you should be able to troubleshoot, debug and support your solution using the logging options described in the chapter, Chapter 14, “Logging and debugging,” on page 171. However, when you contact IBM Support for whatever reason, they may ask you to change some parameters related to the tracing functionality described here to aid the support process.

Tracing Enhancements

Currently most Connectors and Parsers have entry and exit trace statements. For 7.0, a number of classes on the TDI server have trace statements added to:

- Method entry and exit points.
- Interactions with third party software.
- Thread creations.

Understanding Tracing

Tracing is done in the code of Tivoli Directory Integrator using JLOG's PDLogger object. PDLogger or the Problem Determination Logger logs messages in Logxml format (a Tivoli standard), which IBM Support understands and for which they have processing tools.

The basic level of information traced, as handled by the PDLogger APIs, is:

Date | Time | ClassName | methodName | MachineName | IP | {Entry/Exit/Exception} | [Parameter]

Basic tracing information means Time, Level (Min, Mid, Max), Location in code, that is Method name and Entry/Exit. The “|” character serves a documentation purpose only, it is not part of the actual log.

Tracing is not performed using Log4J Appenders for the following reasons:

1. Trace is always to be enabled
2. You wouldn't want multiple traces enabled in the server (could be several for each AL if Appenders were used).

The PDLogger is attached to the JLOG **SnapMemory** handler and the **JlogSnapHandler**.

The **SnapMemory** Handler logs trace messages to memory. On the trigger of a LogEvent (that is, an occurrence of a specific Log level Trace message, as defined by the `jlog.levelflt.level` filter, or an application crash or on the occurrence of a specific TMS XML messageID) the Trace memory buffer is written to a file by the **JlogSnapHandler**.

To make Tracing and Log messages in TDI unique across all IBM products, they are prefixed with a unique prefix: **CTGDI**.

All error messages are prefixed with a unique TMSXML messageID that indicates the cause of the error and an operator response.

All info messages are also prefixed with a unique TMSXML messageID that may or may not provide the operator response.

Configuring Tracing

The `jlog.logger.level` property in the `jlog.properties` file can be used to set the desired trace level. The trace level can be set to any of the following JLOG log level (hierarchy, from most severe to least severe):

- FATAL
- ERROR
- WARNING
- INFO
- DEBUG_MIN
- DEBUG_MID
- DEBUG_MAX

The default level is FATAL.

Default Trace level as well as whether Tracing is done to file or memory is defined in the default `jlog.properties` file. This file is placed in the `TDI_install_dir/etc` folder. If you use a solution directory, it is placed in the `TDI_Solution_dir/etc` folder.

Setting trace levels dynamically

Tivoli Directory Integrator ships a `LogCmd.bat` (for Windows) and `LogCmd.sh` (for Unixes) scripts. By using them the Trace properties can be set dynamically. JLOG logger starts a command server on the default port (9992) to listen for log commands sent by the `logcmd` command line utility.

For the `logcmd` scripts to work, the command server needs to be started first. To start the log command server, you need to set `jlog.noLogCmd=false` in the `jlog.properties` file.

The listen port of this server can be changed by setting the `jlog.logCmdPort` property in the `jlog.properties` file to the desired value. For more information about these properties read the comments in the `jlog.properties` file.

Usage of the `logcmd` command is as follows:

```
logcmd -o port_number { [-h] | [help] |  
    [list {node_name} ] |  
    [config node_name] |  
    [set node_name key_name=value |  
    [remove node_name {key_name} ] |  
    [dump handler_name] | [save {all} ] }
```

where

-o port_number

The port number to use to connect to the log command server. If not specified the default port (9992) is assumed.

-h | help

Displays syntax information for the command.

list Lists the names of all known logging objects (nodes).

list node_name

Lists the names of the children of the node name. Not all logging objects have children.

config node_name

Lists the all the configuration properties for the node.

set node_name key_name=value

Sets a property key for the node name. If the logging object, node_name, does not exist, the logging object is created and the property is added.

remove node_name

Removes the configuration object node_name. A logging object that has been instantiated from this configuration is not affected by removing the configuration node.

remove node_name key_name

Removes the configuration property key_name from the logging object node_name. If the object supports a hierarchical inheritance of properties, a subsequent logcmd config node_name command may show the key just removed. In that case, it was inherited from an ancestor.

save {all}

Saves the logging configuration to persistent store. If **all** is specified, the entire configuration is saved; otherwise, only those configuration nodes that were originally loaded from the file are saved.

Useful JLOG parameters

Property	Value	Description
jlog.snapmemory.queueCapacity	Default 10000	The number of logevents that can be stored in the snapmemory handlers queue.
jlog.snapmemory.dumpEvents	true	The handler immediately sends all the queued events to its output listeners when the property is set to true. The property can then be reset to false.
jlog.snapmemory.userSnapDir	CTGDI/FFDC/user/	The directory to place the trace dump file when a user triggers an FFDC action by using the logcmd scripts.
jlog.snapmemory.isSync	Default false	The log events are dumped to the snap shot file synchronously when the property is set to true. This does not spawn a new thread, and causes the logger to block until the snapshot is complete.
jlog.snapmemory.userSnapFile	userTrace.log	
jlog.snapmemory.triggerFilter	jlog.levelflt	The level filter to be used to take JFFDC action.
jlog.snapmemory.msgIds	*E	The TMSXML message filter to be used for JFFDC action.
jlog.snapmemory.mode	PASSTHRU or BLOCK. Default is PASSTHRU.	The listed IDs are blocked when the msgIDs property is set to BLOCK. When set to PASSTHRU, the listed IDs are sent to the filter.
jlog.snapmemory.msgIDRepeatTime	10000 (in milliseconds)	The minimum time in milliseconds, after passing a logEvent with a given TMS message ID, before another logEvent with the same id can be passed.

The default value for jlog.snapmemory.triggerFilter sets up a trigger filter named jlog.levelflt. An attribute of such a filter is the message severity, which takes one of the JLOG Log values as described above. By default, the entries

```
jlog.levelflt.className=com.ibm.log.LevelFilter
jlog.levelflt.level=FATAL
```

set up the FFDC code to cause the memory buffer to be dumped to the trace log when a trace message of severity FATAL occurs. The jlog.levelflt.level property can take any of the other Log level values as well, but only values of ERROR or FATAL make much sense as otherwise the amount of FFDC dumping is very high, causing huge slowdowns of the TDI Server.

Chapter 16. Administration and Monitoring

The IBM Tivoli Directory Integrator (TDI) 7.0 Administration and Monitoring Console (AMC) user interface is deployed into the Integrated Solutions Console (ISC). Use the AMC to start, stop and manage Tivoli Directory Integrator Configs and AssemblyLines remotely.

IBM Tivoli Directory Integrator 7.0 also ships an Action Manager with the AMC. The Action Manager is a stand-alone Java application that interacts with the AMC database and uses the Remote Server API to manage remote AssemblyLines.

The Administration and Monitoring Console is comprised of a Java WAR file (Web Archive) and a WAB file (Web Bundle) that can be deployed on ISC SE and ISC AE.

The current Action Manager, bundled with IBM Tivoli Directory Integrator 7.0 AMC, supports Tivoli Directory Integrator 7.0, TDI 6.1.X and TDI 6.0. Note that the Action Manager for IBM Tivoli Directory Integrator 7.0 supports TDI 6.1.X and TDI 6.0 with some restrictions. Versions of Tivoli Directory Integrator prior to version 6.0 are not supported.

Note: IBM Tivoli Directory Integrator 7.0 and solutions developed and deployed with it can also be monitored with IBM Tivoli Monitoring (ITM) Server and Portal, by virtue of Tivoli Directory Integrator's Java Management Extension (JMX) interface. You can find an example that show how you can accomplish this on IBM's Tivoli Open Process Automation Library (OPAL): Managing TDI with ITM using JMX.

Installation and Configuration

See “Installing a new version of Tivoli Directory Integrator 7.0” on page 13 for information about installing Tivoli Directory Integrator and the Administration and Monitoring Console. Installing AMC also installs the Action Manager. If you choose a custom ISC AE to deploy AMC or defer deployment of AMC during installation, see “Deploying AMC to a custom ISC SE/AE” on page 27 for information on additional deployment requirements.

Deploying AMC into the Integrated Solutions Console (ISC)

These instructions require that you be familiar with the IBM Tivoli Directory Integrator 7.0 Installation procedures. See “Using the platform-specific TDI installer” on page 12 for information about Tivoli Directory Integrator Installation. Installing AMC also installs the Action Manager.

If you want to deploy Tivoli Directory Integrator Administration and Monitoring Console into the ISC automatically, select one of the following options:

- Embedded instance of ISC SE
- Existing instance of ISC

If you do not want to deploy AMC into ISC automatically at installation time, select "Do not specify. I will manually deploy AMC at a later time."

The installer automatically installs ISC and deploys AMC in it if you select "Embedded instance of ISC SE" or "Existing instance of ISC" during IBM Tivoli Directory Integrator 7.0 installation.

To install and deploy the Administration and Monitoring Console into the ISC SE:

1. Invoke the Tivoli Directory Integrator 7.0 installer.
2. During installation, select the **Custom** installation. (Typical installation does not offer the AMC option.)

3. On the "Select Features" window of the installation, select **AMC: Administration Monitoring Console** and **embedded Web platform** (includes Integrated Solutions Console, Standard Edition (ISC SE)).
4. Finish installing Tivoli Directory Integrator 7.0.

Deploying AMC as a Windows service or UNIX process using the TDI installer

You can register AMC as a Windows service or UNIX process if the following conditions are satisfied:

Note: You cannot register AMC as a service on i5/OS®.

- The person installing Tivoli Directory Integrator must have administrative permissions (Administrators group on Windows or root on UNIX).
- You have selected to install AMC into an "Embedded instance of ISC Standard Edition."
- You have selected "Register AMC as a system service" and given the service a name. The default service name is `tdiamc`

Starting the Administration and Monitoring Console and Action Manager and logging in

You can start and stop the Administration and Monitoring Console and Action Manager by running the following scripts shipped in the `TDI_install_dir/bin/amc` folder:

- To start AMC, run the `start_tdiamc` script.
- To start AM, run the `startAM` script.

For more information about these scripts, see "AMC and AM Command line utilities" on page 220.

The above will start AMC and AM using a Derby database configured locally on the machine running in network mode on localhost at port 1528. For more information on alternate settings and configurations for both AMC and AM, see sections "Enabling AMC" and "Enabling Action Manager" on page 193.

Once the Administration and Monitoring Console is started, you can access it from the following URL: <http://localhost:1528/ibm/console>; for more information, see "Log in and logout of the console" on page 198.

Stopping AMC and AM can be accomplished by running the `stop_tdiamc` and `stopAM` scripts, respectively.

Notes:

1. For information on adding users and user roles, see section "Console user authority" on page 188.
2. For information on AM, see section "Action Manager" on page 189.
3. For information on usage of individual panels in AMC, see the online panel help, or section "Administration and Monitoring Console User Interface" on page 198.
4. You will be registering Tivoli Directory Integrator Servers in AMC so that solutions from Tivoli Directory Integrator configurations can be administered. AMC will only be able to find Configs that each Tivoli Directory Integrator server has loaded when its started up. By default, Tivoli Directory Integrator 7.0 Servers have the remote server API enabled. Ensure that your `TDI_install_dir/configs` folder has the Configs you want to administer and monitor (or put them in the config folder of your solution directory if your server is using a solution directory).
5. For an example walk through of on using AMC and Action Manager for a few simple tasks, see "Example walkthrough of creating a Solution View and Rules" on page 225.

Enabling AMC

The configuration file for the Administration and Monitoring Console is the `amc.properties` file that is located at the same level as the `WEB-INF` directory. This file contains the AMC's database configuration properties, LDAP properties, SSL related properties and help server details.

By default, the Administration and Monitoring Console makes use of Derby version 10 to store data. When AMC is started for the first time, AMC creates a `tdiamcdb` folder inside the Web server directory and creates the tables needed for AMC to function. The Derby database can be accessed in either the network mode or embedded mode. By default, AMC is shipped with Derby configured in network mode. The following properties in `amc.properties` are applicable to Derby configured for network mode:

```
com.ibm.di.amc.jdbc.database=jdbc:derby://localhost:1528/tdiamcdb;create=true
com.ibm.di.amc.jdbc.driver=org.apache.derby.jdbc.ClientDriver
com.ibm.di.amc.jdbc.urlprefix=jdbc:derby:
com.ibm.di.amc.jdbc.user=APP
com.ibm.di.amc.jdbc.password=APP
com.ibm.di.amc.jdbc.start.mode=automatic
com.ibm.di.amc.jdbc.host=localhost
com.ibm.di.amc.jdbc.port=1528
com.ibm.di.amc.jdbc.sysibm=true
```

The property `com.ibm.di.amc.jdbc.database` points to Derby in network mode, running on `localhost:1528`. The database name being accessed is `tdiamcdb`, and `create=true`, indicating that AMC should create the database if not found.

You should change the `create=true` to `create=false` once your environment is set, so that in case the database path gets modified, AMC does not re-create the database, but instead throws a "Database not found" exception. You should also ensure that the database path be set to an absolute path to avoid any confusion about the database path later.

Other databases than Derby can be configured by setting the appropriate properties; see "JDBC Properties" on page 204.

AMC provides a separate startup and shutdown script for Action Manager. AMC allows the Action Manager to run remotely and provides a separate Derby start or shutdown script.

The Administration and Monitoring Console can also be configured to connect to the Derby database in Embedded Mode. In this case, the Action Manager, a separate application that also talks to the AMC database, is unable to connect to AMC's database. This is because in Embedded Mode, only one JVM at a time is allowed to connect to the Derby database. The following example shows the `amc.properties` file with Derby configured for embedded mode:

```
##Location of the database (embedded mode)
configured for embedded mode:
##Location of the database (embedded mode)
com.ibm.di.amc.jdbc.database=tdiamcdb
com.ibm.di.amc.jdbc.driver=org.apache.derby.jdbc.EmbeddedDriver
com.ibm.di.amc.jdbc.urlprefix=jdbc:derby:
com.ibm.di.amc.jdbc.user=APP
com.ibm.di.amc.jdbc.password=APP
```

The `com.ibm.di.amc.jdbc.database` property points to the location of the AMC database. We suggest that this value be set to an absolute path to avoid any confusion about the database path later.

Running Action Manager remotely

Beginning with Tivoli Directory Integrator 7.0, you can run Action Manager remotely without starting the AMC first. The database for AMC, **Derby**, must be running in Network mode in order for Action Manager to connect to it. Tivoli Directory Integrator 7.0 also provides start and shutdown scripts for the Derby data store so that a user can start Action Manager remotely without starting the AMC.

Notes:

1. Before you start Action Manager for the first time, you must have run AMC at least once. This is because AMC creates the necessary database tables required for AM.
2. You can find the scripts in this section in the following folder of the Install Directory of the remote computer: `TDI_install_dir\bin\amc\ActionManager\`.

3. The instructions in the startup and shutdown sections that follow are for Action Manager and Derby running on different remote computers, and with AMC not running.
4. To verify that AMC, Action Manager or Derby has stopped, check the logs.

AMC and Action Manager startup: If you want to run Action Manager and Derby with AMC running, start AMC by typing `start_tdiamc.bat(sh)` and start the Action Manager by typing `startAM.bat(sh)`. The `tdiamc` script calls the `startNetworkServer.bat(sh)` script, thereby starting the Derby database in network mode.

Note: The `startAM.bat(sh)` script has the Classpath defined for all the jars required by the Action Manager. There are two variables namely `CLASSPATH` and `DB_CLASSPATH`. The `DB_CLASSPATH` has the path separated list of JAR files required for achieving JDBC Connectivity with the database. When AMC is configured to use Oracle, MS SQL Server or DB2 the corresponding JDBC JAR files of these databases should be added to the `DB_CLASSPATH` variable.

AMC and Derby shutdown: The `stop_tdiamc.bat(sh)` script calls the `stopNetworkServer.bat(sh)` script. This ensures that the Derby Network server is stopped when AMC is shutdown.

Note: If Action Manager (AM) is running, this should be shut down first.

Action Manager remote startup: This section assumes that Action Manager and Derby are running on different computers.

1. Start Derby using the script `startNetworkServer.bat(sh)`.
2. Start Action Manager using the script `startAM.bat(sh)`.

The `startNetworkServer` script is used for starting the Derby database server in Network mode. The Derby server starts in Network mode on port 1528. The port selected is different from the default port for Derby.

Action Manager shutdown: The Action Manager is stopped using the `stopAM.bat(sh)` script located in the `TDI_install_dir/bin/amc` directory. This script uses the processID of the started AM to kill it. The processID is obtained by the `startAM` script and is stored in a file, which in turn is read by the `stopAM` script.

To stop the Derby database, type `stopNetworkServer`, which stops the Derby database server in Network mode. This should be done after AM is stopped, not before.

AMC Logs

The Administration and Monitoring Console logs are stored in the ISC log in the environment in which AMC runs:

- for ISC SE, the log file is created under `${LWI_HOME}/logs`;
- for ISC AE, the logs are logged in `${WAS_HOME}/profiles/${profileName}/logs/${serverInstance}/SystemOut.log`

The configuration of AMC logs can be done by modifying the `WEB-INF/classes/logging.properties` file. AMC logging follows the Java logging standard (`java.util.logging`).

You can view and delete AssemblyLine logs in the **AssemblyLine Logs** window. To reach **AssemblyLine Logs** on the **Monitor Status** window: (**Monitor Status** → **Solution View Details** → **View Logs**). In the **Solution View Details** window, select the AssemblyLine whose logs you want to view. In the **AssemblyLine Logs** window, select any logs you want to delete, and select **Delete**. You can delete one or multiple logs. To view a log, click its hyperlink.

The **Solution View Details** window also contains the **Action Manager Logs** table. You can select and delete logs from **Action Manager Logs**.

You can manage all of your logs in the **Log Management** window. You can specify criteria for displaying logs, and you can delete logs for all AssemblyLines or for a single AssemblyLine. You can select to delete all logs for AssemblyLine(s) specifying a date range, or you can delete the most recent logs, where you enter the number of most recent logs.

Backward Compatibility with previous versions of Tivoli Directory Integrator

When Administration and Monitoring Console for Tivoli Directory Integrator 7.0 is used to manage older versions of Tivoli Directory Integrator (TDI), some functionality is limited.

- The concept of Tivoli Directory Integrator Properties was changed in TDI 6.1. For this reason, AMC does not support editing/viewing/management of properties when working with a TDI 6.0 server. Also, the properties related triggers and actions are unavailable when working with a remote TDI 6.0 server.
- The **sendEventNotification** API for sending events remotely to a Tivoli Directory Integrator Server was introduced in Tivoli Directory Integrator 6.1. Therefore AMC does not support sending of Tivoli Directory Integrator events to remote TDI 6.0 servers. The Send Event Notification Action (in Action Manager) windows are not available when working with a Tivoli Directory Integrator 6.0 server.
- The Tombstone Manager Feature was introduced in Tivoli Directory Integrator 6.1. Hence viewing of tombstones, and seeing the last run time/stop time of an AL is not possible for remote TDI 6.0 servers. In addition, the "View Tombstones in chronological order" feature (and view last stop time/run time) was added in Tivoli Directory Integrator 7.0, by adding an API in Tombstone Manager. Hence, this feature is available only in Tivoli Directory Integrator 7.0.
- The remote config folder and the ability to view Configs in the remote config folder was a feature introduced in TDI 6.1. For this reason, in the "load-reload Configs" window, only those Configs that are already loaded are shown on the remote TDI 6.0 server. A Load operation is not available for TDI 6.0 servers. Only Stop and Re-load operations are available for TDI version 6.0.
- The AL Operation feature was introduced in TDI 6.1. Therefore starting an AL with AL Operation is not supported for TDI 6.0 servers.
- You are able to Add a Solution View (minus properties), select ALs to expose, select health AL, configure users, start or stop ALs, stop or reload Configs, create rules on ALs when working with a TDI 6.0 server.
- Custom Authentication (using ldap or JavaScript) was introduced in TDI 6.1. AMC's add server window supports username and password fields for this purpose. If a user attempts to pass username and password to a TDI 6.0 server, AMC is not able to connect to the remote Tivoli Directory Integrator server. There is no way for AMC to find out that the remote server is TDI 6.0, since it cannot connect with "incorrect" settings.
- Creation of Quick Solution View:
 - Publish Solution: is enabled only for TDI server version 6.1.1 or later, and if there is actually a published solution defined for the selected config instance.
 - Create Solution View with All Assembly lines exposed: this feature is available with TDI 6.0, 6.1 and 6.1.1 servers.
 - Creates a Solution View with all AssemblyLines from the config instance exposed, and all properties and no Health AL defined. This is similar to a quick start type of option. This option is disabled for TDI 6.0 servers (because TDI properties are not available in TDI 6.0 Servers)..
 - In the View Tombstone window only the 30 most recent tombstones entries are displayed.

AMC in the Integrated Solutions Console

The Integrated Solutions Console (ISC) is designed to offer a common console to organize administrative console functions using industry-standard technologies. Starting with TDI 7.0, integration of the AMC into the Integrated Solutions Console (ISC) comes with the following changes. The primary navigation links for AMC are:

- **Administration and Monitoring Console**
 - Servers
 - Solution Views
 - Monitor Status
 - Action Manager
- **Advanced**
 - Log Management
 - Console Properties
 - Preferred Solution Views
 - Property Stores

Console user authority

Using ISC **Console User Authority**, you can add and remove users to AMC. In the AMC for Tivoli Directory Integrator 7.0, the following are the roles:

Table 22. AMC roles

User Role	Description
administrator	Users with this role assigned are able to configure the roles other users are assigned to.
iscadmins	Users with this role assigned have the ability to control the settings of the ISC console itself.
TDI AMC Admin	This role is considered by the TDI AMC application deployed in the ISC console. Users with this role assigned are able to administer Servers, Solutions View roles, Manage Console Properties. (This was the superadmin role in the AMC prior to Tivoli Directory Integrator 7.0)
TDI AMC User	This role is considered by the TDI AMC application deployed in the ISC console. Users with this role assigned are able to use the provided by the TDI AMC Admin resources.

Within the *TDI AMC user* role, user privileges are assigned roles found in the Solution View:

- Admin
- Config Admin
- Execute
- Read

The roles control access to functions on the console. You can see only those functions for which you have roles assigned. For example, users with the TDI AMC Admin role automatically have administrator privileges over all Solution Views. Administrators can configure the properties required for the Web administration tool (modifying properties related to `amc.properties` file, available from Console Properties in the left navigation pane). A user with the TDI AMC User role in ISC is the same as the current non-admin AMC user. TDI AMC Users cannot access any administrative windows such as TDI Servers and Console Properties.

ISC features the AMC Admin Group or `iscadmins`. A user in the `iscadmins` group has the same privileges as the administrator.

Administrator and the iscadmins group

The administrator, defined as the user who has installed the application, can manage AMC users. The administrator can add or remove users from the local OS Registry to the AMC application and assign or edit roles. Tivoli Directory Integrator 7.0 offers a new **TDI AMC Admin** group. The superadmin role does not exist after TDI version 6.1.1.

Action Manager

The Action Manager is a standalone Java application that allows you to monitor multiple TDI Servers and AssemblyLine execution using user-defined rules, triggering conditions and actions defined in AMC. The Administration and Monitoring Console (AMC) has an Action Manager window that allows users to configure various Action Manager rules.

A rule is a combination of a trigger type and a set of associated actions. A rule specifies that when a triggering condition is detected, then the associated set of actions must be executed. The various trigger types available in AMC are described below:

Table 23. Action Manager triggers

Trigger Type	Trigger Details	User Input for trigger
No trigger	A rule with this triggering type has no triggering condition, and as a result never gets triggered by itself. The only way this rule can be executed is if some other rule executes this rule	No details required
On AssemblyLine start	A rule with this triggering type gets triggered when the Action Manager receives an AL start event for this particular AL.	"AssemblyLine name"
On AssemblyLine termination	A rule with this triggering type gets triggered when the Action Manager receives an AL termination event for this particular AL.	"AssemblyLine name"
On config load	A rule with this triggering type is triggered when a Config is loaded.	"Name of Config to load"
On config unload	A rule with this triggering type is triggered when a Config is unloaded. The Config must be loaded to be unloaded.	"Name of Config to unload"

Table 23. Action Manager triggers (continued)

Trigger Type	Trigger Details	User Input for trigger
On query AssemblyLine result	<p>Note: A rule with this triggering type should not be used with a short-running AL. This is because Action Manager stores the handle of the AL object on receiving the Start AssemblyLine event. Later on, receiving the Stop AssemblyLine event, Action Manager uses this handle to query the final work entry attributes. If the AL terminates before Action Manager can store the handle, then Action Manager is not able to query the work attributes. Usually an execution time of 10 seconds is sufficient (this can be achieved by putting a <code>system.sleep(10)</code> before the AL terminates, for example in the epilog hook).</p> <p>When running On query AL result, Action Manager polls the AL continuously for the specified polling interval. The trigger first checks for the attribute value, starting the AL after the specified polling interval. Next, the trigger checks the AL result entry.</p> <p>On query AL result is a rule that is triggered when the last "work" entry of the specified AL contains the specified "Attribute" matching the given "condition" and "value". This condition is checked only when the ActionManager receives a Stop AssemblyLine event. The user can specify a time interval. The specified AL run periodically depending upon the time interval specified. A rule with this triggering type is triggered when the last "work" entry of the specified AL, contains the specified "Attribute" matching the given "condition" and "value". This condition is checked only when the Action Manager receives a Stop AssemblyLine event.</p> <p>To configure the Query on AL result trigger, enter values for the following fields:</p> <ul style="list-style-type: none"> • AssemblyLine name • Attribute • Condition • Value • Polling Interval • Polling Unit 	"AssemblyLine name", "Attribute", "Condition", "Value", "Polling Interval," and "Polling Unit."
On server API failure	A rule with this triggering type is triggered when the Action Manager is unable to connect to the remote server using the Server API. You can configure different polling time intervals for each Assembly Line depending upon AL execution time.	"Polling Interval" and "Polling Unit."
On received Event	<p>A rule with this triggering type is triggered when the Action Manager receives an event which satisfies the criteria mentioned.</p> <p>Note: If any of the criteria are to be ignored, just leave it blank.</p>	"Event type", "Event Source", "Event Data". Event Data is optional. Event type or source – one of them must be specified.

Table 23. Action Manager triggers (continued)

Trigger Type	Trigger Details	User Input for trigger
On Property	<p>A rule with this triggering type is triggered when the specified property meets the specified condition. The Action Manager periodically checks for this property. You can configure a polling interval and polling units when configuring this trigger.</p> <p>Note: This rule gets triggered only once, and gets reset back to ready state only when Action Manager detects that this property does not meet the specified criteria any longer. This is done so that the rule does not repeatedly get triggered for a single occurrence of the triggering condition.</p>	"Polling Interval," "Polling Unit"."Property name", "Condition", and "Value".
On local variable	<p>A rule with this triggering type is triggered when the specified variables meet the specified condition. The Action Manager periodically checks for this property .</p> <p>Note: This rule gets triggered only once, and gets reset back to the ready state only when Action Manager detects that this variables does not meet the specified criteria any longer. This is done so that the rule does not repeatedly get triggered for a single occurrence of the triggering condition.</p>	"Local Variable" ," Condition", "Value".
Inspect AssemblyLine exit code	<p>A rule with this triggering type is triggered when an AssemblyLine terminates with an error.</p> <p>Inspect AL Exit Code also searches for an error object string for every abnormal AL termination. Under Configure Trigger, if the trigger is Inspect AL Exit Code, you can enable Inspect Error Object. In the Value field, type the string you want for Error Object. Note that if the Value field is empty, then the rule triggers for every abnormal termination of an AL. If Inspect Error Object is not selected, the trigger waits for the AL to terminate and inspects the exit code for an attribute value (entered by the user). Type values for both the Attribute Name and Value.</p> <p>In the Inspect AL Exit code trigger, the Action Manager no longer starts the AL, and there is no polling. The trigger only checks the AL result one time after the AL runs.</p>	"AssemblyLine name"; if "Inspect Error Object" is enabled, you only need supply "Value." If "Inspect Error Object" is disabled, values for "Attribute," "Condition," and "Value" are needed.
Time since last execution	<p>A rule with this triggering type get triggered when the Action Manager detects that the specified assembly line has not run for the specified period. Note: This rule is triggered only once. After that Action Manager wait for receiving a Start AssemblyLine event before resetting the Rule back to Ready mode. This is done so that the rule does not repeatedly get triggered for a single occurrence of the triggering condition.</p>	"AssemblyLine name", "Not Run Since" and "Unit".
Timer	<p>A rule with this triggering type is triggered continuously within an interval defined by a number of units and the measure of seconds, minutes, hours or days.</p>	"Interval" and "Unit."

When a rule gets triggered, the Actions associated with the rule are executed by the Action Manager sequentially. The following are the various types of Actions that are available in AMC:

Table 24. Action Manager actions

Action	Action Details	User Input for action
Start AssemblyLine	This action starts the specified AL of the specified config file on the specified TDI server. The Config field should mention the complete path of the configuration on the remote server. The Config Password field is optional and is required only if the remote config is password protected.	"AssemblyLine", "Of Configuration", "On Server", "Config Password".
Stop AssemblyLine	This action stops the specified AL of the specified configuration on the specified TDI Server. The Config field should mention the complete path of the configuration on the remote server.	"AssemblyLine", "Of Configuration", "On Server".
Enable/Disable Rule	This action Enable or Disable the chosen rule.	"RuleName" "State"
Execute Rule	This action cause the execution of the specified rule, which in-turn imply execution of all the actions specified in that particular rule.	"RuleName"
Notify Event	This action cause the Action Manager to emit an event with the specified details to the Server associated with the current Solution View. See the Session.sendCustomNotification() API for details.	"Event type", "Source", "Data".
Modify Property	This action cause the Action Manager to modify the selected property based on the specified operation.	"Property", "Operation", "Value".
Copy Property Value	This action cause the Action Manager to copy the value of the Source property to the Destination property.	"From Property", "To Property".
Write to Log	This action causes a log of the specified Severity/Message/Description to be logged into the Action Manager logs and the AMC database. The same log is shown when the user goes to the Monitor Status -> Solution View Details -> Action Manager Results table. It is advised to always have at least one Log action (containing descriptive text) in every rule.	"Severity", "Message", "Description".
Send Email	This action causes an email to be sent to the recipient you specify. You supply the content of the email. Along with the content, the Action Manager provides other details before sending the mail. In the content input area as well as in the subject line, you can specify the variable %EVENT_DATA% value. Specifying %EventData% inserts the actual value of the Eventdata variable when the mail is sent. %Action_Error% can also similarly be substituted here. If Attach Action Manager Log is enabled, the Action Manager logs (as specified in the am_logging.properties file) are sent as an email attachment.	"To", "From", "Subject", "Attach Action Manager Log" (Selected/Not Selected), "Content".
Modify local variables	This action causes the action manager to increment, decrement or set the value of the specified variable to the specified value.	"Variable", "Operation", "Value".

Table 24. Action Manager actions (continued)

Action	Action Details	User Input for action
Execute command	This action causes the specified command to execute on the target computer. The command can be any generic command or a TDI specific command.	"Target Machine", "Port", "Username", "Password", "Keystore", "Keystore Password", "Protocol", and "Command".

Rules that are configured for Solution Views in AMC, are stored in AMC's Derby Database. When the Action Manager is run, it connects to the AMC database in network mode, reads the Action Manager-related tables, and creates threads in memory for every rule specified. Each of these threads listens/polls for its respective triggering conditions. The moment any thread detects the occurrence of its respective triggering condition, it queries the database for the set of actions associated with the rule, and executes them sequentially.

The Action Manager runs the following threads in addition to the rule threads that are listening for trigger conditions:

1. HealthAssemblyLine – The Health AssemblyLine thread periodically triggers the Health ALs for querying the status of the solutions, and logging the status back into the AMC database. The health AL must store the status in the "healthAL.result" and "healthAL.status" attributes of their final work entry.
2. ServerStatusListener - The ServerStatusListener thread is created for every server registered with AMC. This thread checks for the server accessibility. If the server has become inaccessible, all rules threads created for the server are terminated (except for those with triggering type 'On Server API failure'). Similarly if the server becomes accessible, rule threads are created for any rules associated with this server.
3. ConfigLoadReloadListener – The ConfigLoadReloadListener thread is created for every running server registered with AMC. It is registered to the remote server for any config load unload events. Rule threads are appropriately terminated, created or refreshed depending on the config event.
4. ServerModificationListener – The ServerModificationListener thread checks for any updates to the set of servers registered in AMC. Depending on the type of change (added, removed, and so on.) rule threads are terminated, created or refreshed.
5. DatabaseModificationListener - This database listener thread continuously monitors addition, modification or deletion of rules. Whenever any changes in the rules are detected, the Action Manager threads are added/recreated appropriately at runtime.

The Action Manager also updates the AMC database with its run details. Whenever an Action Manager rule is triggered, Action Manager logs an entry into the AMC database, registering the rule name that got triggered, and the triggering time. Also, if any Log action is configured for the rule, then that also gets logged into the AMC database. These database entries are used to show the appropriate status in Monitor windows of AMC.

Enabling Action Manager

The Action Manager is installed in the *TDI_install_dir/bin/amc/Action Manager* folder. It contains the following files:

- am_logging.properties - This file controls Action Manager logging properties. Just like AMC, it also follows the java.util.logging logging standard.
- am_config.properties - This is the configuration file for the Action Manager.
- testadmin.jks - This is the ActionManager's truststore and keystore file.

Note: This is a sample truststore and keystore file; for added security, you should generate your own.

The Action Manager connects to AMC's Derby database using the Network Mode driver.

The following properties (in `am_config.properties`) must point to the Administration and Monitoring Console's database:

```
com.ibm.di.amc.am.jdbc.database=jdbc:derby://localhost:1528/tdiamcdb;create=false
com.ibm.di.amc.am.jdbc.driver=org.apache.derby.jdbc.ClientDriver
com.ibm.di.amc.am.jdbc.urlprefix=jdbc:derby:
com.ibm.di.amc.am.jdbc.user=APP
com.ibm.di.amc.am.jdbc.password=APP
com.ibm.di.amc.am.jdbc.start.mode=automatic
com.ibm.di.amc.am.jdbc.sysibm=true
com.ibm.di.amc.am.jdbc.networkserver.host=localhost
com.ibm.di.amc.am.jdbc.networkserver.port=1528
```

Note: Both AMC and AM support alternative databases, like MS SQL, Oracle and so forth. In order for AMC and AM to connect to one of those alternative databases, the configuration statements in `amc.properties` and `am_config.properties` will look very different.

When the Action Manager is started, it attempts to connect to AMC's database. If it fails in performing any initial setup tasks, it exit with an exception message. Check the `am_config.properties` file to ensure it points to the correct database. If the database settings appear to be correct, then ensure that the database that Action Manager is attempting to connect to is running in network mode and that AMC can connect to the same database. You may use the `startNetworkServer.bat` (sh) to start the Derby DB in network mode.

Further configuration of run-time rules, triggers and actions is described under “Action Manager” on page 211.

Action Manager status in real time

When you login to AMC, a one line Action Manager status displays in the **Welcome to AMC** panel. The Welcome to AMC panel displays the Action Manager status in a link, for example, "Action Manager is running" or "Action Manager is not running." To launch the **Action Manager Status** window, click the hyperlink. The Action Manager Status window displays Action Manager Status in real time, as well as thread details and trigger details. This window displays status information in real time. This window shows:

- Action Manager Status, for example the Boot Time
- Action Manager Thread Details
- Action Manager Trigger Details

The AMC directly queries the Action Manager using the APIs exposed by Action Manager. If AMC cannot establish a session with the Action Manager, the AMC concludes that the Action Manager is not available because the Action Manager has stopped. In addition to Action Manager status, AMC displays details of thread information and trigger details.

Action Manager creates a number of threads. Some Action Manager threads monitor the essential functionality of the Action Manager such as the Database Modification Listener and the `ServerStatusListenerThread`. Moreover, from these threads the Action Manager creates threads for each of the trigger rules that is configured in AMC. With the Remote Method Invocation (RMI) Layer, AMC can query the status of the various trigger-related threads. Using the RMI based query, AMC knows the state of these threads, thread priority, and so on. AMC can also query the triggers that have been executed over a period of time.

Two new properties belong to the **Display real time Action Manager Status** requirement. The properties that allow AMC to display the Action Manager status in real time are `am.api.host` and `am.api.port`. Action Manager status used an RMI layer around the Action Manager that exposes an API to be used by AMC for querying the Action Manager for its status.

AMC force trigger for a given rule

AMC allows users to force a trigger for a specific rule. Forcing a trigger gives the user an idea of what the Action Manager does when the rule is triggered. When you select **Disable Rule**, the selected rule is disabled. When you select **Force Trigger**, actions configured for the selected rule are executed.

Action Manager can execute a set of actions configured for a particular trigger (rule) explicitly. The AMC user does not have to wait for the triggering condition of a rule to be satisfied before the configured actions are carried out. Users can define actions that are to be executed so that they can test those actions. Users can execute all of the supported actions using **Force Trigger**. However, the **Revert** action is effective for only some (a subset of) the supported actions.

- Modify Property
- Copy Property
- Write to Log
- Enable Disable Rule

AMC and Action Manager security

Introduction

The Administration and Monitoring Console (AMC) is a web-based application for monitoring and managing remote Tivoli Directory Integrator solutions. The following features of AMC have been improved:

- Encryption or concealment of passwords that are stored in the `amc.properties` file
- Use of stash files to store keystore passwords
- Enablement of the BUILTIN authentication scheme in the Derby database

AMC uses the Remote Server API to communicate with Tivoli Directory Integrator. For this reason, all the security restrictions and configuration settings that are applicable to Tivoli Directory Integrator Remote Server API clients (as mentioned in previous sections) are valid for AMC too.

Action Manager is installed with AMC. Action Manager configures itself and behaves based on rules set in the AMC database by AMC users. To monitor remote AssemblyLines and to take action based on configured rules, Action Manager, just like AMC, uses the Tivoli Directory Integrator Remote Server API to communicate with Tivoli Directory Integrator servers.

Note: Communication between AMC and AM using RMI is not protected in any way.

AMC and SSL

Multiple Tivoli Directory Integrator servers can be registered with AMC. Each TDI server may be configured differently; one Tivoli Directory Integrator server could be running with SSL off, one with SSL on, one with Custom Authentication on and SSL on – and various other combinations. AMC can be used to connect and administer any of these servers simultaneously. As mentioned earlier, to configure Tivoli Directory Integrator to run in SSL mode the `api.remote.ssl.on` property should be set to `true` in `global.properties` (or `solution.properties`).

As AMC is a web application running inside a Web Container it automatically inherits some properties and security restrictions from the Web Container. For instance, if the Web Container has an SSL keystore or SSL truststore configured, then that would be automatically applicable to AMC. But AMC can also override that – and specify its own keystore and truststore.

For being able to communicate with Tivoli Directory Integrator Remote Server API running on SSL, AMC must have a keystore configured which contains the certificate that is trusted by the Tivoli Directory Integrator remote Server API (that is, it must be present in Tivoli Directory Integrator's truststore's

trusted certificates section) and AMC must have a truststore configured which contains the certificate that is sent by the Tivoli Directory Integrator remote Server API. In other words – the certificate that is present in Tivoli Directory Integrator server's keystore must be present in AMC's truststore and the certificate that is present in Tivoli Directory Integrator truststore must be present in AMC's keystore.

For example, the default installation of Tivoli Directory Integrator is shipped with certain stores (.jks files). When you run Tivoli Directory Integrator in SSL mode, then to connect to AMC its keystore and truststore must both be set to the same value: *TDI_install_dir/serverapi/testadmin.jks* and the password being "administrator". Since testadmin.jks contains both trusted certificates and signer certificates, a connection gets established. It is recommended that you set up your own SSL keystores and truststores.

In AMC, the path of the truststore and keystore can be set by logging into AMC as "TDI AMC Admin" (Console Administrator) and navigating to the following window: **Advanced -> Console Properties -> SSL settings**. The settings for truststore and keystore are written to **amc.properties** file inside the tdiamc folder in Web Container. You can alternatively choose to edit the **amc.properties** file directly. With Tivoli Directory Integrator 7.0, AMC can be deployed in ISC Standard Edition (SE) or in ISC Advanced Edition (AE). Depending on the ISC runtime, the location of the testadmin.jks file varies. For example, if AMC is deployed in ISC SE, then the location will be *ISC_RUNTIME_INSTALL_DIR/runtime/isc/eclipse/plugins/AMC_7.0.0*. On the other hand, if AMC is deployed in ISC AE, then the location is *ISC_RUNTIME_INSTALL_DIR/systemApps/isclite.ear/tdiamc.war*. The keystore and truststore password are set to "administrator" by default. To establish an SSL based connection with a remote Tivoli Directory Integrator server, you must start the server in "SSL enabled" mode, and for a Non SSL based connection, start the server in "SSL disabled" mode.

Attention: Default SSL settings are provided. However, using the default certificates does not increase the security more than just using a plain connection, so after installation, you should replace the default SSL certificates and update the keystores and truststores accordingly in order to increase security.

For each Tivoli Directory Integrator server running over SSL that you wish to register with AMC, you must import the necessary certificate into AMC's truststore and the necessary AMC's key certificate into Tivoli Directory Integrator's truststore. The idea here is that AMC must trust TDI and TDI must trust AMC to be able to make a secured two-way SSL connection.

Since AMC runs inside a Web Container, the URL for AMC is *http://hostname:port/ibm/console*.

Action Manager monitors running Configs and AssemblyLines on remote TDI Servers based on rules configured in AMC. Action Manager ships with the keystore and truststore required to connect to a remote TDI server. The SSL properties are defined in the *am_config.properties*. See details on how to configure AMC for SSL in previous sections - the same is applicable for Action Manager.

AMC and remote TDI server

AMC can connect to multiple Tivoli Directory Integrator Servers remotely. Each Server can be configured in one of the following ways:

- Non SSL
- SSL
- Custom Authentication with Non-SSL
- Custom Authentication with SSL

This section looks at each of these cases in detail.

When a remote Tivoli Directory Integrator server is configured for non SSL (that is, *api.remote.ssl.on=false*) then the keystore or truststores of AMC do not come into play, even if correctly configured – since no SSL connection is being attempted. In this case the AMC Server's

computer IP address must be registered with the TDI server. This is done by editing the `global.properties` (or `solution.properties`) file. The property to update is: **api.remote.nonssl.hosts**. Once the AMC computer's IP address is entered in the `global.properties` file of the remote TDI server, AMC is able to connect to that particular server. It is a way of saying – I trust remote server connections (AMC connections) from only those computers whose IP addresses I have mentioned in my **api.remote.nonssl.hosts** property.

Note: If the Tivoli Directory Integrator server is running on the same computer as AMC, then editing this property is not required.

When a remote Tivoli Directory Integrator server is configured for SSL (that is, **api.remote.ssl.on=true**), then the SSL keystore and truststore for AMC must be setup appropriately.

For details on this, see the previous section on AMC and SSL. In addition to being configured for SSL or Non-SSL, a remote TDI server may also require Custom Authentication – in which a username and password must be passed while making a connection to the remote TDI server. The remote TDI server validates this user name and password against some third party repository like LDAP, file, database, script, and so on and then make a decision on whether to allow the Server API client to make a connection or not. In such cases, while registering a server with AMC (**Servers -> Modify Server**) in the Authentication mode window – select **LDAP or Custom Authentication** and enter the Username and Password that AMC must pass every time it attempts to connect to the specified remote TDI Server.

Note: If the Username or Password (in case of custom authentication) or SSL keystores or truststores (in case of SSL) are not set up correctly, then AMC is unable to connect to the remote TDI Server and show that server as "Stopped" or "Not running."

AMC and role management

Every user (or group) in AMC can be assigned one of the following roles in AMC for a particular Solution View. This role assignment can be done in the **Solution Views** window by selecting a particular Solution View and by clicking **Configure Access Control Lists (ACLs)**. The **Configure ACLs** window displays. Select the Name of the user you want to configure and click **Configure Users** on the toolbar. The **Configure Users** window displays. Select the **User ID** and select one of the available roles:

- Read
- Execute
- Admin
- Config Admin

Note: You must reload Solution Views created using the Auto Update option. Use the **Refresh Solution View** in the **Solution Views** window. For Solution Views marked for auto update, you must reload the config file and refresh the Solution View by clicking the **Refresh Solution View**. If a user fails to refresh a Solution View created using the **Simple** option and flagged for auto update, the Solution View may cause inconsistencies in the AMC database. Inconsistencies in Solution Views that are not updated could result in incorrect behavior by the Action Manager.

These roles are in increasing order of privilege – indicating that Config Admin is the highest privilege and Read is the lowest. Any functionality that is available to a user with "Read" role for a Solution View, definitely is available to a user with "Execute" privilege on that Solution View. Any functionality that is available to a user with "Execute" privilege on a Solution View, is available to a user with "Admin" privilege, and so on.

The following is the meaning of these roles

Read This means that this user can only read the "details" of this Solution View – such as what are the ALs inside this view, what are properties inside this view, what is the status of these ALs, and so on. This user cannot modify, start, stop, or change any detail of this Solution View.

Execute

This is essentially a Read user with one extra privilege – the ability to Start and Stop AssemblyLines.

Admin

This user can administer the Solution View, without being able to modify the Solution View itself. This user can do everything that the "Execute" privilege user can do, and additionally he can modify properties, delete logs, configure rules, and so on, for this Solution View.

Config_Admin

This user can virtually do anything to the Solution View – including modifying the view itself, modifying the permissions of other users on this view, and so on. This is the highest privilege that can be given to a user for a particular Solution View.

The above roles can be assigned to any Group too. Therefore, if a user "test" and "tdi" are part of the "DBAdmin" group, and the "DBAdmin" group is given "ConfigAdmin" privilege over a Solution View "SynchDatabase", then both "test" and "tdi" automatically get ConfigAdmin privilege over the "SynchDatabase" Solution View.

Notes:

1. If the "test" user is explicitly given "read" privilege for the same Solution View, then "Read" get precedence over the privilege he gets from being part of the "DBAdmin" group. This is done so that "specific" role assignment gets priority over role assignment from groups. This allows people to restrict or give higher access to individuals – without worrying about inherited access from being part of some groups.
2. If the "test" user is part of two Groups – where Group1 has "read" access and Group2 has "admin" access over the same Solution View – then in this case the test user get the higher of the two privilege – in this case being "admin", unless a specific role is already assigned to "test" for the same Solution View – in which case the specific role assigned to "test" is given precedence [point 1 above].

AMC and passwords

Any password field that is stored in amc.properties file, such as LDAP Bind password, keystore password, and so on, are all encrypted before being written to amc.properties file. Also, AMC never displays any Password fields or protected fields on console. All such fields are masked out.

AMC and encrypted configs

AMC allows users to load and connect to password protected configs. On the Load Reload window of AMC, a password text box has been provided – where the users must enter the password of the config they are attempting to start before clicking **Start**. Similarly, in the Action Manager Screen – for the Start AssemblyLine action, a password field has been provided where the user can enter the password of the config. Action Manager passes this password while attempting to start the Config.

Note: AMC cannot detect that the remote config being started is a password protected config. For this reason, if the password is not specified or incorrectly specified, then the user just see an error message saying – "Unable to start the config". The user can see the TDI Server logs where an exact message is provided.

Administration and Monitoring Console User Interface

Log in and logout of the console

Open a Web browser and type the following address:

`http://hostname:port/ibm/console`

Where *port* is the port where your Web server is running. The default port number is 1528.

The IBM Tivoli Directory Integrator Administration and Monitoring Console login page window is displayed.


Logging on to the console as the console administrator:

The console administrator is a user who can:

- Configure the properties required for the AMC
- Set the authentication mechanism used for AMC logins
- Add new users and configure users' roles

The bundled Integrated Solutions Console is configured to use the operating system repository for authentication. The user who has installed the Administration and Monitoring Console receives IBM Tivoli Directory Integrator Administration and Monitoring Console Admin rights.

To log in to the Integrated Solutions Console, type your user name and password in the boxes provided in the login window and click the **Log in** button.



The **Logout** button is in the upper right hand corner of the console, next to **Help**. When you click **Logout**, you are returned to the Log in page.

AMC Console Layout

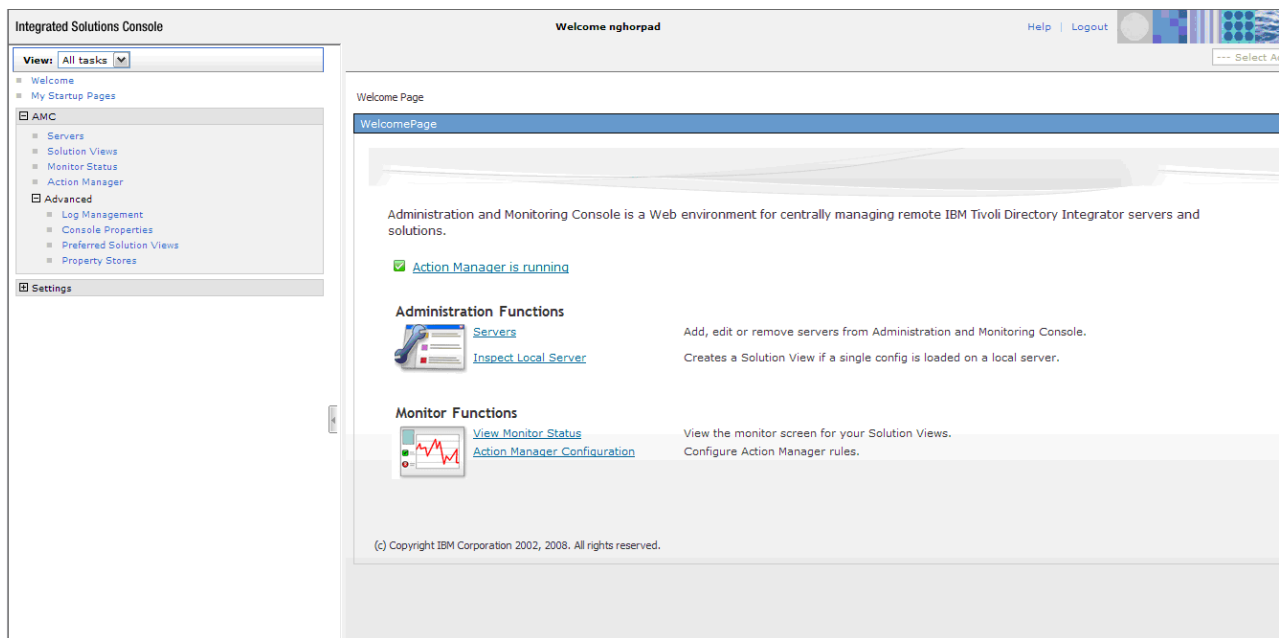
The IBM Tivoli Directory Integrator Administration and Monitoring Console includes the following components:

Navigation Area

The Navigation area provides a tree view that allows users to navigate through the tasks available to the user in the console. You can open and close folders in the navigation area and select tasks (non-folders) to launch in the Work Area of the console framework.

Work Area

The Work Area contains the necessary information and input fields to complete the task you are currently working on.



Logging off the console

To log off of the console, click **Logout** in the navigation area.

Using AMC tables

The IBM Tivoli Directory Integrator Administration and Monitoring Console displays certain information, such as lists of attributes and entries, in tables. Tables contain several utilities that allow you to search for, organize and perform actions on these table items.

Tivoli Directory Integrator Administration and Monitoring Console tables provide icons to help you organize and find information in the table. Some icons appear on some tables and not on others, depending on the current task. The following is a comprehensive list of the icons you might encounter:

- Click the **Show Filter Row** icon to display filter rows for every column in the table. See “Filtering” on page 202 for more information about filtering.
- Click the **Hide Filter Row** icon to hide filter rows for every column in the table. See “Filtering” on page 202 for more information.
- Click the **Clear all filters** icon clear all filters set for the table. See “Filtering” on page 202 for more information.
- Click the **Edit sort** icon to sort the information in the table. See “Sorting” on page 201 for more information.
- Click the **Clear all sorts** icon to clear all sorts set for the table. See “Sorting” on page 201 for more information.
- Click the **Collapse table** icon to hide the table data.
- Click the **Expand table** icon to display the table data.
- Click the **Select all** icon to select all items in the table.
- Click the **Deselect all** icon to deselect all selected items in the table.
- Click the **Export** icon to export the table data.

Select action drop-down menu

The **Select action** drop-down menu contains a comprehensive list of all available actions for a selected table. For example, instead of using the icons to display and hide sorts and filters, you can use the **Select action** drop-down menu. You can also use the **Select action** drop-down menu to perform operations on the table contents; for example, on the **Manage attributes** window, actions such as **View**, **Add**, **Edit**, **Copy** and **Delete** appear not only as buttons on the toolbar, but also in the **Select action** drop-down menu. If the table supports it, you can also display or hide the **Show find** toolbar using the **Select action** drop-down menu. See Finding for more information on finding table items.

To perform an action using the Select action menu:

1. If necessary, select an item from the table.
2. Click the **Select action** drop-down menu.
3. Select the action you want to perform; for example **Shutdown server**.
4. Click **Go**.

Paging

To view different table pages, use the navigation controls at the bottom of the table. You can enter a specific page number into the navigation field and click **Go** to display a certain page. You can also use the **Next** and **Previous** arrows to move from page to page.

Sorting

To change the way items in a table are sorted:

1. Do one of the following:
 - Click the **Edit sort** icon on the table.
 - Click the **Select action** drop-down menu, select **Edit sort** and click **Go**.

A sorting drop-down menu appears for every column in the table.

2. From the first sort drop-down menu, select the column on which you'd like to sort. Do the same for any of the other sortable columns on which you'd like to sort.
3. Select whether to sort in ascending or descending order by selecting **Ascending/ Descending** from the drop-down menu. Ascending is the default sort order. You can also sort using column headers. On every column is a small arrow. An arrow pointing up means that column is sorted in ascending order. An arrow pointing down means that column is sorted in descending order. To change the sort order, simply click on the column header.
4. When you are ready to sort, click **Sort**.

To clear all the sorts, click the **Clear all sorts** icon.

Finding

To find a specific item or items in a table:

Note: The Show find toolbar option is available on some tables and not on others, depending on the current task.

1. Select **Show find toolbar** from the **Select action** drop-down menu and click **Go**.
2. Enter your search criteria in the **Search for** field.
3. If desired, select a condition upon which to search from the **Conditions** drop-down menu. The options for this menu are:
 - **Contains**
 - **Starts with**
 - **Ends with**
 - **Exact match**
4. Select the column upon which you want to base the search from the **Column** drop-down menu.

5. Select whether to display results in descending or ascending order from the **Direction** drop-down menu. Select **Down** to display results in descending order. Select **Up** to display results in ascending order.
6. Select **Match case** if you want search results to match the upper and lower case criteria in the **Search for** field.
7. When you have entered the desired criteria, click **Find** to search for the attributes.

Filtering

To filter items in a table, do the following:

1. Do one of the following:
 - Click the **Show filter** row icon. Click the **Select action** drop-down menu, select **Show filter** row and click **Go**.
2. Filter buttons appear above each column. Click **Filter** above the column on which you want to filter.
3. Select one of the following conditions from the **Conditions** drop-down menu:
 - Contains
 - Starts with
 - Ends with
4. Enter the text you want to filter on in the field; for example, if you selected **Starts with**, you might enter **C**.
5. If you want to match case (upper case text or lower case text) select **Match case**.
6. When you are ready to filter the attributes, click **OK**.
7. Repeat the above steps 2-6 for every column on which you want to filter.

To clear all the filters, click the **Clear all filters** icon.

To hide the filter rows, click the **Show filter** icon again.

Servers

This window allows you to view the registered server. Additionally, the console administrator can add, edit, delete and shut down IBM Tivoli Directory Integrator servers from this window, as well as launch the Config Files window.

When AMC is started , it automatically has a local Tivoli Directory Integrator server, registered on port 1099 . Therefore, in the **Servers** window, one entry in under LOCAL SERVER is already present with its state depicted as running or unavailable depending on its status.

To load or reload a config, select **Servers** and click **Config Files** in the toolbar of the Servers window. The **Config Files** window appears.

You can choose the operations you want to perform from the tool bar at the top of the table or using the Select action drop-down menu, such as:

Add Click Add on the toolbar.

Delete Select the radio button next to the server you want to delete and click **Delete** on the toolbar.

Modify

Select the server for which you want to modify information and click **Modify** on the toolbar.

Config Files

Select the server for which you want to list configuration files. When you click the **View Config Files** link in the Solution Views window, it launches the Config Files window. Each configuration file is labelled as loaded or not loaded. The toolbar provides a variety of load, unload, and reload options.

Shutdown server

Select the server you want to shut down and click the Shutdown Server on the toolbar.

Add a server

This window allows you to add an IBM Tivoli Directory Integrator server to the Administration and Monitoring Console (AMC). Once you have added a IBM Tivoli Directory Integrator server to the AMC, you can then use features on other AMC windows to add Solution Views to the TDI server and to create and define views for the Solution Views associated with the IBM Tivoli Directory Integrator server.

To add a new TDI server:

1. Enter a name for the IBM Tivoli Directory Integrator server in the **Name** field.
2. Enter the host name or IP address of the computer on which the IBM Tivoli Directory Integrator is running in the **Hostname** field.
3. Enter the port number on which the IBM Tivoli Directory Integrator server is configured to run.
4. Select the desired authentication mode. If you selected the Custom or LDAP authentication authentication method, enter the username and password to be used for authentication.
5. Click **OK**.

Modify a server

This window allows you to edit the information for an existing IBM Tivoli Directory Integrator server. To edit an existing server:

1. Look at the displayed Server ID. If you want to change the Server ID, click **Change Server ID**.
2. Type the Name for the server.
3. Enter the host name or IP address of the computer on which the IBM Tivoli Directory Integrator server is running in the **Hostname** field.
4. Enter the port number on which the IBM Tivoli Directory Integrator server is configured to run.
5. Select the desired authentication mode. If you selected the Custom or LDAP authentication authentication method, enter the username and password to be used for authentication.
6. Click **Cancel** to exit the window without making any changes, or click **OK** to save the changes.
7. Click **Test Connection** to see whether the connection to the server succeeds or not based on the current settings.

Console Properties

From these windows, you can set General, SSL and Miscellaneous console settings.

General

This window allows you to set general properties such as refresh rates and session timeouts for the Administration and Monitoring Console. From this window you can:

- Set the monitor window refresh rate in minutes
- Set the frequency (in days) at which Action Manager logs are rotated.
- Set the Administration and Monitoring Console session timeout in minutes.

Note: The session time-out parameter is set when a user logs into AMC. Therefore, this parameter only takes effect from next user login.

LDAP

The Administration and Monitoring Console provides a way for you to authenticate users to an LDAP backend. If LDAP authentication is enabled, you must configure the properties of the LDAP server to which you want to authenticate.

SSL

This window allows you to set up the console so that it can communicate with other directory servers using the Secure Sockets Layer (SSL) encryption, if necessary.

From Tivoli Directory Integrator 7.0 the keystore location depends on the ISC runtime in which AMC is deployed. For example, if AMC is deployed in ISC SE, the default keystore value is `ISC_RUNTIME_INSTALL_DIR/runtime/isc/eclipse/plugin-ins/AMC_7.0.0`. If the AMC is deployed in ISC AE, the keystore value is `ISC_RUNTIME_INSTALL_DIR/systemApps/isclite.ear/tdiamc.war`.

JDBC Properties

JDBC properties are used to define the connections settings to the Derby database, or to other databases compatible with the Administration and Monitoring Console, such as Oracle and MS-SQL Server. The AMC database stores AMC configuration information, connection details, and Action Manager rules and results.

The IBM Tivoli Directory Integrator AMC supports alternative databases in addition to Derby. AMC bundles the Derby database, or another database that you select. AMC communicates with its database using the Java Database Connectivity (JDBC) protocol. JDBC is a generic protocol and can be easily extended to other databases. AMC support for alternate databases enables you to have AMC installed and communicating to an existing database. The database stores Action Manager logs, results, and so forth. The Integrated Solution Console **Advanced -> Console Properties** section groups the **JDBC properties** to Derby or to another database. In the case of Derby, you can configure the database to run in both embedded as well as network mode. The default database is Derby and the default mode is network mode.

From this window you can:

- Select a database from the **Database Type** field, options are Derby, MS SQL Server, Oracle and DB2.
- Type the a value for the JDBC URL in the **JDBC URL** field.
- Type the user name for the database in the **Username** field.
- Type the password for the database in the **Password** field.
- Type the JDBC driver name in the **JDBC Driver** field.

As for the JDBC URL and JDBC Driver parameters, the following table provides some guidance:

Table 25. Driver parameters

Database	JDBC URL	JDBC Driver	Driver .jar file
Derby	<code>jdbc:derby://host:port/database [;create=true create=false]</code>	<code>org.apache.derby.jdbc.ClientDriver</code>	<code>derby.jar</code>
MS SQL Server (2005)	<code>jdbc:sqlserver://host:port; databasename=database</code>	<code>com.microsoft.sqlserver.jdbc. SQLServerDriver</code>	<code>sqljdbc.jar</code>
Oracle	<code>jdbc:oracle:thin:@host:port:database</code>	<code>oracle.jdbc.driver.OracleDriver</code>	<code>ojdbc14.jar</code>
DB2	<code>jdbc:db2://host:port/database</code>	<code>com.ibm.db2.jcc.DB2Driver</code>	<code>db2jcc.jar</code>

Notes:

1. Depending on the database selected the corresponding driver .jar file must be copied to `TDI_install_dir\lwi\libs`.
2. Configuration of the Action Manager is also needed in order to specify the new database from where it will work. The same .jar file must be added to `TDI_install_dir/bin/amc/ActionManager/jars` and adjustments must be made to the `am_config.properties` file.
3. If you decide not to use Derby, but one of the alternatives, keep in mind that the database specified in the JDBC URL must already exist before you start AMC (otherwise AMC won't be able to create one

and populate it). This is not needed if Derby is used because it supports the "create=true" option in the JDBC URL, thus causing AMC to automatically create the database (if it does not exist) when started.

Solution Views

Use the Solution Views window to view, Add, Modify, and Delete Solution Views.

- To add a Solution View, click the **Add** button on the toolbar.
- To modify an existing Solution View, select the Solution View and click **Modify**. Follow the steps in the **Modify Wizard**. Under **Modify Solution View**, click **Next** to go to the next step, and click **Finish** when you have completed the steps.
- To configure Access Control Lists for a Solution View, select the Solution View for which you want to configure ACLs and select **Configure ACLs...** on the toolbar.
- To delete an existing Solution View, select the Solution View you want to delete and click the **Delete** button on the toolbar.
- To launch a separate panel to Add / Edit / Modify local AM variables for that Solution View, click **Local Variables...**

Note: You must reload Solution Views created using the Auto Update option.

When you **Modify a Solution View** AMC checks to see if the Solution View was created using **Auto Update**. If the Solution View selected for modification was created using Auto Update, a message appears, saying:

The selected Solution View is marked for auto update. Ensure that auto update is disabled to modify the Solution View.

Solution Views are listed in the Solution Views table. If a specific Solution View was created using Auto Update, a >> short menu appears when you click on the arrows up and to the right of the Solution View name. You can select **Refresh Solution View** or **Disable Auto Update**. For Solution Views marked for auto update, you must reload the config file and refresh the Solution View by clicking the **Refresh Solution View**. If a user fails to refresh a Solution View created using the **Simple** option and flagged for auto update, the Solution View may cause inconsistencies in the AMC database. Inconsistencies in Solution Views that are not updated could result in incorrect behavior by the Action Manager.

Configure ACLs

From this window you can set the Access Control Lists (ACLs) for a user and associate that user with a specific Solution View.

- To configure a user or users, select the user or users you want to configure and click **Configure Users** on the toolbar.
 1. Select the user you want to assign a role to from the **User ID** drop-down menu.
 2. Select the radio button next to the role or roles you want to assign the selected user:
 - **Read** - Allows the user to read Solution View details like ALs, Tombstones, logs, properties belonging to the Config, and so on.
 - **Execute** - Allows the user to read and start/stop AssemblyLines
 - **Admin** - Grants the user Reader and Execute roles. This role also allows users to delete logs and tombstones.
 - **Config Admin** - Grants the user the ability to start and stop a Config, modify the Solution View, and assign and modify ACLs for other users.
 3. Click **Apply**.
- To remove an existing user, select the user from the table and click **Remove**.

When you are finished making changes, click **Apply**.

Local variables

Select **Solution Views** from the AMC left hand navigation pane. The **Solution Views** window appears. Select **Local Variables** from the toolbar. In the **Local Variables** window, you can select and **Add**, **Modify**, or **Delete** local variables for a Solution View.

The Action Manager triggers and actions must provide support for local variables that you can set or increment using rules and actions. Local variables can be used as triggering conditions for other rules. For example, a local variable can be set to a value of 1 and then can be incremented for every occurrence of the event and the local variable (in this example, the number 1 set to increment for every occurrence of the event) – the local variable can trigger the rule "Terminate AssemblyLine". When the variable reaches a value of 10, you can configure a new rule to be triggered. The new rule could start a new AssemblyLine on a different server. Set these "local", AM-specific variables to a "Solution View". That means that the one variable created in a rule belonging to one Solution View can only be used in that Solution View's rules and is not accessible to rules of another Solution View.

Add a Solution View

The purpose of a Solution View is to give users access to information in the configuration file without granting them the ability to edit the configuration file directly. Administrators can use a Solution View to filter a configuration file for specific information so that only certain information within the configuration file is displayed. You can create multiple Solution Views for each Config, with each view exposing different information contained in the configuration file.

To Add a Solution View, select **Solution View** and select **Add** on the toolbar of the **Solution Views** window.

1. Enter view details:
 - a. Enter a name for the Solution View in the **Solution View Name** field.
 - b. Enter a description of the Solution View in the **Description** field.
2. Select the **Server** and **Configs** (configuration file) you want to use to create a Solution View:
 - From the **Server** menu, select the IBM Tivoli Directory Integrator server containing the configuration file you want to use to create a Solution View. This menu is empty if no IBM Tivoli Directory Integrator servers have been added to the Administration and Monitoring Console.
 - Select the configuration file you want to use to create a Solution View from the **Configs** list. The menu contains all currently loaded Configs.

Note: Click the **View config files** button to go to the Config files window. You can perform load or unload operations for the configs in this window.

3. Click **Add** on the **Solution Views** toolbar.
 - a. Type the name of the solution view you want to create in the **Solution View Name** field.
 - b. Type an optional **Description** for the Solution View you are creating.
 - c. Select the **Server** that contains the configuration file and AssemblyLines you want to use for creating a Solution View.
 - d. Select the configuration file you want to use from the **Configs** list.
 - e. Enable or disable **Auto Update**.

When the AssemblyLines or properties for a configuration change, **Auto Update** automatically changes the Solution View.

Note: When Auto Update is selected, you cannot edit the Solution View you created with Auto Update on, nor can you create Rules and Triggers for Solution Views made while Auto Update is on. If you want to edit the Solution View or add Rules and Triggers, you must disable Auto Update. The users would have to disable the auto update functionality in order to be able to create a Rules and triggers for Solution Views marked for auto updation. Review any changes to the config in Solution View by using the Refresh button on

the Solution View window. This button is only be visible to configs with auto-update set to true. Any config created manually using the Create Solution View wizard has the auto-update flag set to false.

Note: You must reload Solution Views created using the Auto Update option. Use the **Refresh Solution View** in the **Solution Views** window. For Solution Views marked for auto update, you must reload the config file and refresh the Solution View by clicking the **Refresh Solution View**. If a user fails to refresh a Solution View created using the **Simple** option and flagged for auto update, the Solution View may cause inconsistencies in the AMC database. Inconsistencies in Solution Views that are not updated could result in incorrect behavior by the Action Manager.

4. Use the following options in creating a solution view:

Simple

Create a Solution View with common default options.

Auto Update

For Solution Views marked for auto update, you must reload the config file and refresh the Solution View.

Create Solution View from published solution.

Creates the Solution View from the published solution as specified in the TDI Configuration Editor (CE). This option requires that your active configuration instance have a published solution associated with it, and also requires a TDI 7.0 server.

Create Solution View with all AssemblyLines exposed.

Creates a Solution View with all AssemblyLines from the config instance exposed, and no properties and no Health AL defined. Use this option for a quick start (useful for development purposes). Available for TDI 6.0 and later servers.

Create Solution View with all AssemblyLines exposed and all properties exposed.

Creates a Solution View with all AssemblyLines from the config instance exposed, and all properties and no Health AL defined. This option does not expose the Java properties. Available for TDI 6.1 and later servers. Use this option for a quick start (useful for development purposes).

Create Solution View with all AssemblyLines exposed and all User properties exposed.

Creates a Solution View with all AssemblyLines from the config instance exposed, and all properties and no Health AL defined. This is similar to a quick start type of option. This option is disabled for TDI 6.0 servers (because TDI properties are not available in TDI 6.0 Servers)

5. Click **OK** to finish creating the Solution View.

Config files (allows loading/reloading of configurations)

To reach the Config Files window, and to access options for loading, reloading, unloading, and refreshing of config files, select **Solution Views** in the left navigation area. Select a server and a config file, then click the **View Config Files** button. This launches the Config Files window. This window displays loaded Configs and the Configs in the configs folder of the remote IBM Tivoli Directory Integrator server. When AMC is connected to a Tivoli Directory Integrator server, the Config Files window shows a listing of all files in the remote config folder (whether the files are valid TDI config files or not). You should perform Load operations on valid Tivoli Directory Integrator Config files only, otherwise an error message displays in AMC. The status Loaded or Unloaded displays with green (Loaded) and red (Unloaded) icons in the Status column. You can select one or more configs from the **Select** column of the Confi files table. Once you have selected a config, you can **Load**, **Load As...**, **Unload**, **Reload**, or **Refresh** using the buttons at the top of the table. If you want to load a password protected Config, select the Config and type the password in the Password field.

Whether an action is successful or unsuccessful, a message displays after the action (Load, Load with Run name, Reload, Unload, and Refresh) executes, describing the outcome. For Load, Reload, and Unload, the new status for the configs that you selected displays in the Status column.

Note: You must have superadmin or config admin privileges to perform these actions.

- To load Configs, select the configurations you want to load and click **Load**.
- To load multiple instances of one Config, select the config you want to load and click **Load As....** The **Custom Load** window opens, allowing you to specify **Config File**, **Config Run name**, **Config Password**, and **Property Store Value**.
- To unload Configs, select the configurations you want to unload and click **Unload**.

Note: Loading a server does not automatically start the AssemblyLines associated with the selected Config. Only those AssemblyLines designated as AutoStart starts upon loading.

- To reload Config, select the loaded configurations you want to reload and click **Reload**. You can only reload a configuration that has the status of Loaded.
- To refresh Configs, click **Refresh**. Information for all of the configs in the table is redisplayed.
- Click **Close** when you are finished making changes.

Custom load: The Tivoli Directory Integrator server supports loading multiple instances of the same config with different run names. If you load config instances using **Load As...**, you can use these configs to create Solution Views and Rules. Use AMC to load multiple config instances by performing these steps:

1. From the **Welcome** page, select **Servers -> Config Files**.
2. Click **Load As...**
The **Custom Load** window appears.
 - a. Select the **Config File** from which you want to create multiple instances and click **Go**.
 - b. Type the **Config Run name**.
 - c. Type the **Config Password**.
 - d. Type the **Property Store Value** for each Property Store Name.
3. Click **OK** to use the values you have entered to create an instance of the config with the Run name you have specified. After an instance of the config is created, you are returned to the **Load Reload** window.
4. Click **Cancel** if you do not want to create the config with the values you have specified in the **Custom Load** window.

Note: Users must maintain data integrity.

- For example, if a Solution View and rules have been created for a config named config1.xml, and with a run name of ABC, do not load a different config, for example, config2.xml, with the name ABC either as a solution name or a run name.
- If you want to reuse Solution Views that you created using a specific run name and set of property files, you must unload this config using the same run name and property files.

Monitor Status and Action Manager

If you have not done so already, expand the **Monitor Status** category in the main navigation area of the Administration and Monitoring Console.

Do one of the following:

- To view information about each Solution View, see the Monitor Status table. Information regarding the Solution Views, such as Action Manager Status, Health Check Result and Health Check Status, display. You can also display **Solution View Details**, **Server Information**, and **Show Preferred Views**.
- To add, edit or delete Action Manager rules, click “**Action Manager**” on page 211.

Monitor Status

This window displays the views selected on the Preferred Views window accessed from **Advanced** -> **Preferred Solution Views**. It displays high level information about each preferred Solution View, such as:

Action Manager Status

Displays the status of the Action Manager rules for the selected Solution View: A blue exclamation mark indicates that no Action Manager rules have been triggered recently. An yellow triangle containing an exclamation mark indicates that an Action Manager rule has been triggered recently.

Health Check Result

Displays the health check result obtained from the healthAL.result final work entry attribute in the Solution View's Health AssemblyLine. This value is displayed as text.

Health Check Status

Displays the health check status obtained from the healthAL.status attribute in the Solution View's Health AssemblyLine.

Additionally, if you have designated a .gif file with the same name as the returned status value in the Administration and Monitoring Console's resources/amc_images/healthAL directory, the .gif image is also displayed in this column. For example, if the healthAL.result is returned as "Error", and you have created an "Error.gif" in the above mentioned directory, the Error.gif image displays in the table column.

From this window you can:

- View Solution View details - To view the details of a specific Solution View, select the desired Solution View and click **Solution View Details**
- View Tivoli Directory Integrator Server Information - To view the details of the server to which the Solution View belongs, click **Server Information**.
- Show Preferred Solution Views - Click **Show Preferred Views** to view preferred Solution Views. This button is visible only if Preferred Solution Views are defined. You can define preferred Solution Views on the "Preferred Solution Views" window under **User Preferences**.

Solution View Details: The Solution View details panel in turn provides deeper view of the details specific to a Solution View which an administrator can take a look at and take action upon.

This window contains two tables. The top table displays the AssemblyLines associated with the selected Solution View and the status of each Solution View. The bottom table displays log information about recently triggered Action Manager rules.

When you are through making changes, click **Close**.

Solution View Details Table:

Columns: The **Solution View Details** table contains the following columns:

Select Select the radio button next to the AssemblyLine on which you want to perform an action.

AssemblyLines

Displays the name of the AssemblyLine.

Status Displays the AssemblyLine's status; for example, **Running** or **Stopped**.

Start Time

AssemblyLine is running

Start Time is when the running AL started. Start Time is based on the running AL.

AssemblyLine is stopped

The time when the last run of the AL started. Start Time is based on the most recent tombstone entry for the AL. (Available only with Tivoli Directory Integrator 7.0 servers).

Last Stop Time

The time when the last run of the AL terminated. Stop Time is based on the most recent tombstone entry for the AL. (Available only with Tivoli Directory Integrator 7.0 servers).

Statistics

Displays the current statistics of the running AssemblyLine.

Actions: You can choose the operations you want to perform from the tool bar at the top of the table or using the **Select action** drop-down menu, such as:

- View Tombstones - Select the AssemblyLine you want to view and click the **View Tombstones** button
- View Logs - Select the AssemblyLine you want to view and do one of the following:
 - Click the **View Logs** button on the toolbar.
 - Select **View Logs** from the **Select action** drop-down menu and click **Go**.
- Manage Properties - Select the radio button next to the AssemblyLine with properties you want to manage and click the **Manage Properties** button on the toolbar.
- Start AssemblyLine -
 1. Select the AssemblyLine you want to start
 2. Click the **View pop-up** button
 3. Click **Start AssemblyLine**.
- Stop AssemblyLine - Select the AssemblyLine you want to stop and do one of the following:
 1. Select the AssemblyLine you want to stop
 2. Click the **View pop-up** button
 3. Click **Stop AssemblyLine**.
- Solution View Details - Click the **Solution View Details** button. Select the component you would like to view, for example, AssemblyLines.

Start AssemblyLine: Run the selected AssemblyLine.

Start AssemblyLine synchronously: AMC waits for the AL to terminate and shows the status of the run AL periodically. The output schema attributes of the AssemblyLine after its termination are viewable for synchronous AL runs.

Start AssemblyLine in simulate mode: The Assembly Line executes all components except for the connectors in the add, update, and delete modes. In essence, putEntry, modEntry and deleteEntry methods of connectors are not invoked in simulate mode. As a result, an Assembly Line running in simulate mode does not perform any additions, modifications, or deletions on third party repositories. For more information on simulate mode, see the corresponding section in *IBM Tivoli Directory Integrator V7.0 Users Guide*.

View Tombstones: If you have tombstones enabled on the remote IBM Tivoli Directory Integrator server, the Administration and Monitoring Console can display the tombstone entries for terminated AssemblyLines. This window displays useful information about tombstone entries, such as when the entry was changed to the tombstone state.

Delete Tombstones: On the **Monitor Status** window, select an AssemblyLine. Select the arrow to the right of the AssemblyLine and select **Delete Tombstones** from the menu. This launches the **Delete Tombstones** window. The component details section of this window identifies the Solution View and AssemblyLine that are being worked on. In the choose delete criteria section, select one of the options to specify which tombstones you want to delete:

- Select **All Tombstones** to delete all of the tombstones for the selected AssemblyLine.

- Use **Start Date** and **End Date** to specify the date range from which the tombstones are to be deleted. AMC calculates the number of days from the selected date to the current date. AMC then deletes the tombstones generated for the calculated number of days.
- Use **Number of entries to return** to indicate a whole number indicating the number of recent tombstones to delete. When you click **Delete**, a confirmation message appears. When you confirm, AMC executes the delete command.

View Logs: Logs for a given AssemblyLine are displayed on the View Logs window. **Monitor Status** -> **Solution View Details** -> **View Logs** to view the list of log files for the selected AssemblyLine, click the radio button next to the log you want to view and click **View Logs**.

Note: In order to view an AssemblyLine log in the Administration and Monitoring Console, the AssemblyLine must log using the SystemLog logger.

Action Manager results table: When a rule set in the Action Manager is triggered, information about the violation is logged, such as the source of the violation, a description of the error and the time at which the violation occurred. These details are displayed in the **Action Manager Results** table.

The following sections contain information about the **Action Manager Results** table columns and how to perform operations on Action Manager Results.

Columns: The **Action Manager Results** table contains the following columns:

Select Select the radio button next to the message on which you want to perform an action.

Source

Displays the name of the Action Manager rule that was triggered.

Severity

Displays the severity of the message.

Message

Displays the message associated with the Action Manager action.

Description

Displays additional information about the message.

Timestamp

Displays the time at which the Action Manager rule was triggered and the message was generated.

Actions: Select the result or results you want to delete and click **Delete**.

Server Information: This window displays the IBM Tivoli Directory Integrator server information of the server to which the currently selected Solution View belongs. The information on this window is read-only, although administrators have the capability to shut the server down from this window.

Show Preferred Solution Views: Preferred Solution Views are the default Solution Views that are displayed on **Monitor Status** window.

Action Manager

This window allows you to add, delete or modify rules, triggers and actions to be performed as a result of rules execution and triggering conditions.

Add/Edit configuration rules: Using the settings on this window you can create an “Action Manager” on page 189 (or modify an existing one) for the current Solution View.

A rule consists of two parts:

- The condition under which the rule is to be invoked, called a "trigger."
Some examples of triggers are Server API failure, AssemblyLine failure, or failure of the AssemblyLine to run at the specified intervals.
- The set of alternate actions to be performed when the trigger is encountered.

Configuration rules settings: This window is concerned with the first part of the rule: defining triggers. From the window you can select a name, description, and trigger type.

Name

Enter a name for the rule. If you are adding a rule, this field is required.

Description

Enter an optional description of the rule.

Trigger type

The trigger type defines the conditions under which a rule is invoked. From the drop-down menu, select a trigger type:

No trigger

Rule has no triggering condition.

On AssemblyLine termination

Rule is triggered when the specified AssemblyLine is terminated.

On Config Load

Rule is triggered when the Action Manager receives a Config load event for this particular config.

On Config Unload

Rule is triggered when the Action Manager receives a Config Unload event for this particular config.

On Query AssemblyLine result

Rule is triggered when the last "work" entry of the specified AssemblyLine contains an attribute matching a given condition and value.

On server API failure

Rule is triggered when the Action Manager is unable to connect to the remote server using the Server API. This rule is triggered only once. The rule resets when it detects that it can reconnect to the server using the Server API.

On received Event

Rule is triggered when the Action Manager receives an event that meets the criteria specified in the Event type, Event Source and Event Data fields.

On Property Trigger

Rule is triggered when the specified property meets the determined Property name, Condition and Value specifications.

On Local Variable

Rule is triggered when the specified variables meet the specified condition. The Action Manager periodically checks for this property.

Note: This rule gets triggered only once, and gets reset back to ready state only when Action Manager detects that this variables does not meet the specified criteria any longer. The rechecking ensures that the rule is not repeatedly triggered for a single occurrence of the triggering condition.

Inspect AssemblyLine Exit Code

Rule is triggered when an AssemblyLine terminates abnormally. You can define an error object that Action Manager searches for in the AssemblyLine Exit Code.

Time since last execution

Rule is triggered when the specified AssemblyLine has not run for the determined period of time.

Timer Trigger

Rule is triggered continuously within the given interval.

Configure trigger: Each trigger type has a different selection of settings. If you do not see some of the fields listed below on your window, it is because the trigger type you currently have selected does not support them.

Source

Enter the source you want to monitor.

Data Enter the data you want to monitor.

Property name

From the drop-down menu, select the property name you want to monitor.

Condition

Select the condition you want to use to compare the property and value. Possible options are:

- equals
- not equals
- greater than
- less than

Value Enter the value you want to monitor.

Configured actions: From this table you can add, delete, and modify actions. You can also move actions up and down in the table. For every action in the configured actions table that you can select, there is a column where you can enable the special trigger **Execute on Error**. **Execute on error** performs the action you have selected when an error condition occurs.

- To select an action to manage, enable the radio button that precedes each action that is listed.
- To add an action, click **Add**.
- To delete an action, select the action you want to delete and click **Delete**
- To modify an action, select the action you want to modify and click **Modify**.
- To move an action up one position in the table, select the action you want to move and click **Move Up**.
- To move an action down one position in the table, select the action you want to move and click **Move Down**.

Selecting **Execute on Error** carries out actions only if an error has occurred during the execution of any of the previous actions. You can use such actions to take corrective measures for handling any error that might have occurred during the execution of any previous actions. Action Error variables: AMC and Action Manager allow you to make the action error available in the various actions. At any point of time, if an error occurs during the execution of any configured actions, this error becomes available to you in the form of special reserved variables. You can then use these reserved variables in other actions you have configured. When the following actions are executed, Action Manager replaces the string `%Action_Error%` by the actual error that occurred during the execution of the previous actions. If no error occurs, the variable `%Action_Error%` is not be replaced and stays as it is.

- Send Email
- Execute command
- Send event action
- Write Log action

Add/Modify Action: When a rule is triggered, the Action Manager executes the actions associated with the rule. This window allows you to specify or modify the actions you want Action Manager to take when the rule is triggered.

From the drop-down menu, select an action type, and configure it. Click **OK** when you are finished.

Start AssemblyLine

This action starts an AssemblyLine. If you select this action, you must specify the name of the AssemblyLine you want to start and its associated Config (and possibly the Config's password).

Server This is a drop-down list of configured Servers. LocalServer means the Server on the computer Action Manager is executing.

Select from remote config folder

Check box; if enabled, queries the remote Server for available Config files. The Config files displayed are those present in the folder whose path is specified for the `api.config.folder` property in the `global.properties` file.

Config name

Enter the Config to which the AssemblyLine in the AssemblyLine field belongs. If **Select from remote config folder** is checked, you are presented with a list of available Config files on the remote Server, if unchecked, you must fill in the name of a locally-available Config file.

This field is required.

Config password

If required, enter the Config password for the selected Config file. This field is applicable only if the config is password protected.

AssemblyLine

Enter the name of the AssemblyLine to start.

Configure AssemblyLine Operation

This hyperlink launches the 'Select Operation' dialog. If the AssemblyLine has been defined with one or more custom Operations, this dialog enables you to select such an Operation. Subsequently, you are prompted for the AssemblyLine's Initialization attributes and Operation attributes for this Operation. This label is shown only for TDI 6.1.X and Tivoli Directory Integrator 7.0 servers if configured and is not applicable for TDI 6.0.

Stop AssemblyLine

This action stops an AssemblyLine. If you select this action, you must specify the name of AssemblyLine you want to stop and its associated Config.

Server This is a drop-down list of configured Servers. LocalServer means the Server on the computer Action Manager is executing.

Select from remote config folder

Check box; if enabled, queries the remote Server for available Config files.

Config name

Enter the Config to which the AssemblyLine in the AssemblyLine field belongs. If **Select from remote config folder** is checked, you are presented with a list of available Config files on the remote Server, if unchecked, you must fill in the name of a locally-available Config file.

This field is required.

AssemblyLine

Enter the name of the AssemblyLine to stop.

Enable/Disable Rule

Select the Enable/Disable Rule to enable or disable an Action Manager rule.

Rule name

Select the name of the rule-Solution View pair that you want the action "Enable/Disable Rule" to execute. In versions before TDI 7.0, you selected the rule name instead of a rule-Solution View pair, which is a new feature for 7.0. This option belongs to the action "Enable/Disable Rule."

State Select the desired state from the drop-down menu. If you want to enable the rule in the **Rule name** field, select **Enabled**. If you want to disable the rule, the select **Disable**.

Execute Rule

This action causes the Action Manager to execute the specified rule. Action Manager then executes the actions associated with the specified rule. The trigger condition associated with the specified rule is not required to be satisfied.

Rule name

Select the name of the rule-Solution View pair that you want the action "Execute Rule" to execute. In versions before TDI 7.0, you selected the rule name instead of a rule-Solution View pair, which is a new feature for 7.0. This option belongs to the action "Execute Rule."

Execute Command

The Execute Command action can execute the command entered in the Command field on the target computer specified under **Target Computer Name**. The command can be any generic command or a IBM Tivoli Directory Integrator specific command. The Execute Command can be used when a user configures a rule to execute commands that are specific to the target computer or to execute Tivoli Directory Integrator commands that are not exposed by AMC. For example, in AMC we do not have actions that can restart a server or load a config. The user has to perform the restart or reload commands using either the TDI Server or Config Files windows. If any error occurs while executing the command, it is captured in the %ACTION_ERROR% variable, which can be further used by the Action Manager,

Target Computer Name

Name or IP address of the target computer. Action Manager connects to the computer specified in this field. If neither a computer hostname nor an IP address is specified, the command executes on the computer where the Action Manager is running.

Port Port specifies the channel over which the Action Manager can connect to the target computer where the command is to be executed.

Username

The user name is verified for authentication and authorization when establishing a connection with the target computer.

Password

The password is verified for authentication and authorization when establishing a connection with the target computer.

Keystore

Keystore path is entered and used in case certificate authentication is required when connecting to the target computer.

Keystore Password

Keystore password is required when certificate authentication is mandatory for connection to the target computer.

Protocol

The protocol that is to be used for establishing a connection with the remote machine. Protocol can have the following values, WINDOWS, RSH, SSH OR REXEC (Windows, remote shell, secure shell, or remote execution protocols).

Command

Command that is to be executed.

Notify Event

This action causes the Action Manager to send an event with the specified details to the IBM Tivoli Directory Integrator server associated with the current Solution View. To add this action to the rule, select **Notify event**. If you select this action, you must specify an event type.

Event type

Enter an event type. This field is required.

Source

Enter a source for the event type.

Data Enter data for the event type.

Modify property

This action causes the Action Manager to modify a property based on a specific operation and value. If you select this action, you must also select a value.

Property name

Select the property you want to modify from the drop-down menu.

Operation

From the drop-down menu, select the operation you want to use to modify the property. Possible options are:

- Set
- Increment
- Decrement

Value Enter the desired value. This is a required field.

Copy property value

This action causes the Action Manager to copy the value of the source property to the destination property.

From property

From the drop-down menu, select the property you want to copy from.

To property

From the drop-down menu, select the property you want to copy to.

Write to log

This action creates a log of the Action Manager rules that have been invoked, according to the specified severity, message and description. This log can be viewed under **Monitor Status**, on the "Solution View Details" window in the **AM results** table. Having at least one log action for every rule is recommended. If you select this action, you must enter a message in the **Message** field.

Severity

Select the desired severity from the drop-down menu. Possible options are:

- Severe
- Warning
- Info
- Fine

Message

Enter the desired message.

Description

Optionally, enter a description.

Send Email

This action causes an email to be sent to the recipient you specify. You supply the content of the email. Along with the content, the Action Manager provides other details before sending the mail. In the content input area as well as in the subject line, you can specify the variable `%EVENT_DATA%` value. Specifying `%EventData%` inserts the actual value of the Eventdata variable when the mail is sent. `%Action_Error%` can also similarly be substituted here. If Attach Action Manager Log is enabled, the Action Manager logs (as specified in the `am_logging.properties` file) are sent as an email attachment. In the content input area, you can specify the variable `%EVENT_DATA%` value. Specifying `%EventData%` in the content puts the actual value of the Eventdata variable when the mail is sent. `%Action_Error%` is also similarly be substituted here. If **Attach Action Manager Log** is enabled, the Action Manager logs (as specified in the `am_logging.properties` file) are sent as an email attachment.

Substitute variable for event data: Select **Action Manager** from the left navigation pane or select **Action Manager** from the **Welcome** screen. Under Action Manager, you can Add a rule to a Solution View. You can name a new rule, and edit or delete an existing rule. You can make Event Data available when configuring or sending that data to other actions.

Use Action Manager to make event data available when configuring actions for a trigger. In Action Manager, you can **Add**, **Modify**, or **Delete** a rule. When you add a rule, you name the rule and select the **Trigger type**. Variable substitution allows you to select data that is output from certain Action Manager triggers, and use that data in certain actions that are triggered by a rule. AMC and Action Manager make the data available to the triggered actions in the form of a reserved variable. The action then uses the data that is stored in the variable. You can use this reserved variable in any of the actions you have configured for this trigger.

Triggers that can produce event data: The following trigger types can produce event data that can be consumed by actions:

- On Start AssemblyLine - Event data is available as `%Event_Data%`.
- On AssemblyLine Terminate - Event data from On AssemblyLine Terminate is available as `%Event_Data%`.
- On Received Event - Event data from the received event is mapped as `%Event_Data%`.
- On Local Variable - Event data from the Local variable event is mapped as `%Event_Data%`.
- On Config Load - Event data from the On Config Load event is available as `%Event_Data%`.
- On Config Unload - Event data from the trigger would be available as `%Event_Data%`.
- On query AssemblyLine result - Event data is available as `%attribute_name %`. The `%attribute_name%` variable is replaced with the details about the actual attribute from the last work entry.
- Inspect AssemblyLine Exit Code - Event data is available as `%attribute_name %` and `%Event_Data%`.
 - Inspect Error Object set to enabled - While configuring the Inspect AssemblyLine exit code trigger, if the user enables Inspect Error Object (sets the option to true), the `%Event_Data%` variable is replaced with actual error data. The `%attribute_name%` variable is not available for actions.
 - Inspect Error Object set to disabled - While configuring the Inspect AssemblyLine exit code trigger, if the user sets the Inspect Error object to disabled (sets the option to false), the `%attribute_name%` variable is replaced with the details about the actual attribute from the last work entry. The `%Event_Data%` variable is not be available for actions.

Actions that can access event data: The actions executed for each of the above triggers can access the event data produced by the triggers using the `%Event_Data%` variable. Every occurrence of `%Event_Data%` is replaced with the actual event data for that trigger. The following action types can use event data available from their respective triggers:

- Notify Event - Users can specify the `%Event_Data%` variable in the Data text field only.

- **Write to log** – Users can see a log message that is logged to a database. If the log message, after substitution for the %Event_Data% variable, exceeds 500 characters, the log message is truncated to the first 500 characters. This is because the database has a limit of 500 characters only.
- **Send E-mail** – Any event data specified by %Event_Data% or error data specified by %Action_Error% is substituted in the subject line of the email. Action Manager appends other data about execution before sending the mail. You can specify the variable %EVENT_DATA% value in the content textbox. Specifying %EventData% in the content substitutes the actual value of the Eventdata variable when the mail is sent. You can also similarly substitute %Action_Error% here. If **Attach Action Manager Log** is enabled, the Action Manager logs (as specified in the am_logging.properties file) are sent as an email attachment.

For any action being executed, such as the Send E-mail action, the Execute command action, the Log action, the Start AssemblyLine action, and so on, executes in response to the same trigger, the string of %Event_Data% automatically gets replaced by the event data generated by that trigger.

View Rules Summary: To view the current Action Manager for the selected AssemblyLine, click **View Rules Summary**. The table lists all the defined rules, triggers and actions associated with the Solution View. When you are done viewing, click **Close**. Only those rules which are in Enabled state are listed here.

Property Stores

If you have not done so already, expand the **Property Stores** category in the navigation area of the Administration and Monitoring Console. To add or edit Java, Solutions, Global, System, User Property and Password Store properties, click **Advanced → Property Stores**.

When you are done entering the desired the property values, click **OK** to save your changes.

The order in which these Property Stores are listed is significant. The Property Stores are evaluated from top to bottom, but the last definition of a given Property is the one that is used. By default, the system is set up such that properties defined in a solution-specific properties file called `solution.properties` (residing in the Solution Directory) override corresponding ones in the system-wide `global.properties` file.

Note: Certain System Properties and Java Properties are read-only. These read-only properties are shown in the respective Property Stores. Trying to modify these properties has no effect.

Select Solution View

This window allows you select a Solution View. The menu contains only those Solution Views for which you have any type of access rights amongst *Read*, *Execute*, *Config_Admin* or *Admin*. Nothing is displayed if you do not have access to any of the Solution Views being created. Once you have selected a view, click **Set**.

After you select a Solution View, you can manage properties by clicking on the other property tabs, such as **Solution Properties** and **Global Properties**.

Solution Properties

This window allows you to add, edit and delete properties in the Solution Properties list.

Global Properties

This window allows you to add, edit and delete Global Properties.

Java Properties

This window allows you to add, edit and delete Java properties.

System Properties

This window allows you to add, edit and delete System properties.

Password Store

This window allows you to add, edit and delete properties in the Password Store.

User Property Store

This window allows you to add, edit and delete properties in the User Property Stores list.

The Property Stores drop-down menu contains a list of property stores configured by the user. Global, Solution, Java and Password Stores properties are not included. Select the property store whose associated properties you wish to view, add, edit or delete.

Log Management

If you have not done so already, expand the **Advanced** category in the navigation area of the Administration and Monitoring Console. To delete log files for all AssemblyLines, for a particular AssemblyLine, or to delete by date, click **Log Management**. When you select a new Solution View, you can click **Refresh**. Clicking **Refresh** lists all of the AssemblyLines that belong to the Solution View you just selected.

This window allows you to select the name of a Solution View. The AssemblyLines listed for deletion are taken from the Solution View you choose. You can choose to delete log files for all AssemblyLines, or for a particular AssemblyLine. You can also specify logs to delete by date. To manage display and deletion of logs:

1. Select the Solution View with the AssemblyLines whose logs you want to clean up from the **Solution View** menu.
2. In the **Choose Component** section, do one of the following:
 - Select the **All AssemblyLines** radio button to delete the logs of all AssemblyLines within the selected Solution View.
 - Select the **Specific assembly line** radio button to delete only those logs associated with a specific AssemblyLine.
3. If you selected **Specific assembly line**, select the AssemblyLine with logs you want to delete from the menu.
4. In the **Display Log Files** section, do one of the following:
 - To delete all logs belonging to the selected AssemblyLine(s), select **All**.
 - To delete logs within a date range, use the **Start Date** and **End Date** options. Logs created within the two dates specified will be deleted. Enter the desired dates in the date field; its format is locale-dependent. You can also use the Calendar button, which lets you specify a date by choosing from a calendar.
 - For pre-Tivoli Directory Integrator 7.0 servers, to delete logs older than a certain date, select the **End Date** option. All logs older than the date specified are deleted.
 - To preserve your most recent logs, select the **Display first** radio button. Enter the number of recent logs you want to save. Used to specify that keep those log files that are recent and list the others. You draw the line on *recent* using the edit box. If you type 20, you are telling AMC to keep the most recent 20 log files and list the rest of the files in the table so that they are available for deletion. If you type the number 10, the 10 most recent logs are saved.
5. In the **Logfiles** table, a list of log files displays. In the Select column, select any logs you want to delete and click **Delete**. In the **Select Action** menu, you can choose any of the following options:
 - Export data
 - Change all selected
 - Collapse table
 - Restore

When you have selected one of these options, click **Go**.

6. From the display that results from the criteria you have selected, choose the logs you want to delete and click **Delete** to remove the specified logs. When you are finished deleting logs, click **Close** to exit this window.

Preferred Solution Views

You can select the Solution Views that you want to be loaded by default in the monitor window, using the **Preferred Solution Views** panel.

The "preferred" Configs are shown by default in the monitor status page when it is opened. If there are no views defined then this panel will simply display a message saying that there are no views. Once a set of views are defined then a user can set what he would like to see as the default views. This panel can be viewed by any user who has a set of views assigned to him by the superadmin.

You can make a Solution View preferred by selecting its checkbox in the **Select** column, and click **Enable As Preferred**.

Conversely, you can disable the Preferred status for a Solution View by selecting its checkbox in the **Select** column, and click **Disable As Preferred**.

AMC and AM Command line utilities

A number of command line utilities are included with AMC and its associate product, the Action Manager (AM). These command line utilities help in installing, uninstalling or re-installing the AMC war file. There are also scripts for backup and restore, as well as a migration script. The migration script is for migrating to future versions of AMC and AM, and not for migrating from previous version to the current version. All these scripts get installed in the *TDI_install_dir/bin/amc* directory.

install The install.bat (.sh) script is used to deploy the AMC console module on ISC SE or ISC AE. The script relies on the setupCmdLine script for setting up the necessary environment variables and the tdiISCHome script for determining the location of the ISC runtime and also the type of runtime being used, that is, whether it is LWI or WAS. This script is called by the installer.

Usage: install

This script does not take any parameters.

uninstall

The uninstall.bat (.sh) script uninstalls the AMC console module from ISC SE or ISC AE. The script relies on the setupCmdLine script for setting up the necessary environment variables and the tdiISCHome script for determining the location of the ISC runtime and also the type of runtime being used, that is, whether it is LWI or WAS.

Usage: uninstall

This script does not take any parameters.

backupamc

The backupamc.bat (.sh) script backups all the configuration related information of AMC (configuration files, logs, and so forth.) A backup_tdiamc folder will be created inside the backup directory.

Usage: backupamc [-d *folder_to_create_backup_in*]

If the -d option is not specified, the files are copied to the *TDI_install_dir/bin/amc/ActionManager/backup_tdiamc* directory.

The following files are backed up:

1. amc.properties
2. logging.properties
3. amcdbschema.xml
4. amcdbhandler.properties

restoreamc

The restoreamc.bat (.sh)script will restore the backed up files to a fresh AMC deployment. The backed up files need to be first obtained by using the backupamc script for this to work.

Usage: restoreamc

This script does not take any parameters.

migrateamc

This provides a single backup, restore, uninstall and install command. This will backup the old AMC data, uninstall the old AMC plug-in archive, installs the new AMC and restores the old AMC configuration data.

This script requires the new AMC plug-in archive to be copied into the *TDI_install_dir/amc* directory.

Usage: migrateamc.bat [-d *backup_directory*]

start_tdiamc

This script is a convenient wrapper utility to start AMC. This script internally starts the ISC runtime. If the runtime is LWI then it calls the lwistart command, else if the runtime is WAS then it calls the startServer server1 command. Before starting the ISC runtime the script calls the startNetworkServer command, which is used for starting the Derby database in secured network mode. If the database type is anything other than Derby, then this script only starts the ISC Runtime.

On Windows platforms:

Usage: start_tdiamc [*Service name*]

If a service name is passed, the service will be started instead of calling lwiStart.

On Unix Platforms:

Usage: start_tdiamc

stop_tdiamc

This script is a convenient wrapper utility to stop AMC. This script internally stops the ISC runtime. If the runtime is LWI then it calls the lwistop command, else if the runtime is WAS then it calls the stopServer server1 command. After executing the command the script makes a call to the stopNetworkServer script to stop the Derby database. If the database type is anything other than Derby this scripts only stops the ISC Runtime.

On Windows platforms:

Usage: stop_tdiamc [*Service name*]

If a service name is passed, the service will be stopped instead of calling lwiStop.

On Unix Platforms:

Usage: stop_tdiamc

startAM

The Action Manager is started using the startAM.bat(.sh) script located in the *TDI_install_dir/bin/amc* directory.

Note: The script has the Classpath defined for all the jars required by the Action Manager. There are two variables, CLASSPATH and DB_CLASSPATH. The DB_CLASSPATH has the path separated list of .jar files required for achieving JDBC Connectivity with the database.

When AMC is configured to use Oracle, MS SQL Server or DB2 the corresponding JDBC .jar files of these databases should be added to the DB_CLASSPATH variable. On Windows, the script accepts an optional service name parameter that can be used to start an already registered service:

```
startAM.bat [service name]
```

stopAM

The Action Manager is stopped using the stopAM.bat(.sh) located in the *TDI_install_dir/bin/amc* directory. This script uses the processID of the started AM to kill it. The processID is obtained by the startAM script and is stored in a file, which in turn is read by the stopAM script.

On Windows, the script accepts an optional service name parameter that can be used to stop an already registered service:

```
stopAM.bat [service name]
```

startNetworkServer

This script is used for starting the Derby database server in network mode, on port 1528. The port selected is different from the default port of Derby.

Usage: startNetworkServer

stopNetworkServer

This script is used for stopping the Derby database server in network mode.

Usage: stopNetworkServer

setDBType

This script is used for setting the type of database that you are using. The script sets the property namely DB_TYPE. If the DB_TYPE is set to derby then on executing the startNetworkServer script the Derby database will be started on the host and port that you specified in the startNetworkServer script file. The setDBType also sets the database user name and password. The database user name and password are required by startNetworkServer to enable the BUILTIN security mechanism and to add the user to the list of authorized users.

The setDBType script is called internally by the startNetworkServer and stopNetworkServer scripts for setting the DB_TYPE and the DB_USER and DB_PASSWORD properties.

backupamcdb

This script is used during the migration of an AMC database. The script backs up the AMC database and has the data exported in a Tivoli Directory Integrator defined XML format. This script is called by the installer when you choose the migration path.

Usage: backupamcdb -d *folder_which_contains_AMC_backup* -p *location_of_the_amc.properties_file*

restoreamcdb

This script is used to restore the AMC database during migration. The scripts are called by the installer when you choose the migration path.

Usage: restoreamcdb -d *folder_which_contains_AMC_backup* -p *location_of_the_amc.properties_file*

backupam

This script is used for backing up the Action Manager properties files. The script backs up the am_config.properties and am_logging.properties file.

Usage: backupam [-d backup_directory]

The archived info is created in the backup folder. If the -d option is not specified the files are copied to the *TDI_install_dir/bin/amc/ActionManager/backup_tdiamc* directory.

restoream

This script is used for restoring the properties files of Action Manager which were backed up using the backupam script. The restore script restores the `am_config.properties` and `am_logging.properties` files.

Usage: `restoream [-d backup_directory]`

If the `-d` option is not specified, the files are copied from the `TDI_install_dir/bin/amc/ActionManager/backup_tdiamc` directory.

setAMCRoles

This script is used for mapping the user who is installing TDI AMC to the ISC admin and TDI AMC Admin roles. This script is introduced in TDI 7.0.

Once these roles are granted to the install user, that user has the authorization to add new users and grant them with the necessary roles. The install user becomes the administrator for the AMC console module.

Usage: `setAMCRole username [OS Group]`

The OS Group is an optional parameter while deploying AMC on ISC SE.

tdimigam

This script is used for migration of the `am_config.properties` file.

The usage for this command is:

`tdimigam -f propfile [-b backfile] [-n newfile] [-v] [-?]`

where:

- `-f propfile` - The name of the file to migrate
- `-b backfile` - Backup the original file with the specified name
- `-n newfile` - Name to give the file that is migrated
- `-v` - Enable verbose mode
- `-?` - Prints the usage statement

Logging for this command is controlled by the `tdimigam-Log4J.properties` file.

tdimigamc

This script is used for migration of the `amc.properties` file. The options of this script are similar to those of `tdimigam` and `tdimiggb1` that are used for migration of the `am_config.properties` and `global.properties` files respectively.

The usage for this command is:

`tdimigamc -f propfile [-b backfile] [-n newfile] [-v] [-?]`

where:

- `-f propfile` - The name of the file to migrate
- `-b backfile` - Backup the original file with the specified name
- `-n newfile` - Name to give the file that is migrated
- `-v` - Enable verbose mode
- `-?` - Prints the usage statement

Logging for this command is controlled by the `tdimigamc-Log4J.properties` file.

addAMCService

This script is used for adding AMC as a service on a system.

Usage: `addAMCService Service_Name`

On Windows the script registers the Generic Windows Service executable (*TDI_install_dir/bin/amc/amcwin-service.exe*) from the IBM Platform Integration Toolkit. The Generic Windows Service uses the configuration file *TDI_install_dir/bin/amc/amcwin-service.ini*. That file specifies the name of the service and the start/stop commands. The file is automatically populated by the installer or the "addAMCService" script.

By default the file looks like this:

```
[Service]
ServiceName=$service_name$
WorkingDirectory="$install_dir$\bin\amc"
StartCommand="$install_dir$\bin\amc\amcservice.bat" start amc am"
StopCommand="$install_dir$\bin\amc\amcservice.bat" stop amc am"
```

This means that by default the AMC service runs both the AMC and the Action Manager.

After you call "addAMCService", you can edit the .ini file to customize which components are run by the service (both AMC and AM, just AMC or just AM).

For example to run only the AMC, specify start and stop commands like the following:

```
StartCommand="$install_dir$\bin\amc\amcservice.bat" start amc"
StopCommand="$install_dir$\bin\amc\amcservice.bat" stop amc"
```

To start/stop the service use the GUI "Services" utility (**Control Panel -> Administrative Tools -> Services**) or the Service Controller command line tool:

```
sc start <service name>
sc stop <service name>
```

Beware that in the "Services" utility the display name of the registered service looks like this:

IBM Tivoli Directory Integrator Administration and Monitoring Console – myamc

where "myamc" is the service name that you specified as an argument to "addAMCService.bat".

On UNIX the script appends a line like the following at the end of the /etc/inittab system file:

```
<service name>::once:<install_dir>/bin/amc/amcservice.sh" start amc am
```

To stop the service use the "amcservice.sh" script from the *TDI_install_dir/bin/amc* folder:

```
amcservice.sh stop amc am
```

or just

```
amcservice.sh stop amc
```

if the service runs only AMC.

deleteAMCService

The script is used for removing AMC as a service from a system.

Usage: deleteAMCService *Service_Name*

setDerbyProps

The script sets the necessary Derby database properties that are used by the startNetworkServer and stopNetworkServer scripts.

Usage: setDerbyProps

amcservice

This script starts/stops the whole Administration and Monitoring configuration. The following configurations are supported:

- both AMC and AM

- only AMC
- only AM

Internally the script calls "start_tdiadc"/"stop_tdiadc" and "startAM"/"stopAM". It is intended to be used when registering an Operating System service.

Usage: amcservice [start|stop] [amc] [am]

Examples:

```
amcservice start amc
amcservice stop amc am
```

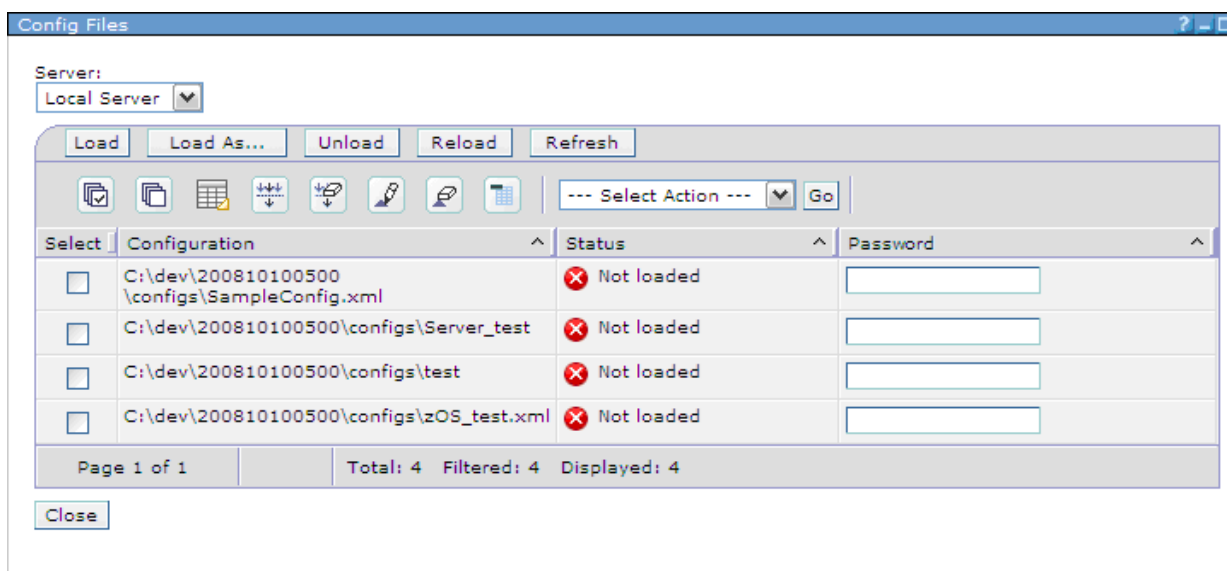
Example walkthrough of creating a Solution View and Rules

This section illustrates all the steps to create a Solution View, configure a rule and trigger it in Action Manager. It is assumed that IBM Tivoli Directory Integrator is installed along with AMC. The configuration SampleConfig.xml used in this example is the one that you are supposed to create following the tutorial in the *IBM Tivoli Directory Integrator V7.0 Getting Started*, "Introducing IBM Tivoli Directory Integrator" -> "Creating your first AssemblyLine". It should be copied into the *TDI_Install_dir*/configs folder, where *TDI_Install_dir* is the installation directory of TDI. This solution reads data from the 'examples/Tutorial/People.csv' file and writes it to 'examples/Tutorial/Output.xml'.

This section illustrates all the steps to create a config view, configure rules, and trigger this rule in Action Manager. It is assumed that Tivoli Directory Integrator is installed along with AMC. The sample config (SampleConfig.xml) and associated files are available in the downloads section and should be copied into the *TDI_install_dir*/configs folder. This solution reads data from the sample.csv file and writes to sample.xml.

Steps:

1. Start the Tivoli Directory Integrator server in daemon mode.
2. Start AMC using the following command - *TDI_install_dir*\bin\amc\start_tdiadc.bat
3. Logon to the AMC console using the URL with the following syntax - <http://host:port/ibm/console> using the default username and password.
4. After logging on to the AMC console, select the Servers on the navigation panel and then choose the Server you wish to use. Press the **Config files** button and the following panel is shown.



Select the SampleConfig.xml file and click the **Load** button.

5. Select the Solution View link of the navigation panel to create a solution view for the loaded config. Select the Add button and the following panel will be shown.

Add Solution View

Solution View Name
SampleSolutionView

Description:
|

Server:
Local Server ▼

Configs:
C__dev_200810100500_configs_SampleConfig.xml ▼ View Config Files

☒ **Simple** Add a Solution View with common default options

☐ Auto Update

☐ Add Solution View from published solution.

☒ Add Solution View with all AssemblyLines exposed.

☐ Add Solution View with all AssemblyLines and all properties exposed.

☐ Add Solution View with all AssemblyLines and all User properties exposed.

☐ **Advanced** Configure the properties to create a Solution View.

OK Cancel

Add a suitable name for the config, for example, SampleSolutionView. On selecting OK, a message indicating that the Solution View SampleSolutionView was created successfully will be shown.

6. Select the Action Manager link on the navigation panel to reach the Action Manager rules configuration screen. Select "SampleConfigView" from the Select solution views dropdown box. When clicking the **Add** button in the Configured Rules section, the following panel will be shown.

Add Rule - SampleSolutionView

Name:

Description:

Trigger type:

Configure trigger - On AssemblyLine termination

AssemblyLine:

Configured Actions:

Select	Action
<input type="checkbox"/>	Execute On Error

OK Cancel

Enter a name, for example "rule 1". Select trigger type "On AssemblyLine termination" and click the **Add** button in the Configured Actions section.

Select Action

Select action type:

- * Start AssemblyLine
- Stop AssemblyLine
- Enable/Disable rule
- Execute rule
- Notify event
- Modify property
- Copy property value
- Write to log

Next >> Cancel

In the above panel, select the "Write to log" option and click **Next**.

Add Action - Write To Log

Severity:
Severe ▼

Message:
*Data copied to out.xml

Description:
[Empty text box]

OK Cancel

Add the text in the **Message** text box and click **OK**. The following rule panel will be shown.

Add Rule - SampleSolutionView

Name:
*rule1

Description:
[Empty text box]

Trigger type:
On AssemblyLine termination ▼

Configure trigger - On AssemblyLine termination

AssemblyLine:
read ▼

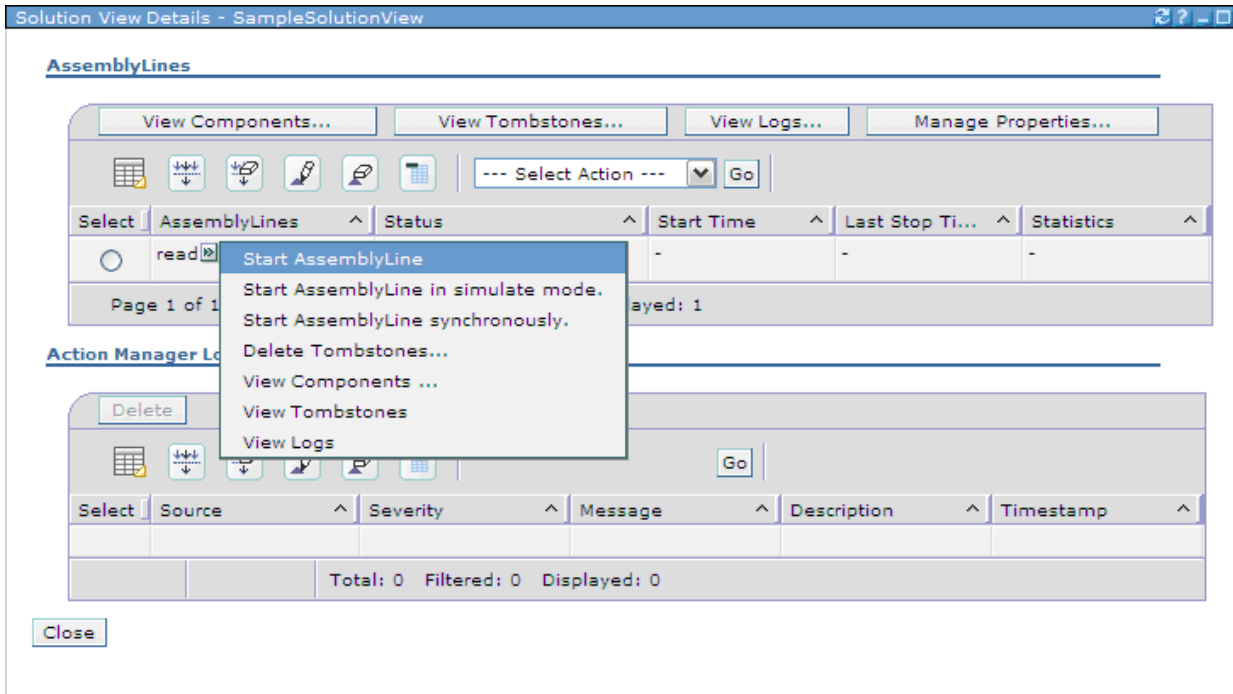
Configured Actions:

Select	Action	Execute On Error
<input type="radio"/>	Write to Log (Data copied to out.xml)	<input type="checkbox"/> Execute if previous actions fails

OK Cancel

Clicking **OK** completes the creation of this rule and the AMC configuration required for this example.

7. Start Action Manager using `TDI_install_dir\bin\amc\startAM.bat`. A thread is created for the rule "rule1". which should be waiting for the termination of the specified AL.
8. To trigger this rule, the "read" AssemblyLine of SampleConfig.xml needs to be executed. Select **Monitor Status** on the navigation panel. On the Monitor Status panel, select the SampleSolutionView and click on the **Solution View Details** button. The following panel will be shown.



9. Start the AL using the option shown above. The rule will be triggered and the following status will be displayed on the Action Manager console.

```

C:\WINDOWS\system32\cmd.exe - startAM.bat
2008-10-16 17:05:01 com.ibm.di.amc.actionmanager.ServerConfigHandler startAMCServerModificationListe
nerThread
INFO: CTGDB659I Started AMCServerModificationListener thread.
2008-10-16 17:05:01 com.ibm.di.amc.actionmanager.api.AMService initialize
INFO: AM.RMI.REGISTRY.SYSTEM.SECURITYMANAGER.NULL
2008-10-16 17:05:01 com.ibm.di.amc.actionmanager.api.AMService initialize
INFO: AM.RMI.REGISTRY.RESETTING.SYSTEM.SECURITYMANAGER.NULL
2008-10-16 17:05:01 com.ibm.di.amc.actionmanager.api.AMService initialize
INFO: CTGDB720I The Action Manager RMI registry is started on port 13104.
2008-10-16 17:05:01 com.ibm.di.amc.actionmanager.AMHandler startHealthALManagerThread
INFO: CTGDB512I Started Health AssemblyLine Manager.
2008-10-16 17:05:01 com.ibm.di.amc.actionmanager.AMHandler main
INFO: CTGDB511I Action Manager initialization completed.
2008-10-16 17:05:01 com.ibm.di.amc.actionmanager.AMHandler startDatabaseModificationListener
INFO: CTGDB532I Database modification listener started.
2008-10-16 17:06:23 com.ibm.di.amc.actionmanager.AssemblyLineTerminateListener triggerRule
INFO: CTGDB549I Rule rule1 triggered.
2008-10-16 17:06:24 com.ibm.di.amc.actionmanager.RuleExecutionManager writeToLog
SEVERE: CTGDB622I [Message]: Data copied to out.xml
[Description]: None
2008-10-16 17:06:24 com.ibm.di.amc.actionmanager.RuleExecutionManager execute
INFO: CTGDB615I Action successfully executed.
[Solution View]: SampleSolutionView
[Rule]: rule1
[Action Order]: 2
[Action Type]: WRITE_LOG

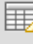





```


On AMC, the Action Manager logs table on the Solution View Details panel will be shown as below.

Solution View Details - SampleSolutionView

AssemblyLines

View Components... View Tombstones... View Logs... Manage Properties...

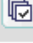
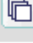
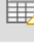

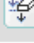

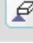
      --- Select Action --- Go


Select	AssemblyLines	Status	Start Time	Last Stop Ti...	Statistics
<input checked="" type="checkbox"/>	read	 Stopped	-	-	-

Page 1 of 1 Total: 1 Filtered: 1 Displayed: 1

Action Manager Logs

Delete

       --- Select Action --- Go

Select	Source	Severity	Message	Description	Timestamp
<input type="checkbox"/>	rule1	 SEVERE	Data copied to out.xml	None	16.10.2008 17:06:24

Page 1 of 1 Total: 1 Filtered: 1 Displayed: 1

Close

Chapter 17. Tombstone Manager

Introduction

Tivoli Directory Integrator 7.0 can keep track of configurations or AssemblyLines that have terminated. This way, you can tell when your AssemblyLines last ran, without going into the log of each one.

This is accomplished by TDI's "Tombstone Manager" that creates "tombstones" for each AssemblyLine and configuration as they terminate. Tombstones contain exit status and other information that you can request through the Server API. Tombstone Manager also:

- Displays the status of an entire Tivoli Directory Integrator configuration in an AMC status window.
- Ensures repeated runs of AssemblyLines within Action Manager, for example, every 24 hours.
- Provides status information to Server API clients about AssemblyLines that run asynchronously.

The Tombstone Manager API is documented in the Java API documentation; look for class `com.ibm.di.api.Tombstone`.

Configuring Tombstones

The creation of Tombstones for AssemblyLines and Config Instances is configured by means of check boxes in a number of screens in the Configuration Editor (CE), as well as a number of options in the `global.properties` or `solution.properties` files.

Once configured, your Config contain the following switches:

At the configuration level:

- Config switch: specifies whether tombstones are created or not for the Config Instance itself.
- All AssemblyLines switch: specifies whether tombstones are created for all AssemblyLines from this configuration.

At the AssemblyLine level:

A switch that specifies whether tombstones are created for that particular AssemblyLine. This switch is only taken into account when the "All AssemblyLines switch" at the configuration level is switched off.

Configuration Editor Configuration screen

You can configure tombstones creation for an AssemblyLine using the following AssemblyLine Configuration window. The Create Tombstone option is near the bottom of the window.

AssemblyLine settings

AssemblyLine settings

AssemblyLine Settings

Load task parameters from

Save task parameters to

Log statistics interval

Max number of reads (Iterator)

Max number of errors

Max entries returned by lookup

Include All Global Prologs ☒

Include Additional Prologs

Automatically map all attributes ☒

Define ALPool Options

Null Value Behavior

Detailed Log ☐

Create Tombstones ☒

Checking **Create Tombstones** cause tombstones to be generated for this AssemblyLine when runs, even when the master switch for AssemblyLines is disabled.

AssemblyLine Configuration screen

When AssemblyLine Tombstones are disabled using the Configuration option shown above, tombstone generation can still be enabled individually per AssemblyLine by using the appropriate option in the AssemblyLine configuration screen, **Create Tombstones**.

A sample tombstone record could look like this:

Table 26. Tombstone record

Field Name	Value
Component Type ID	1
Event Type ID	0
StartTime	11.11.2005 11:11:54
TombstoneCreateTime	11.11.2005 17:22:45
Component Name	"ActiveDirectoryChangeLogSynchronizer"
Configuration	"C:\TDI_SOL_DIR\rs.xml"
Exit Code	0
Error Description	""
GUID	"432640786324026346432"
Statistics	[get:571] [add:571] [err:0]

The statistics returned can be one or more of the following attributes:

Table 27. Statistics returned in a Tombstone record

Attribute	Description
add	Total number of entries the AssemblyLine has added (performed by Connectors in AddOnly mode)
mod	Total number of entries the AssemblyLine has modified (performed by Connectors in Update mode)
del	Total number of entries the AssemblyLine has deleted. (performed by Connectors in Delete mode)
get	Total number of entries the AssemblyLine has retrieved (performed by Connectors in Iterator mode)
request	Total number of requests accepted when there is a Server mode Connector in the AssemblyLine
callReply	Total number of Call/Reply operations the AssemblyLine has executed (performed by Connectors in CallReply mode)
err	Total number of errors encountered
skip	Total number of entries the AssemblyLine has skipped entries
lookup	Total number of Lookup operations the AssemblyLine has executed (performed by Connectors in Update/Delete/Lookup mode)
ignore	Total number of entries the AssemblyLine has ignored (performed by Connectors in Update/Delta mode)
reconnect	Total number of times the AssemblyLine has attempted to reconnect to another client
exception	The exception text if the AssemblyLine terminated with an exception
getTries	Total number of times the AssemblyLine has attempted to retrieve an entry (performed by Connectors in Iterator mode)

Table 27. Statistics returned in a Tombstone record (continued)

Attribute	Description
getClientTries	Total number of times the AssemblyLine has attempted to get the next connected client (performed by Connectors in Server mode)
nochange	Total number of entries the AssemblyLine processed but left unchanged
branchtrue	Total number of Branch components executed by the AssemblyLine because their expression evaluated to true
branchfalse	Total number of Branch components skipped by the AssemblyLine because their expression evaluated to false
loopstart	Total number of Loop components executed by the AssemblyLine
loopcycles	Total number of cycles executed for all Loop components that had more than one cycle in an AssemblyLine
reconnectTime	Time in ms after last reconnect was attempted by the AssemblyLine

The Tombstone Manager

The Tombstone Manager monitor the number of tombstone records at runtime and delete old records as per the values of the `com.ibm.di.tm.autodel.age`, `com.ibm.di.tm.autodel.records.trigger.on`, `com.ibm.di.tm.autodel.records.max` configuration properties (see “Tombstone Manager”).

- The Tombstone Manager tracks Config Instances and AssemblyLines stop events.
- The Tombstone Manager uses the local Server API calls for registering for event notifications and receiving stop events for Config Instances and AssemblyLines.
- The Tombstone Manager uses the TDI System Store for data persistence.
- The Server API (documented in the Java API documentation) interfaces contain calls for querying the Tombstone Manager for various data, like AssemblyLine and Config Instance Tombstones.
- The Tombstone Manager provides options for deleting old tombstone records.

A possible AssemblyLine tombstone lifecycle could look like:

- The Tombstone Manager receives a Server API event that an AssemblyLine has terminated (this assumes the Server API and the Tombstone Manager are turned on and the configuration file specifies that tombstones are created for this AssemblyLine).
- The Tombstone Manager extracts the required data from the Server API event and creates a corresponding database tombstone record in the System Store.
- While the tombstone record is available in the System Store, queries can be executed through the Server API calls that provide all the information contained in the tombstone record.
- The tombstone record is deleted from the System Store either when an explicit cleanup Server API call is executed that deletes it, or when the logic for automatic deletion of old tombstone records collects it. Neither of these events is required, so theoretically, the tombstone record may live forever.

Tombstone Manager

The Tombstone Manager task is configured by means of properties in the `global.properties` or `solution.properties` file for your Config instance.

Note: In order for the Tombstone Manager to function, the Server API must be switched on; that is, the property `api.on` must be set to **true**.

The relevant properties are:

`com.ibm.di.tm.on`

Master switch for the Tombstone Manager. Values are **on** and **off** - if set to off, no Tombstones are generated even if specified in the Config file; neither are they managed (nor can they be queried using the Server API, or AMC).

The default value for this property is **false**.

com.ibm.di.tm.autodel.age

The number of days a Tombstone live. When this property is present and contains an integer value greater than 0 the Tombstone Manager automatically delete all tombstone records that are older than the specified number of days.

The logic for tombstone record deletion is triggered on TDI Server startup and once a day on a long running TDI Server.

The default value for this property is **0**.

com.ibm.di.tm.autodel.records.trigger.on

Specifies the total number of tombstone records that trigger the logic for trimming the number of tombstone records to a certain number.

The default value for this property is **10000**.

com.ibm.di.tm.autodel.records.max

The number of Tombstones to be retained once the trigger specified by the previous parameter, `com.ibm.di.tm.autodel.records.trigger.on` is exceeded.

The default value for this property is **5000**.

com.ibm.di.tm.create.all

This property acts as an override switch for the values specified in the Config files. When this property is set to **true**, Tombstone Manager create Tombstones for every AssemblyLine and Config Instance regardless of the values specified in the configurations. This is useful to turn on Tombstone creation for pre-6.1 configurations that do not have tombstone values without modifying the configurations.

The automatic cleanup logic determined by the `com.ibm.di.tm.autodel.age` property is independent of the automatic cleanup logic determined by the `com.ibm.di.tm.autodel.records.trigger.on` and `com.ibm.di.tm.autodel.records.max` properties.

The Tombstone Manager uses the TDI logging framework and logs its messages in the TDI Server main log.

An example section in the `global.properties` or `solution.properties` file could look like:

```
com.ibm.di.tm.on=true
com.ibm.di.tm.autodel.age=90
com.ibm.di.tm.autodel.records.trigger.on=50000
com.ibm.di.tm.autodel.records.max=25000
com.ibm.di.tm.create.all=false
```

This set of configuration properties specifies that: The Tombstone Manager is turned on. Tombstones older than 90 days are automatically deleted. Also when the total number of tombstone records reaches 50000, the oldest 25000 tombstone records is automatically deleted.

Chapter 18. Multiple TDI services

IBM Tivoli Directory Integrator as Windows Service

Introduction

In IBM Tivoli Directory Integrator 7.0 there is a mechanism that allows multiple Tivoli Directory Integrator server instances to be registered as Windows services. Each instance requires a separate solution directory. After creating a solution directory, a utility program should be copied in it. The name of the program is `ibmdiservice.exe`. The configuration of the utility program and the Windows service is made with a properties file named `ibmdiservice.props`. Each solution directory should contain a configuration properties file.

Each Windows service must have a different name. A property called "servicename" in the property file specifies a name that is used in creation of the Windows service name and the Windows service display name. The Windows service name is formed by prefixing the value of the "servicename" property with the "ibmdisrv-" prefix. The Windows service display name is formed by inserting the value of the "servicename" property between the brackets of "IBM Tivoli Directory Integrator ()". For example if the "servicename" property is set to "test" the Windows service name is "ibmdisrv-test" and the Windows service display name is "IBM Tivoli Directory Integrator (test)". If the "servicename" property is not present or has no value default names are used. The default names for the Windows service name and the Windows service display name are "ibmdisrv" and "IBM Tivoli Directory Integrator".

A property exists so it can be configured whether the Windows service is started automatically on Windows startup or has to be started manually. The name of the property is "autostart" and the valid values for it are "true" and "false".

Note: This property is used during installation and uninstallation as well as while the service is running. That is why the property value must not be changed after the Windows service has been installed.

For more information about the TDI Windows service configuration properties file see the "Configuring the service" on page 238" section.

Installing and uninstalling the service

Installing the service

Do the following to install the IBM Tivoli Directory Integrator service:

1. Make sure the IBM Tivoli Directory Integrator is installed. The installation folder of the IBM Tivoli Directory Integrator is referred to as `root_directory`. See 11.
2. Choose a solution folder that is used by IBM Tivoli Directory Integrator when it is started as a Windows service - this can be any folder of your choice. Once IBM Tivoli Directory Integrator is installed as a service the solution folder used by the service cannot be changed until it is uninstalled as a service. Note that choosing the solution folder for the Windows service does not prevent from running IBM Tivoli Directory Integrator with any other solution folder.
3. Once the solution folder is chosen copy into that folder all files from the `root_directory/win32_service` folder: these are "ibmdiservice.exe", "ibmdiservice.props" and "Log4J.properties".
4. Execute the following command from the solution folder chosen for the Windows Service:
`ibmdiservice.exe -i`

Uninstalling the service

Note: In order to use the Tivoli Directory Integrator 7.0 version of the "ibmdiservice.exe" utility program any registered pre-Tivoli Directory Integrator 7.0 Windows service must be uninstalled and then the Tivoli Directory Integrator 7.0 windows service must be installed. This is necessary because the Tivoli Directory Integrator 7.0 windows service uses a different default name for the Windows service name – "ibmdisrv" as opposed to the pre-Tivoli Directory Integrator 7.0 default name of "IBM Tivoli Directory Integrator".

Do the following to uninstall the IBM Tivoli Directory Integrator service:

1. Make sure the IBM Tivoli Directory Integrator service is stopped.
2. Execute the following command from the solution folder chosen when installing the service:

```
ibmdiservice.exe -u
```

Notes:

1. Uninstalling the IBM Tivoli Directory Integrator service does not uninstall the IBM Tivoli Directory Integrator itself. You can still use the IBM Tivoli Directory Integrator but it is not registered and run as a Windows service. You can install IBM Tivoli Directory Integrator service again later.
2. If the IBM Tivoli Directory Integrator service is installed and you wish to completely uninstall the IBM Tivoli Directory Integrator (not just the service), do the following:
 - a. Uninstall the Windows service.
 - b. Uninstall the IBM Tivoli Directory Integrator (Refer to 30).

Starting and stopping the service

The IBM Tivoli Directory Integrator service automatically starts the IBM Tivoli Directory Integrator at system boot. The IBM Tivoli Directory Integrator is not, however, automatically started when the service is installed. After installing the service you have two options to start the service:

- Restart the computer.
- Start the IBM Tivoli Directory Integrator service from the Windows Services window.

Manual start and stop

You can manually start and stop the IBM Tivoli Directory Integrator service from the Windows Services window.

In the **Services** window you must select the service IBM Tivoli Directory Integrator and, depending on the Windows version, either click the **Start/Stop** button, or right-click on the service name and select **Start/Stop**.

Changing service startup type

By default, the IBM Tivoli Directory Integrator service is configured to start automatically on system boot.

You can manually change the service startup mode from the Windows Services window to **Manual** or **Disabled**.

Logging

The IBM Tivoli Directory Integrator service logs all messages (**error**, **info** and **debug**) in the Application Windows system log. You can view these messages with the Windows Event Viewer.

Configuring the service

The IBM Tivoli Directory Integrator service is configured through the `ibmdiservice.props` file placed in the solution folder chosen during installation of the log service.

Note: Before running the service, make sure this file is properly configured as described in this section. The service could fail if the file contains incorrect values.

The following properties are specified in the `ibmdiservice.props` file:

path Specifies the PATH environment variable used for running the IBM Tivoli Directory Integrator process (this property is usually the same as the PATH variable from `ibmdisrv.bat`, but you can change it). This is an optional property.

ibmdirroot

Specifies the root folder of the IBM Tivoli Directory Integrator (for example, `C:\Program Files\IBM\IBMDirectoryIntegrator`). This is a required property.

configfile

Specifies the file path to the IBM Tivoli Directory Integrator configuration file. This is an optional property.

assemblylines

Specifies in a comma-delimited format the AssemblyLines that are started automatically when the IBM Tivoli Directory Integrator service is started. This is an optional property.

cmdoptions

Specifies other options that are directly passed to the IBM Tivoli Directory Integrator on service startup (see "IBM Tivoli Directory Integrator options" in *IBM Tivoli Directory Integrator V7.0 Users Guide* for the full list of IBM Tivoli Directory Integrator options).

One such option could be the `-c` option; here you could specify multiple config files (separated by commas), something which is not allowed by the **configfile** parameter. This is an optional property.

jvmcmdoptions

Specifies options to the Java Runtime Environment executable, used to deploy the Tivoli Directory Integrator Server that runs the Windows service. This is an optional property.

servicename

Specifies a name that is used to form the Windows service name and the Windows service display name. The windows service name is set to the value of the **servicename** property prefixed with the "ibmdisrv-" prefix. The windows service display name is created by inserting the value of the **servicename** property between the brackets of the "IBM Tivoli Directory Integrator ()" expression.

For example, if the property value is "test" the Windows service name will be "ibmdisrv-test" and the Windows service display name will be "IBM Tivoli Directory Integrator (test)". If the **servicename** property is not present or has no value, default names are used. The default Windows service name is "ibmdisrv" and the default Windows service display name is "IBM Tivoli Directory Integrator".

Note: This property is used during installation and uninstallation as well as while the service is running. That is why the property value must not be changed after the Windows service has been installed.

autostart

Specifies whether the Windows service starts automatically on Windows start-up or whether it has to be started manually. The valid values for this property are **true** and **false**. A value of **true** specifies that the Windows service is started on Windows start-up and a value of **false** specifies that the service has to be started manually. If this property is not present or has no value, then the default value of **true** is used.

This property is used during Windows service installation and changing it after the Windows service has been installed has no effect.

debug Specifies **true** or **false** to correspondingly turn debug information on or off. When debug information is turned on, detailed trace messages are dumped in the Application Windows system log. This is an optional property.

Note: When specifying properties in the configuration file, specify each property on a single line and use the following format:

<property_name>=<property_value>

There must be no spaces around the equals (=) sign.

An example of a completed `ibmdiservice.props` file looks like the following:

```
path=C:\Program Files\IBM\TDI\V7.0\jvm\jre\bin;  
C:\Program Files\IBM\TDI\V7.0\libs;  
ibmdiroot=C:\Program Files\IBM\TDI\V7.0  
configfile=rs.xml  
assemblylines=AssemblyLine1,AssemblyLine2  
cmdoptions=-d  
debug=false
```

Note: If you change any of the properties in `ibmdiservice.props`, you must restart the service for the changes to take effect.

IBM Tivoli Directory Integrator as Linux/UNIX Service

Deployment methods

On Linux and UNIX platforms, there are two different ways of ensuring that certain system jobs or 'daemons' start and stop at respectively system initiation and system termination:

1. Using a script in `/etc/init.d` containing the logic to start and stop the daemons you are interested in. This script you then (hard)link to scripts in `/etc/rc3.d`: their names beginning with `SXX...` and `KXX...` - the `XX` being a numeral which causing the files to show up in the right sequence in the `/etc/rc3.d` directory. The scripts starting with `S` are called when the system reaches run phase 3 at system startup, and the scripts starting with `K` are called when the system terminates.
2. By editing the `/etc/inittab` file.

The latter process is what we describe here. Some of the information presented could be used to construct scripts utilizing the first deployment method.

Tailoring `/etc/inittab`

In order to start up TDI daemon processes when the UNIX/Linux OS starts appropriate entries must be added to the `/etc/inittab` file. The registering of TDI as a windows service on Windows translates to adding a line of text to the `/etc/inittab` file on UNIX/Linux. The un-installation of the TDI windows service on Windows translates to removing the corresponding entries from the `/etc/inittab` file. For each TDI daemon process that needs to be started on system startup one line of text must be added to the `/etc/inittab` file. The format and meaning of the entries in this file is described below. Each entry in the `/etc/inittab` file has the following format:

Identifier:RunLevel:Action:Command

A description of each of these fields is as follows:

- The **Identifier** field is a string (at least a single character in length) that uniquely identifies an object. This string is used to uniquely identify the corresponding command.
- The **RunLevel** field is the run-level in which this entry can be processed. Run-levels effectively correspond to a configuration of processes in the system. Each process started by the `init` command is assigned one or more run-levels in which it can exist. A run-level is represented by the numbers 0

through N, where N is a positive integer different for the different UNIX/Linux operating Systems (for example on some AIX computers N is 9, on RedHat Linux N is 6, and so on.). If the OS is running in run-level 3, for example, then only processes specified for run-level 3 are started.

The **RunLevel** field can define multiple run-levels for a process by selecting more than one run-level in any combination from 0 through N. For example, if Tivoli Directory Integrator needs to run in run-level 3 and 6, then the run-level must be specified as "36". If no run-level is specified, the process is assumed to be valid at all run-levels.

It is recommended that no run-level numbers are specified, unless the specific TDI solution specifically needs to.

- The **Action** field is a value from a set of predefined actions which tells the **init** command how to treat the process specified in the **Command** field. There are many actions recognized by the init command, but for running the TDI server as a daemon process it is recommended that the **once** action be used. The semantics of the **once** action are:

When the init command enters a run-level that matches the entry's run level, start the process, and do not wait for its termination. When it dies, do not restart the process. When the system enters a new run level, and the process is still running from a previous run level change, the program not be restarted. All subsequent reads of the `/etc/inittab` file while the init command is in the same run level cause the init command to ignore this entry.

- The **Command** field specifies the shell command to run.

Here are three example TDI-related entries in `/etc/inittab`:

```
tdi1::once:/opt/IBM/TDI611_1/ibmdisrv -c "/opt/IBM/TDI611_1/myconfigs/rs1.xml" -r "testAL1"
tdi2::once:/opt/IBM/TDI611_2/ibmdisrv -c "/opt/IBM/TDI611_2/myconfigs/rs2.xml" -r "testAL2"
tdi3::once:/opt/IBM/TDI611_3/ibmdisrv -c "/opt/IBM/TDI611_3/myconfigs/rs3.xml" -r "testAL3"
```

This example starts three TDI server instances which are installed in different folders.

Note: There are some differences in the different UNIX/Linux operating systems for system startup. That is why the information provided here covers the main issues of starting IBM Tivoli Directory Integrator on a UNIX/Linux system and does not refer to any specific UNIX/Linux system.

As an example of an `/etc/inittab` file, detailed information about the `/etc/inittab` configuration file for an AIX system can be found at http://publib16.boulder.ibm.com/pseries/en_US/files/aixfiles/inittab.htm

IBM Tivoli Directory Integrator as z/OS Service

USS process

In order to start an instance of the TDI server when the Unix System Service of z/OS is started, you must have the HFS data set for Tivoli Directory Integrator mounted before the Unix Service is initialized.

This could be done by adding a record for mounting the HFS data set to the BPXPRM01 data member of the USER.PARMLIB data set.

For example the record for mounting the HFS data set might look like this:

```
MOUNT FILESYSTEM('TDI.V6R1M0.HFS') TYPE(HFS) MODE(READ) MOUNTPOINT('/usr/lpp/itdi')
```

After mounting the HFS data set permanently, you should add a record in the `/etc/rc` file that starts an instance of the Tivoli Directory Integrator Server. The record in the `/etc/rc` file must specify the solution directory.

For example:

```
/usr/lpp/itdi/ibmdisrv -s /u/musala/tdi_solutions -c rs.xml -r al
```

In order for the changes in the USER.PARMLIB(BPXPRM01) data member and the /etc/rc file to take effect, the z/OS computer needs to be re-IPLeD (rebooted).

Normal z/OS started task

Tivoli Directory Integrator ships with an example in order to illustrate how a Tivoli Directory Integrator instance can be started as a normal z/OS task. It consists two parts: shell script and JCL.

Shell script: iditask.sh:

The shell script iditask.sh can be found in the <TDI_install>/bin folder (by default usr/lpp/itdi/bin). It contains the actual shell command for starting a TDI instance as well as exports several environment variables, which are needed by the further processing of the script by the z/OS native program.

The script sets the JAVA_HOME environment variable to the path of the Java 5.0 JVM, which is the required version to run the TDI server, and exports the LIBPATH variable by adding the TDI installation directory to it.

In order to run any child processes in the same address space as the started task the environment variable _BPX_SHAREAS is set to **YES** in the script.

The script implements logic to distinguish the actual installation directory of TDI under which it is visible in USS. As distributed, the server is started in daemon mode with the same solution directory as the installation directory. When TDI is installed on the default location the command is interpreted as follows:

```
/usr/lpp/itdi/ibmdisrv -s /usr/lpp/itdi -d
```

In case you want to change the default behavior, for example by changing the solution directory, starting Assembly line instead of daemon mode and so forth, this script should be edited correspondingly before the z/OS task is started.

JCL: ITDITASK:

The second part of the mechanism is the ITDITASK JCL, which invokes the shell script described above and starts Tivoli Directory Integrator as a normal z/OS task in its own address space.

It resides in MVS in the adi.HADI700.F1 library, where adi stays for the high level qualifier of the product chosen by the user (usually ADI or TDI), as well as in the corresponding AADIJCL and SADIJCL libraries (for example, TDI.V7R0M0.AADIJCL). From a legacy perspective it is better to first copy it into an user's CNTL and submit from there.

The ITDITASK uses the BPXBATCH native utility to start the shell script.

Note that it specifies the predefined location of "iditask.sh" - /usr/lpp/itdi/bin/iditask.sh. In case TDI is installed in a location different from the default (/usr/lpp/itdi) or the script is moved to another location after the installation ITDITASK JCL must be edited in order to function.

The ITDITASK redirects STDOUT, STDERR and LOGOUT to the standard output of the task (SYSOUT). Thus every message sent to the console is logged under the task's name in SDSF. "Configuring the TDI task to log to its SYSOUT" on page 249 The behavior can be changed, so that the relevant messages are stored in an intermediate USS file or another MVS dataset.

Below is an example as to how STDOUT can be redirected from SYSOUT to the USS file /usr/lpp/itdi/stdout.txt. In the ITDITASK just replace the statement:

```
//STDOUT DD SYSOUT=*
```

with the following ones:

```
//STDOUT DD PATH='/usr/lpp/itdi/stdout.txt',  
// PATHOPTS=(OWRONLY,OCREAT,OTRUNC),  
// PATHMODE=SIRWXU
```

TDI autostarting:

There is an additional step in order to automatically start the task with the start of the system. For this purpose the COMMNDxx member in the SYS1.PARMLIB dataset, which is a mechanism of having console start commands invoked during the bring-up of a system image, must be edited. By having a "started task" named for example ITDITASK defined as an ITDITASK member in the SYS1.PROCLIB dataset, the user can get that task started during system startup by having a CMD='S ITDITASK' in the COMMNDxx member.

An alternative is to add the task to USER.PROCLIB and editing the COMMNDxx member of USER.PARMLIB or to use another library concatenated to the Job Entry Subsystem's (JES2 or JES3) PROCLIB allocation in case the user does not have write access to SYS1 members.

By this means the ITDITASK can be started from the SDSF menu with the standard z/OS START command. For example:

```
/S ITDITASK
```

IBM Tivoli Directory Integrator as i5/OS Service

Tivoli Directory Integrator can be started as a service on i5/OS (OS400) by defining it as a custom TCP/IP server. The server can then either be auto started when TCP/IP is started or manually started via the **STRTCPSVR** command. To define a TCP/IP server, the **ADDTCPSVR** (Add TCP Server) command is used which requires a program to be registered that is called by the **STR/ENDTCPSVR** command.

The following are the steps to manually start a Tivoli Directory Integrator service which start an AssemblyLine "AL" in Config file rs.xml .

1. Create Library for Tivoli Directory Integrator:
CRTLIB (TDILIB) TYPE(*TEST) TEXT('TDI Library')
2. Create a Source physical file in the TDILIB library:
CRTSRCPF FILE(TDILIB/TDISRC)
3. Add a file member to the Source file and add the following code using the Source Entry Utility (SEU):

```
PGM PARM(&P) /* START TDI Service Program */  
DCL VAR(&P) TYPE(*CHAR) LEN(10)  
SBMJOB CMD(STRQSH CMD(' <TDI_INSTALLATION >/ibmdisrv -c +  
    <TDI_INSTALLATION> /configs/rs.xml -r AL ')) +  
    JOB(TDIService) ALWMLTTHD(*YES)  
ENDPGM
```

The program is saved as **TDIPGM** in **TDILIB**.

4. Create a program object from the file member in the TDILIB library. (Use **CRTBNDCL** cmd or Option 3 in Programmer menu)
5. Add the Tivoli Directory Integrator TCP/IP server using the **ADDTCPSVR** command with name **TDISERVER**:

```
ADDTCPSVR SVRSPCVL(*TDISERVER)  
          PGM(TDILIB/TDIPGM)  
          SVRNAME(TDISERVER)  
          SVRTYPE(TDISERVER) AUTOSTART(*NO)
```

6. Start the TDI TCP/IP server:
STRTCPSVR SERVER(*TDISERVER)

For more information on the i5/OS (OS400) (V5R4) and commands that have been demonstrated above, see: <http://publib.boulder.ibm.com/infocenter/iseres/v5r4/index.jsp?topic=/clfinder/finder.htm>

Chapter 19. z/OS environment Support

The full set of IBM Tivoli Directory Integrator components is available in the native z/OS environment (version v1.7 and v1.8), with IBM JVM 1.5, except for those components that require native libraries, like the Windows Users and Groups Connector and the Domino Change Detection Connector. Conversely, the z/OS Command Line Connector is available on z/OS® only.

You can start an instance of the Server by executing the startup shell script `usr/lpp/itdi/ibmdisrv` residing under the Unix System Services. Also see “IBM Tivoli Directory Integrator as z/OS Service” on page 241 for information about automatic startup of a Tivoli Directory Integrator Server, either as a USS process or a normal z/OS started task.

The user whose identity the Tivoli Directory Integrator Server runs under needs an OMVS segment definition in his profile specifying that at least 200MB operating memory is available.

The z/OS TSO Command Line Function Component is of particular relevance for the z/OS environment. It is able to run privileged z/OS TSO Commands. This component addresses the need to manage RACF®, TopSecret and ACF2 users – this can be achieved by executing TSO commands.

The Configuration Editor and AMC, however, are not supported natively on z/OS; instead, you should use remote management options, like

- The Remote Configuration Editor. Run the Configuration Editor on a supported platform, and access Config files on z/OS using a configured Config Instance on z/OS.
- The Administration and Monitoring Console
- Applications that use the remote Tivoli Directory Integrator Server API.

For installation instructions and hardware requirements on z/OS, see the Program Directories physically shipped with the product:

- Program Directory for IBM Tivoli Directory Integrator General Purpose Edition (GI11-8734-00: TDI_GPE.PDF)
- Program Directory for IBM Tivoli Directory Integrator Identity Edition (GI11-8733-00: TDI_IE.PDF)

Post install configuration

Using MQe for system queue

In order to use MQe as a system queue which is turned on by default you can invoke:

```
TDI_install_dir\jars\plugins\mqeconfig.bat(.sh) TDI_install_dir\jars\plugins\mqeconfig.props create server
```

If you do not want to use the system queue, set the `systemqueue.on` property of the `global.properties` or `solution.properties` file to false.

Default encoding different than IBM-1047

The text files as installed by SMP/E are encoded using the IBM-1047 encoding. If the default character encoding on your system is different from IBM-1047 you need to perform the following post-installation steps before the IBM Tivoli Directory Integrator installation directory is made read-only in order to convert the text files to your default encoding:

1. Change the current working directory to `"/usr/lpp/itdi/tools"`.
2. Issue the following command:

```
iconv -f IBM-1047 -t YOUR_DEFAULT_ENCODINGenccnvz > ../enccnvz
```

3. The previous command should have created a shell script called "enconvz" in the "/usr/lpp/itdi" folder
4. Run the following command to change the mode of the file:
`chmod u+x enconvz`
5. Run the enconvz script.

JDK 5.0 not located at /usr/lpp/java/J5.0

Ensure that either the JRE_PATH or the JAVA_HOME environment variable is set accordingly. For example: If your Java SDK is installed in "/usr/lpp/java/MyJava5.0", then

```
JAVA_HOME=/usr/lpp/java/MyJava5.0
```

or

```
JRE_PATH=/usr/lpp/java/MyJava5.0/bin
```

Running Tivoli Directory Integrator

Since the install directory of Tivoli Directory Integrator (/usr/lpp/itdi) is read-only, Tivoli Directory Integrator must be run from a solution directory (as explained in the documentation) which is different from the install directory. That is why before running Tivoli Directory Integrator, you should create a solution directory (which is just a normal directory different from the install directory), make this solution directory the current working directory and then launch the Tivoli Directory Integrator server. You can follow these steps in order to accomplish the creation of a solution directory and launching the Tivoli Directory Integrator server:

1. `mkdir solution_dir`
2. `cd solution_dir`
3. `/usr/lpp/itdi/ibmdisrv TDI_PARAMETERS`

where *TDI_PARAMETERS* are just the normal parameters the Tivoli Directory Integrator server reads from the command line. The first time the server runs this way, it will populate the solution directory with a number of configuration files based on those in the installation directory; you can now customize the files in the solution directory for your own needs.

Reading License Files

All license files are UTF-8-encoded so that NLS characters are preserved. Most OS tools can only read files encoded using the default (native) character encoding. That is why before you can read a license file using a general purpose tool, you need to convert its encoding. You can do that using the `iconv` utility in the following way:

```
iconv -f UTF-8 -t YourDefaultCharacterEncoding LicenseFile DestinationFile
```

where *YourDefaultCharacterEncoding* is the native character encoding of your z/OS system, *LicenseFile* is the license file you wish to convert and *DestinationFile* is the text file that will contain the contents of *LicenseFile*, but encoded using *YourDefaultCharacterEncoding*.

After this command completes you should be able to open and read *DestinationFile* (provided you have specified the correct encoding and that encoding supports all characters in *LicenseFile*).

Another option for reading the license files is to transfer them in BINARY mode via FTP to a system which can read UTF-8-encoded text files and supports the NLS characters contained in the license file you wish to read.

Using the Remote Configuration Editor on z/OS

On z/OS the only way of editing and modifying Config files stored on z/OS is by using the Remote Configuration Editor. See “Using the Remote Configuration Editor” on page 104 for some general characteristics for this particular way of editing Config files. In addition to that, there are some additional considerations when using the Remote CE for z/OS:

1. Config files developed locally must be uploaded (using FTP) in binary mode to the z/OS computer.
2. The default encoding on z/OS is EBCDIC (or IBM-1047 as it is also known). It is very different from ASCII or from UTF-8. There are no common character ranges between it and ASCII/UTF-8 (whereas the range 0-127 is the same between ASCII and UTF-8).

This is why any text in EBCDIC viewed as ASCII/UTF-8 looks unintelligible and the other way round.

Any file created on a Windows or Unix computer (for example, a properties file, a text file, and so on.) that needs to be read by TDI running on a z/OS system must follow the native encoding format of z/OS. One way of converting a Windows file to the native z/OS encoding format is to use the `enconvz` utility shipped with TDI in the “tools” directory.

Here is an example (UNIX) usage of the `enconvz` command:

```
./enconvz myfile_win.txt  myfile_z/OS.txt  ISO-8859-1  IBM-1047
```

For more information, see “Handling configuration and properties files.”

3. The `tdisrvctl` CLI remote utility is installed in the `TDI_root_directory/bin` directory. See “Command Line Interface – `tdisrvctl` utility” on page 160.

Note that for z/OS, by default, the `tdisrvctl` utility is configured to create logs in the `TDI_root_directory/logs/tdisrvctl.log` file. Since this could be a read-only location, it is recommended that you edit the `TDI_root_directory/etc/tdisrvctl-Log4J.properties` file to point its log file to a writable location. The property to edit in the `tdisrvctl-log-4j.properties` file is: `Log4J.appender.Default.file`. Also, if the `-h` (hostname) option is skipped, the `tdisrvctl` utility takes `localhost` as the default. This may not work on z/OS for all cases. Always specify the `-h` option with the computer's IP address.

Handling configuration and properties files

Handling of configuration and properties files is important because of the specific default encoding used on z/OS (EBCDIC), which is not compatible with UTF-8 usually used on other platforms.

The TDI Server can read configuration files in any encoding that is supported by the JVM; TDI Configuration files are read with the encoding specified in the header of the XML file. If no encoding is specified in the header of the configuration file, **UTF-8** is used.

The TDI Server can write configuration files in any encoding that is supported by the JVM.

- If the `-n <encoding>` switch is used when starting the TDI Server the encoding specified by `<encoding>` is used for writing configuration files.
- If the `-n` switch is not specified and the system property `com.ibm.di.config.encoding` is non-null then the value of this property is used as encoding when writing configuration files.
- If neither of the `-n` switch nor the `com.ibm.di.config.encoding` system properties are specified, then **UTF-8** is used for writing configuration files.

In all cases the encoding used for saving the XML configuration file is written in the header of the XML file.

This strategy for reading and writing configurations assumes that it is usually the UTF-8 encoding that is used on z/OS for configuration files. If however you want to use a different encoding (for example the system default so that the configuration file can be opened by a text editor like vi) then you are provided with a mechanism that can be used to create and use configuration files with an arbitrary encoding.

You should pay attention on the encoding used whenever you operate with text files on the z/OS system. For example when a file is read with the FileSystem Connector the **Character Encoding** parameter of the Parser used should specify the encoding of the file or should be left empty when the file uses the default EBCDIC encoding.

All *.properties files (that is, global.properties, Log4J.properties, and so on) in the installation directory and/or Solution Directory are read with the system default encoding. This makes it convenient for you to open and edit the properties files directly on the z/OS system.

Using ASCII mode

The ibmdisrv startup script starts the Tivoli Directory Integrator server without altering its default encoding, which on z/OS is EBCDIC (IBM1047). In order to run the server on z/OS in ASCII mode you must start it using the `ibmdisrv_ascii` startup script. This script starts the server with its default encoding set to ASCII (ISO-8859-1).

Note: In ASCII mode, the server ignores the `global.properties` file. Only the `solution.properties` file in your Solutions Directory is used, and this file needs to be encoded in the ASCII character set.

Encoding of `solution.properties`:

Altering the default encoding of the Tivoli Directory Integrator server affects how the `solution.properties` file is read. That is why the encoding of the `solution.properties` file in your Solution Directory must be changed to ASCII before starting the Tivoli Directory Integrator server in ASCII mode. The location of the `solution.properties` file is *your_solution_directory/solution.properties*.

Changing the encoding of a text file on z/OS:

The standard `iconv` utility available on z/OS can be used to convert the encoding of a text file. Starting the `iconv` utility with no parameters on the z/OS prints usage information.

The `global.properties` file:

When the Server on z/OS is run in ASCII mode the `global.properties` file is ignored and only the `solution.properties` file in your Solution Directory is read. That is why you must have all the required properties for your solution in the `solution.properties` file in your Solution Directory.

Log files:

When the Server on z/OS is run in ASCII mode the server log files are encoded in ASCII when being written to the file system. That is why in order to read these ASCII log files you might have to first convert their encoding to the native encoding on z/OS, which is EBCDIC (IBM1047).

Console output:

When the Server on z/OS is run in ASCII mode any text output to the z/OS console by the server appears garbled. This is caused by the output text being encoded in ASCII while the console expects the text to be encoded in EBCDIC. In order to read the server output to the console, the server output can be redirected from the console to a file and then this file can be converted from ASCII encoding to native encoding on z/OS (EBCDIC).

Configuring the TDI task to log to its SYSOUT

The Tivoli Directory Integrator product implementation on z/OS provides the ability to redirect the information from the intermediate USS `ibmdi.log` file within the related HFS file to the TASK SYSOUT of the TDI started task.

Configuring the TDI task to log to its SYSOUT allows z/OS users to have all the relevant information in a single place; it also allows people not fully skilled on Unix System Services to find the logs using SDSF. In addition the full *sysout* (including product messages) is saved and kept using the normal tool. Using this approach ensures that each start of the TDI TASK stores the information in its own *sysout* thus preventing any log replacement.

The redirection of the TDI server logs consists of two configuration steps: editing the `log4j.properties` file and the JCL script starting TDI as z/OS task (ITDITASK; also see “IBM Tivoli Directory Integrator as z/OS Service” on page 241). This approach relies on the ability of newer releases of z/OS including the oldest supported by TDI - version 1.6 - to route STDOUT and STDERR to SYSOUT instead of an 'intermediate' USS file. For this purpose all the needed log information should be first made visible from within the standard console. In order to take advantage of this feature the two steps described below should be performed.

Modifying log4j.properties

This step is required in order to redirect the log messages to the standard console output. For this purpose the `log4j.properties` file used by the TDI instance must be edited to add a console appender and to use it as default logger (to ensure that all log messages are stored in the task SYSOUT) or specify it as a logger for certain TDI objects like ALs, configurations and the like, so that only relevant information is added.

For example, to add and configure a console appender as default logger:

```
# Here is an example on how to make a logger that logs to the console
log4j.appender.CONSOLE=org.apache.log4j.ConsoleAppender
log4j.appender.CONSOLE.layout=org.apache.log4j.PatternLayout
log4j.appender.CONSOLE.layout.ConversionPattern=%d [%t] %-5p - %m%n
log4j.appender.CONSOLE.encoding=IBM-1047
```

Changing the root logger to the CONSOLE appender:

```
log4j.rootCategory=INFO, CONSOLE
```

With this modification the log information is now routed to the standard console and can be further manipulated in the MVS environment.

Modifying the starting JCL

The second step involves modification of the JCL responsible for starting TDI as normal z/OS started task, so that the messages received from the console are routed to the SDSF visible logs of the task. Note that if the ITDITASK JCL shipped as an example with TDI is used, this step is redundant, since its default behavior is to log there.

Edit the JCL, which starts the TDI server as standard z/OS started task, to redirect the log output to the task SYSOUT. This could be done by simply adding the lines:

```
//STDOUT DD SYSOUT=*
//STDERR DD SYSOUT=*
```

or

```
//STDOUT DD SYSOUT=(,)
//STDERR DD SYSOUT=(,)
```

Since the z/OS system takes advantage of the Log4J options, it stores all the messages with status ERROR and FATAL in the STDERR associated with the TASK and the others - DEBUG, INFO,

and so on. in the STDOUT. The length of the created records in this manner is not fixed to 133 characters, therefore extremely long or multiline messages can be saved without truncation.

The information can be easily routed or copied to another dataset or intermediate USS file. By this means it can be stored both in MVS and USS environment.

Below is an example showing how the STDOUT can be redirected back to the `ibmdi.log` file from the JCL:

```
//STDOUT DD PATH='/usr/lpp/itdi/logs/ibmdi.log',  
// PATHOPTS=(OWRONLY,OCREAT,OTRUNC),  
// PATHMODE=SIRWXU
```

Note: By formatting the messages a non-readable character appears between some of the components of the message, for example, the Date and the message ID, which is due to a conversion issue when transferring characters from ASCII to EBCDIC. When using the characters "[" and "]" in the pattern, they are not translated properly in the z/OS native encoding. The "[" character is replaced by the "!" and "]" is converted to ".".

Appendix A. Dictionary of terms

IBM Tivoli Directory Integrator terms

Action Manager (AM)

Action Manager is a stand-alone Java application used to configure failure-response behavior for Tivoli Directory Integrator 7.0 solutions. AM executes *rules* defined with AMC v.3. An AM rule consists of one or more *triggers* that define a "failure" situation – such as the termination of an AL that should not stop running, or if an AL has not been executed within a given time period, and so forth. Furthermore, each rule also defines *actions* to be carried out in case of this "failure". Actions include operations like sending events or e-mail, starting ALs (locally or remotely) and changing configuration settings. Action Manager requires Tivoli Directory Integrator 7.0 and AMC v.3.

Accumulator

A special object that can be set in a Task Call Block (TCB) for use when starting another AssemblyLine either via a scripted call, or a component like the AssemblyLine Connector or the AssemblyLine FC. The Accumulator is either a collection of Work Entry objects handled by the called AL, or it is a component that is called to output each Entry. Accumulator handling is done at the end of each AssemblyLine Cycle.

AES Shorthand for Advanced Encryption Standard. AES is an encryption algorithm for transmitting sensitive (but unclassified) content by U.S. Government agencies.

Adapter

Adapter is a word used in many contexts and with different meanings. A *TDI Adapter* refers to an AssemblyLine that is "packaged" as a single Connector. Creating a TDI Adapter requires setting up an AssemblyLine that is written to perform (and expose) one or more business related tasks. Each task is defined as an AssemblyLine Operation (for example, 'EnableAccount', or 'ReturnGroupMembers'). This AL can then be *published* for sharing, and can be leveraged by the AssemblyLine Connector which offers mode settings reflecting these operations².

AL Shorthand for AssemblyLine.

Administration and Monitoring Console (AMC)

AMC is a browser-based console for managing and monitoring solutions. AMC Version 3, which is part of the Tivoli Directory Integrator 7.0 release, runs inside the Integrated Systems Console (ISC). Each AMC version is designed to work with a specific release of TDI and may be incompatible with other versions. AMC v.3 is designed for Tivoli Directory Integrator 7.0, however, it also works with TDI 6.1.X and TDI 6.0 (albeit with some restrictions). AMC v.2 works with TDI 6.0 and AMC v.1 runs with TDI 5.2.

API Application Program Interface. A way of programmatically (local or networked) calling another application, as opposed to using a command-line or a shell script.

Appender

Appender is a Log4J term (a third party Java library) for a module that directs log-messages to a certain device or repository. In IBM Tivoli Directory Integrator you control logging for your AssemblyLines by creating and configuring *Appenders*, either under the Logging tab of a specific AL, or under Config -> Logging in the Config Browser to control how all AssemblyLines in the Config do their logging.

AssemblyLine (AL)

The basic *unit-of-work* in a TDI solution. Each AL runs as a JVM thread in the Server and is made

2. AL Operations are also accessible via the AssemblyLine FC.

up of a series of AssemblyLine components (one or more Connectors, Functions, Scripts, Attribute Maps and Branches) linked together and driven by the built-in workflow of the AssemblyLine.

AssemblyLine Component

This term denotes an TDI component used to construct AssemblyLines. The possible Components are:

- Connectors
- Function Components
- Script Component
- Attribute Map Component
- Branches (including Loops and Switches)

The components list in an AssemblyLine is divided into two sections: *Feeds* where the Work Entry for each AL cycle is created from input data by a Connector in Iterator or Server mode, and the *Flow* section that holds the Connectors (in any mode except Server), Functions, Attribute Maps and Scripts providing the additional data access and processing.

AssemblyLine Operation

A business task that is implemented by an AssemblyLine and published via its Operations tab. Each Operation can have its own Input and Output Attributes Maps for defining the parameters expected when this Operation is invoked (Input Map), as well as those returned (Output Map). This is also called the *Schema* of the Operation.

AssemblyLine Phases

An AssemblyLine goes through three phases:

Initialization

At this point the TDI Server uses the "blueprint" for the AssemblyLine in the Config to create the various components as well as set up the AL environment, including processing the TCB, starting the AL's script engine and invoking the AssemblyLine's Prolog Hooks. All components that are configured for Initialization At Startup are initialized at this point causing their Prolog Hooks to get run as well.

Cycling

Now the AL workflow drives each of its components in turn, starting each cycle by invoking the On Start of Cycle Hook. Then the currently active *Feeds* Connector reads in data, creates the Work Entry and passes it to the *Flow* section. The Work Entry is passed from component to component until the end of *Flow* is reached, at which time control is returned to the start of the AssemblyLine again³. Cycling continues until an unhandled error occurs or there is no more data available (for example, the Iterator reaches End-of-Data).

Shutdown

When cycling stops then the AssemblyLine goes into Shutdown phase: Epilog Hooks are called and all initialized components are closed down (which flushes output buffers and executes their Epilog Hooks as well). Finally the AssemblyLine closes down its environment and its thread terminates.

AssemblyLine Pool

Actually a collection of AL *Flow* sections that can be configured to allow a Server mode Connector to service more clients. Available for ALs that use Server mode Connectors and set up in the AssemblyLine's Config tab.

Attribute

Part of the TDI Entry data model. Attributes are carried by Entry objects (Java "buckets", like the

3. If the current cycle was fed by a Server mode Connector, then the reply is created by the Server mode Connector's Output Map and sent to the client.

Work Entry) and they can hold zero or more *values*. These *values* are the actual data values read from, or written to connected systems, and are represented in TDI as Java objects.

Attribute Map (AttMap)

An Attribute Map is a list of rules (individual Attribute mapping instructions) for creating or modifying Attributes in an Entry object typically based on the values of Attributes found in another Entry object. Components like Functions and Connectors have an Input Map for taking data read into local cache (the conn Entry) and use this to define Attributes in the Work Entry. These components also have an Output Map that takes Attributes carried by the AssemblyLine (in its Work Entry) and use this to set up the conn Entry that is used by the component's output operation. Attribute Map components use the Work Entry as both the source and target of the mappings.

Attributes can be mapped in one of three ways: Simply (copying values between Attributes), Advanced (using a snippet of JavaScript), or with a TDI Expression.

Attribute Map component

A free-standing list of individual Attribute mappings that take values from the Work Entry and use them to create and update other Attributes in the Work Entry. They can be tied to Connector and Functions to define their Input or Output Maps. Note that Input and Output Maps can be copied to the library as AttMap components for reuse.

Best Practices

Recommended methodology and techniques for working with TDI. These include the ABCs: Automation, Brevity and Clarity:

Automation

Use the automated features of TDI in preference to your own custom scripted logic whenever possible – for example, using Branches/Loops instead of extensive scripting in Hooks. Not only does this make your solution easier to read and maintain (and you can step through with the AL Debugger!), but your solution benefits directly as built-in logic is strengthened and extended with each new release.

Brevity

Keep your AssemblyLines as short and simple as possible, as well as your script snippets. Break complex logic into simpler patterns that can be tested individually and reused in other solutions.

Clarity

Choose legibility over elegance. Write solutions for others to read and maintain.

Branches

A construct used to control the flow of logic in an AssemblyLine. Tivoli Directory Integrator 7.0 provides three types of Branches:

- Simple Branches (IF, ELSE-IF and ELSE)
- Loops (Connector-based, Attribute-based or Conditional)
- Switches (for example, switching on the Work Entry delta operation code, or the Operation an AL is called with).

CBE Common Base Event. A term used in the Common Base Infrastructure. See "Common Base Event" in the chapter about the CBE Generator Function Component in the *IBM Tivoli Directory Integrator V7.0 Reference Guide*.

CEI The IBM Common Event Infrastructure. See "The Common Event Infrastructure", in *IBM Tivoli Directory Integrator V7.0 Reference Guide*.

Change Detection Connector (CDC)

A Connector that returns changes made in the connected system. Typically, a CDC can be configured to return only a subset of Entries: new, modified, deleted, unchanged or a combination of these. Some CDC's provide only the changed Attributes in the case of a modified

Entry, while others return them all. Change Detection Connectors also tag the data with special *delta operation codes* to indicate what has changed, and how.⁴.

CLI Command-line Interface, such as the `tdisrvctl` utility.

cipher A cipher is any method of encrypting text (hiding its readability and its meaning). The resulting encrypted text message is called *ciphertext*.

ciphertext

Ciphertext is encrypted text, the result of applying a cipher, or an encryption.

Components

The architecture of IBM Tivoli Directory Integrator is divided into two parts: generic functionality and technology-specific features. Generic functionality is provided by the TDI *kernel* which provides automated behaviors to simplify building integration solutions. The kernel also lets you extend or override these behaviors as desired, as well as doing the housekeeping for your solution: logging/tracing, Hooks for error handling, API and CLI access, and so forth.

Technology-specific "intelligence" is handled by helper objects called *components*, such as Connectors, Functions, Branches, Scripts and Attribute Map components. Components provide a consistent and predictable way to access heterogeneous systems and platforms, and the kernel lets you "click" together components to build AssemblyLines.

Compute Changes

A special feature of the Connector Update mode that instructs the Connector to compare the Attributes about to be written to the connected system with those that exist in this data source already – in other words, it compares the value of each Attribute in the conn Entry (the result of the Output Map) with the corresponding ones found during the Update mode *lookup* operation (which is stored in the current Entry).

Config or Config File

A collection of AssemblyLines and components that comprise a solution. A Config is stored in XML format, typically in a Config file and is written, tested and maintained using the Configuration Editor.

Config Browser

This is the tree-view window at the top left-hand part of the Configuration Editor screen. It gives you access to Config-wide settings, the AssemblyLines and components that make up the Config, as well as Properties, *included* Configs and custom Java libraries that are to be loaded and made available to your scripts.

Configuration Editor (CE)

The graphical development environment used to write, test and maintain Configs. Configs are stored in XML format and are deployed by assigning them to one or more IBM Tivoli Directory Integrator Servers to run.

Config Instance

A copy of a TDI Config that is running on a Server. Typically loaded only once on a given Server, TDI allows you to start the same Config multiple times if desired. Each running copy is given its own context and can be accessed individually through the API.

conn Entry

This is the local Entry object maintained by a Connector or Function. The conn Entry is used as a local cache for read and write operations, and data is moved between this cache and the AssemblyLine's Work Entry via Attribute Maps (specifically, Input and Output Maps).

Connector

One of the component types available in TDI to build AssemblyLines. Connectors are used to

4. For LDAP there is also a special kind of modify operation where the directory entry has been moved in the tree: *modrdn*, that is, a "renamed" entry.

abstract away the technical details of a specific data store, API, protocol or transport, providing a common methodology for accessing diverse technologies and platforms.

Unlike the other components, Connectors can perform different tasks based on their *mode* setting (for example, Iterate, Delete, Server and Lookup). Modes are provided by the AssemblyLine component part of the Connector. However, the list of modes supported is dependent on the Connector Interface.

Connector Interface

When a component is used in an AssemblyLine, a distinction must be made between the *Connector Interface* (CI), containing the "intelligence" for working with a connected system (for example, LDAP, JDBC, Notes, and so forth), and the *AssemblyLine Connector*.⁵ This latter object is the "AL wrapper" that allows the CI to be plugged into an AssemblyLine and provides them with a consistent set of generic features, like input or output maps, Link Criteria, Hooks and the Delta Engine. See "Objects" in *IBM Tivoli Directory Integrator V7.0 Reference Guide* for more information. See also "Connectors" in *IBM Tivoli Directory Integrator V7.0 Reference Guide*.

Connector Pool

Unlike the AssemblyLine Pool feature available to ALs using Server mode Connectors, a Connector Pool is a global collection of pre-initialized Connectors that can be used in multiple ALs. Note that the Connector Initialization setting "Initialize and terminate every time it is used" means that no AssemblyLine gains exclusive rights to a pooled Connector, giving you detailed control over resources used by your solution.

current Entry

This Entry object is local to a Connector Interface (just like the conn Entry) and contains the Attributes read in from a *lookup* operation (for example, as carried out by Lookup, Update and Delete modes). It is used to provide the Compute Changes feature.

Delta Engine

Available for Connectors in Iterator mode, the Delta Engine provides functionality for detecting changes in data sources that do not offer any changelog or change notification features. See Delta Operation Codes, as well as "Deltas and compute changes" in *IBM Tivoli Directory Integrator V7.0 Users Guide* for more information.

Delta mode (for Connectors)

This Connector mode is used to apply changes specified with delta operation codes in the Work Entry, and to do so as efficiently as possible by performing incremental modifications. Note that Delta mode is only available for the LDAP and JDBC Connectors, and does not work with Entries without a valid delta operation code. See "Deltas and compute changes" in *IBM Tivoli Directory Integrator V7.0 Users Guide*.

Delta Operation Codes

These are special values assigned to Entries, Attributes and their values to reflect change information detected in some data source. An Entry that has delta codes assigned is called a *Delta Entry*, and these are only returned by a limited set of components: Change Detection Connectors, the Delta Engine and the DSML and LDIF Parsers⁶. Delta Operation Codes can be queried and used in Branch Conditions or your own JavaScript code, and are used by Delta mode to apply all types of changes to target systems as efficiently as possible.

See also "Deltas and compute changes" in *IBM Tivoli Directory Integrator V7.0 Users Guide*.

Derby Apache Derby (previously known as Cloudscape) is a small footprint relational database implemented entirely in Java. Derby v10.2 is shipped as the default System Store for TDI.

DES Short for Data Encryption Standard. DES is a widely-used method of data encryption using a secret key. DES is superseded by the Advanced Encryption Standard (AES).

5. Functions are similar to Connectors in that they are divided into two parts: the Function Interface and the AssemblyLine Function. Unlike Connectors, Functions have no mode setting.

6. Note that these Parsers only return Delta Entries if the DSML or LDIF entries read contain change information.

Distinguished Name (DN)

An LDAP term that refers to the fully qualified name of an object in the directory, representing the *path* from the root to this node in the directory information tree (DIT). It is usually written in a format known as the User Friendly Name (UFN). The dn is a sequence of *relative distinguished names* (RDNs) separated by a single comma (,).

ECB Short for Electronic Code Book. Electronic Code Book (ECB) is a method of operation for a block cipher. In ECB, each possible block of plaintext has a defined corresponding ciphertext value and the other way around. The same plaintext value always results in the same ciphertext value. Electronic Code Book is used when a volume of plaintext is separated into several blocks of data, each of which is then encrypted independently of other blocks. Moreover, Electronic Code Book can create a separate encryption key for each block type.

Entry An Entry is a TDI object used to carry data, and forms the core of the TDI Entry model. The Entry object can be thought of as a "Java bucket" that can hold any number of Attributes, which in turn carry the actual data values read from, or written to connected systems. Each Entry corresponds to a single row in a database table/view, a record from a file or an entry in a directory (or similar unit of data), and there are a number of named Entry objects available in the system. The Work Entry and conn Entry are the most commonly used ones, but there is also a current Entry available in some Connector modes, an error Entry that contains the details of the last exception that occurred, and an Operation Entry (Op-Entry) for accessing details of an AL operation.

Epilog A set of Hooks that, if enabled, are run during the AssemblyLine Shutdown phase. Note that the shutdown of components occurs between the two AL Epilog Hooks, which means that the Epilog Hooks of these components are all completed before the AssemblyLine Epilog - After Close Hook is called.

Error Entry

An Entry object that is created by an AssemblyLine during initialization, and contains Attributes like "status", "connectorname" (applies for all types of components) and "exception"⁷. See also Error Handling.

Error Handling

Error Handling in TDI is based on the concept of *exceptions*. Exceptions are a feature of a programming language, like Java, C and C++, that lets you build error handling like a wall around your program. It also lets you fortify smaller parts within any wall, so you can add specific handling where necessary. TDI leverages the power of exception handling so that you can design the error handling in your solution the same way.

First you have the AssemblyLine's On Failure Hook which is called if the AL stops due to an unhandled exception⁸. This is the outer line of defense⁹. The next level is a component, given that it provides Error Hooks. Connectors actually provide two levels of handling: the mode-specific Error Hook, as well as the Default On Error (same goes for Success Hooks as well).

Finally, in your JavaScript code you can do exception handling yourself. Use the try-catch statement, for example:

```
try {
    myObj = someFunctionCallThatCanThrowAnException();
} catch ( excptThrown ) {
    task.logmsg("***ERROR - The call failed: " + excptThrown );
}
```

7. The "exception" Attribute holds the actual Java exception object, in the case of an error – in which case the "status" Attribute would also be changed from a value of "ok" to "error" and "message" would contain the error text.

8. An "unhandled" error is one that is *caught* in an enabled Error Hook (no actual script code is necessary). If you wish to escalate an error to the next level of error handling logic, you need to re-throw the exception:

```
throw error.getObject("exception");
```

9. If you want to share this logic (or that in any Hook) between AssemblyLines, implement it as a function stored inScript and then include them as a Global Prolog for the AL.

ERP Enterprise Resource Planning, usually indicates a software suite of programs that aims to manage enterprise resources, usually after heavy customization by the software vendor.

Exception

See Error Handling.

External Properties

A type of Property Store that uses a flat file for storing configuration settings (like passwords and other component parameter settings) outside the Config itself.

Feeds This is the first section of an AssemblyLine and can only hold Iterator and Server mode Connectors. The Feeds section is where the Work Entry is created from data retrieved from a connected system or client. The Feeds section is like a built-in Loop that drives the Flow section components list, once for each Entry read.

FIPS Short for Federal Information Processing Standard. TDI uses FIPS 140-2, a standard that defines requirements for cryptographic modules that handle sensitive information.

Flow This is the second (and usually the main) section of an AssemblyLine and holds a list of components; any type, except Connectors in Server mode. The Flow section receives a Work Entry from the currently active Feeds Connector and passes it from component to component for processing.

Function component (FC)

One of the component types available in TDI to build AssemblyLines. Functions are used to abstract away the technical details of a specific service or method call. Typical examples are the AssemblyLine FC used to run ALs and the Java Class FC that lets you browse jar files and call class methods. Unlike Connectors, FCs do not have mode settings.

Global Prolog

This is a Script component that is defined in the "Scripts" library folder of the Config Browser, and which is configured to be executed when an AssemblyLine starts up. The simplest way to do this is to select which Scripts to use with the "Include Addition Prologs - Select" button. Note that Global Prologs are executed before the AssemblyLine's own Prolog Hooks.

GUI (ibmditk or ibmditk.bat)

The term "TDI GUI" is sometimes used to refer to the Configuration Editor.

Hook This is a *waypoint* in the built-in workflow of the AssemblyLine, or of a Connector or Function, where you can customize behavior by writing JavaScript. In a Connector, the Hooks available are also dependent on the mode setting.

HTML

HyperText Markup Language. a more or less standardized way of describing and formatting a page of text on the WorldWide Web. Different manufacturers' interpretations of the standard are often the cause of Web Browser's different renderings of a given page.

HTTP HyperText Transfer Protocol. The protocol in use for the WorldWide Web, another protocol on top of TCP.

IEHS IBM Eclipse Help System. Used to host the Tivoli Directory Integrator documentation locally. The documentation hosted by IBM in the Documentation Library also uses IEHS.

Initial Work Entry (IWE)

This is an Entry that is passed into an AssemblyLine by the process that called it (for example, an AssemblyLine Connector or Function, or by using script calls like `main.startAL()`). Note that the presence of an IWE causes any Iterators in the Flow section to skip on this cycle.

Iterator

A Connector mode¹⁰ that first creates a data result set (for example, by issuing a SQL SELECT

10. Connectors running in Iterator mode are often referred to as "Iterators".

statement, a LDAP search operation, opening a file for input, and so forth) and then returns one Entry at a time to the AL for processing. Iterators can reside in the AssemblyLine Feeds section where they drive data to Flow components. If they are placed in the Flow section then they still retrieve the next Entry from their result set for each AL cycle, but they do not *drive* AL cycling in this case.

Java Virtual Machine or JVM

IBM Tivoli Directory Integrator runs inside what is known as a Java Virtual Machine. It has its own memory management and is in most respects a computer within the computer.

Java API documentation (Javadocs)

A set of low-level API documentation, embedded in the product's source code and extracted by means of a special process during the product's build. In IBM Tivoli Directory Integrator the Java API documentation can be viewed by selecting the **Help -> Welcome** screen, **JavaDocs** link from the Configuration Editor.

JavaScript

The language you can use to fine tune the behavior of your AssemblyLines. Tivoli Directory Integrator 7.0 uses the IBM JSEngine.

JMS Java Messaging Service. A standard protocol used to perform guaranteed delivery of messages between two systems.

JNDI Java Naming and Directory Interface. See "JNDI Connector", in the *IBM Tivoli Directory Integrator V7.0 Reference Guide*.

Link Criteria

Link Criteria represent the matching rules defined for a Connector in Update, Lookup or Delete, and they must result in a single entry match in the connected system; otherwise either an Not Found or Multiple Found exception occurs. Note that a Lookup Connector tied to a Loop is an efficient way of dealing with lookup operations where no match (or multiple matches) are expected.

LDAP Lightweight Directory Access Protocol. An easier way of accessing (using TCP) a name services directory than the older Directory Access Protocol. Used in for example querying the IBM Directory Server.

Memory Queue (MemQ)

The MemQ is a TDI object that lets you pass any type of Java object (like Entries) between AssemblyLines running on the same Server. This feature is usually accessed through the MemQueue Connector (or the deprecated Memory Queue FC). See also System Queue for more on how to pass data between running ALs.

Message Prefix

All error messages and Info messages in IBM Tivoli Directory Integrator are prefixed with a unique Message Prefix. The prefix assigned to TDI is **CTGDI**.

Mode Connectors have a mode setting that determines how this component participates in AssemblyLine processing. In addition to the custom modes (implemented through Adapters) there is a set of standard modes:

- Iterator
- AddOnly
- Lookup
- Update
- Delete
- CallReply
- Server
- Delta

Dependent on the features provided by the underlying system or functionality built into the Connector, the list of modes supported by the different Connectors varies. See "Connectors" in *IBM Tivoli Directory Integrator V7.0 Reference Guide* for more information about Connector modes.

Null Value Behavior

This term refers to how TDI deals with Attribute mappings that result in "null" values. Null Behavior configuration can be done for a Server by setting Global/Solution properties. These Server-level settings can be overridden for an Attribute Map by pressing the **More** button in the button bar at the top of the map and selecting **Null behavior**; or for a specific Attribute via the **Null behavior** context menu item in the Assignment column for its mapping.

TDI lets you both configure what constitutes a "null" value situation (for example, missing values, empty string or a specific value) as well as how to handle this.

Op-Entry (Operation Entry)

An entry which contains information about the Operation for the currently executing AL. An Op-Entry persists its value over successive cycles for the same AL run and is available for scripting via the `task.getOpEntry()` method.

Parameter Substitution

A way of specifying patterns based on Java MessageFormat class - for simpler/quicker editing. Available in various places in Tivoli Directory Integrator.

Parser TDI components used to interpret or generate the structure for a byte stream. Parsers are used by attaching them to a Connector that reads/writes byte streams, or to a Function component like the Parser FC which is used to parse data in the Work Entry.

Persistent Object Store

See System Store.

Persistent Parameter Store

See Property Store.

plaintext

Plaintext is unencrypted text. In cryptography, plaintext is ordinary readable text before being encrypted into ciphertext or after being deciphered.

Prolog A set of Hooks that, if enabled, are run during the Assemblyline Initialization. You can also define Global Prologs: Script Components that are run before either of the AL Prolog Hooks. Note that the "At Startup" initialization of components occurs between the two AL Prolog Hooks, which means that the Prolog Hooks of these components are all completed before the AssemblyLine Prolog - after the Initialization Hook is called. See also Epilog.

Properties

This term refers to values maintained in a Property Store and used to configure AssemblyLine and Component settings at run-time.¹¹

Property Store

This is a feature for reading and writing all types of properties. This includes:

- Java-Properties, which are settings of the JVM.
- Global-Properties, Tivoli Directory Integrator Server settings that are kept in a file called `global.properties` in the "etc" folder of your installation directory.
- Solution-Properties, which typically override Global-Properties and are found in a file in your solution directory called `solution.properties`.
- System-Properties, for keeping custom property settings (uses the System Store).

11. Note that an Entry object can also hold *properties* (in addition to Attribute and delta operation codes) and these can be accessed using the `getProperty()` and `setProperty()` methods of the Entry class.

In addition, you can define your own Property Stores using a Connector. The Property Store feature also lets you designate one of your Property-Stores as a *Password Store*, giving you automatic protection of sensitive configuration details.

Raw Connector

Deprecated term; this is now called the Connector Interface and refers to the part of an AL Connector that contains the logic needed to access a specific API, protocol or transport.

Relative Distinguished Name (RDN[®])

In LDAP terms the name of an object that is unique relative to its siblings. RDNs have the form *attribute name=attribute value*. For example,

cn=John Doe

Resource Library

A simple method for sharing AssemblyLines and components between Configs. In the Configuration Editor, the "Resources" navigator appears just below the Config Browser.

RMI **Remote Method Invocation**; a way of making procedure or method calls on a remote system using a network communication channel. In Tivoli Directory Integrator, used by the Remote API functionality.

RSA RSA is an internet encryption and authentication system that uses an algorithm developed by Ron Rivest, Adi Shamir, and Leonard Adleman. The encryption system is owned by RSA Security. RSA is an algorithm for public-key cryptography, suitable both for signing and for encryption.

Sandbox

The feature of the IBM Tivoli Directory Integrator that enables you to record AssemblyLine operations for later playback without any of the data sources being present. See "Sandbox" in *IBM Tivoli Directory Integrator V7.0 Users Guide*.

SAP Used to stand for "Systeme, Anwendungen, Produkte" (Systems, Applications, Products) but today, the abbreviation just stands for itself. A large, German provider of an integrated suite of ERP applications. Mostly known for its R/3 distributed ERP software suite, but also known for its mainframe-based R/2 software.

Script Component (SC)

A Script is a block of JavaScript that is stored as a single component in TDI. In addition to appearing in the Scripts library folder of the Config Browser¹², Scripts can be dropped anywhere in the Flow section of an AssemblyLine.

Script Engine

The component that interprets the Java scripts written inside a TDI Config. The IBM jsEngine is used by Tivoli Directory Integrator 7.0, which replaces Rhino from the previous releases.

Schema

The word "Schema", unfortunately, can mean different although related things, depending on context. In a relational database context, a schema is the collection of tables and objects a user has defined and owns (including content); and each table in a schema is described by a Data Definition. In an LDAP context, the Schema is the actual layout of the LDAP database, with its attributes and objects.

In addition, Connectors and Functions can have Input and Output schemas that represent the data model discovered in a connected system. Furthermore, an AssemblyLine Operation can have an Input and Output schema as well.

12. In order to be used as Global Prologs (which are executed at the very start of Assemblyline Initialization) the Script must be in the *Scripts* library folder and selected for inclusion in the Config tab of an AssemblyLine.

In a product like TDI, which with equal ease can access both relational databases as well as LDAP databases, the word Schema can therefore mean different things, depending on where it is used.

Script Connector

A Script Connector is a Connector where you write the *Interface* functionality yourself: It is empty in the sense that, in contrast to an already-existing Connector, the Script Connector does not have the base methods `getNextEntry()`, `findEntry()` and so forth implemented. Not to be confused with the Script Component.

Server (`ibmdisrv` or `ibmdisrv.bat`)

This is the part of the TDI product that is used to deploy and run Configs.

Server (mode)

This is a Connector mode used for providing a request/response service (like an HTTP server). This mode also provides an AssemblyLine Pool feature to enable support for more connections/traffic.

Solution Directory

The directory in which you store your Config files, Derby databases, properties files, keystores and so forth. The solution directory is selected when you install TDI, and the filepaths used in your solution can be relative to this folder. The solution directory can be explicitly specified when you start the Configuration Editor or Server using the `-s` commandline option. Note that the counterpart of `global.properties` is kept in this folder and called `solution.properties`—unless, of course, your solution directory is the same as your installation directory.

Solution View

This term is used in the context of AMC to describe how a particular Config appears in the management screens of AMC. A Solution View is a selection of the AssemblyLines and properties that are to be visible onscreen (user/role based), providing solution-oriented Config administration and management. Config Views can be combined to define a Monitoring View in AMC.

SSL Secure Socket Layer; a protocol used in Internet communications to encrypt data such that if someone were to eavesdrop on the packets going back and forth he would not be able to see what the packets contain. The protocol was invented by Netscape; and you can see if a Web page uses the SSL protocol to talk to the Web server if it has the 'https://' prefix instead of 'http'. SSL is by no means limited to Web pages; in fact, Tivoli Directory Integrator uses it (if configured that way) to talk between different Tivoli Directory Integrator Servers and AssemblyLines if network access is called for.

State Defines the *level of participation* for an AssemblyLine component. It can be in either *Enabled* State, which means it participates in AL processing, or *Disabled* in which case the component is not used in any way.

Connectors and Functions can be set to a third State: *Passive*. Passive State causes the component to be initialized and closed during the Assemblyline Initialization and Shutdown phases, but never used during AL cycling. However, you can drive these components manually through script calls.

System Queue

A built-in queue infrastructure to facilitate the guaranteed delivery of messages between AssemblyLines, even running on different TDI Servers. By default, the System Queue uses the bundled MQe (WebSphere MQ Everyplace), but can be configured to leverage other JMS-compliant messaging systems. TDI provides a SystemQueue Connector to help you leverage this feature.

For more information about the System Queue and how to enable it, see the "System Queue" chapter in the *IBM Tivoli Directory Integrator V7.0 Installation and Administrator Guide*.

System Store

Called the Persistent Object Store, or POS in older TDI versions, the System Store is a relational database used to store state information, like Delta Tables (used by the Delta Engine) or Iterator state for Change Detection Connectors. It also provides the User Property Store which is accessible through the `system.setPersistentObject()`, `system.getPersistentObject()` and `system.deletePersistentObject()` methods. In the current implementation, the Derby product (previously known as **CloudScape**) is used. See <http://db.apache.org/derby> for more details.

Task By convention, all threads (AssemblyLines, EventHandlers and so forth) are referred to as *tasks* and are accessible from script code via the pre-registered **task** variable.

Task Call Block

A Java structure used to pass parameters to and from AssemblyLines. Often referred to by its abbreviation: **TCB**.

TCP Transmission Control Protocol, a level 4 (transmission integrity) protocol usually seen in combination with its layer 3 (routing) Internet Protocol as in TCP/IP. A stack of protocols designed to achieve a standardized way of communicating across a network, be it local (as in on the premises) or over long distances. Originally invented and specified by DARPA, the (US) Defense Advanced Research Projects Agency. Successor to ARPANET, which was a network of a (small) number of universities and the US Department of Defense, the civil side of which was managed by the Stanford Research Institute (SRI). TCP is related to UDP.

TDI Unofficial monicker for this product, IBM Tivoli Directory Integrator.

TMS XML

Tivoli Message Standard XML. A Tivoli standardized way of formatting messages. Each message is prefixed by a unique TMS code, which can be looked up in the Message Guide for explanation and user response. If the code ends in "E" - it indicates an Error, "W" indicates a warning and "I" indicates an Information message. All Tivoli messages issued by TDI start with this product's unique identifier, which is "CTGDI".

Tombstone

A record or trace showing that an AssemblyLine, an EventHandler or Config has terminated. Configured through the Tombstone Manager in the CE. The trace includes a timestamp and the AL exit status. The Tombstone Manager creates a tombstone for each AssemblyLine as it terminates.

TWiki TWiki as a piece of software is a flexible and easy to use enterprise collaboration system. Its structure is similar to the Wikipedia, except that is not linked into that. It is rather meant as an independent community resource for a group of people with common interest. There is one for IBM Tivoli Directory Integrator as well, at <http://www.tdi-users.org>.

Note: The TWiki site is a volunteer effort, and is not an official Tivoli support forum. If you need immediate assistance, contact your local Tivoli support organization.

Update

One of the standard Connector modes. Update mode causes the Connector to first perform a lookup for the entry you want to update¹³, and if found it modifies this entry. If no match is found then a new entry is added instead. See also Computed Changes.

UDP User Datagram Protocol. A protocol used on top of the Internet Protocol (IP) which, unlike TCP does **not** guarantee that the packet of data sent with it reaches the other end. Also see TCP.

URL Unified Resource Locator. A way of defining where a resource is, be it a fileserver or a HTML page on the WorldWide Web.

13. Data is read into both the conn and current Entry objects. After the Output Map, the contents of conn are now the Attributes to be written. The original entry data is still available in current.

User Property Store

See Property Stores in the *IBM Tivoli Directory Integrator V7.0 Users Guide*.

Value (data values and types)

See Entries, and Attribute.

WikiPedia

A Web-based world-wide encyclopedia, where (registered) users can add articles or pictures, edit them, browse them, search for applicable content, and so forth. For Tivoli Directory Integrator there is one that similar in functionality but not linked into the WikiPedia, a "TWiki" at <http://www.tdi-users.org>. The TWiki is a groupware product.

Work Entry

An Entry object that is used by the AssemblyLine to carry data from component to component¹⁴. This object can be accessed in script code via the pre-defined variable work. The Work Entry is typically built by a Server or Iterator mode Connector in the Feeds section before being passed to the AL Flow section. You can also have an Initial Work Entry (IWE) passed in if the AL was called from another process; or you can create it in the Prolog by using `task.setWork()`:

```
init_work = system.newEntry(); // Create a new Entry object
init_work.setAttribute("uid", "cchateauvieux"); // populate it
task.setWork(init_work); // make it known as "work" to the Connectors
```

Note that an Iterator in the Feeds section does not return any data if the Work Entry is already defined at this point in the AL. So if an IWE is passed into an AssemblyLine, any Iterators in the Feeds section simply pass control to the next component in line. It is also the reason why multiple Iterators in the Feeds section run sequentially, one starting up when the previous one reaches End-of-Data.

XML The Xtensible Markup Language. A general purpose markup language (See also HTML) for *creating* special-purpose markup languages, and also capable of describing many types of data IBM Tivoli Directory Integrator uses XML to store Config files.

14. Note that the "Work Entry" window shown in the Configuration Editor is actually a list of all Attributes that appear in Input Maps or in the Loop Attribute field of Loops in the AssemblyLine.

Appendix B. Example Property files

An installation of IBM Tivoli Directory Integrator is to a large extent customized by means of a set of text files containing one of more **properties**, usually in the form of a keyword or identifier followed by a value. The following global property text files can be found at the root/etc level of the IBM Tivoli Directory Integrator installation directory:

- “Log4J.properties”
- “jlog.properties” on page 266
- “derby.properties” on page 268
- “global.properties” on page 268

Properties set in any of those files form a baseline for the entire IBM Tivoli Directory Integrator installation for all users on that computer. However, if your Solution Directory is different from the installation directory, you can have a set of text files in your Solutions Directory that mirror their counterparts in the installation directory. A property listed in any of those files overrides anything set in any of the global installation property files mentioned above. Furthermore, a Java property set inside a Config file takes the highest precedence, and overrides anything in a global property file or the property files in the Solution Directory.

You can specify the Solution Directory in multiple ways:

- By setting the environment variable *TDI_SOLDIR* before starting the Configuration Editor or the Server
- By specifying the *-s* parameter to the *ibmditk* script to start the Server. This takes precedence over setting *TDI_SOLDIR*.

If *TDI_SOLDIR* equals the installation directory, the behavior is like in older versions of TDI: all property files are read from there, and the remarks about property files in the Solutions Directory do not apply.

In any other case, the first time you run the Tivoli Directory Integrator Server, it makes a copy of all the property files into your Solutions Directory (it does not overwrite these files if they already exist). You can now tailor these files to your particular needs, without affecting the property files in the installation directory. The files remaining in the installation directory continue to form a baseline configuration for other instances of Tivoli Directory Integrator.

Note: The file *global.properties* is copied to a file called *solutions.properties* in your Solutions Directory. Other files, like *Log4J.properties* and the files in the *amc* and *serverapi* folders are copied under their own name.

Log4J.properties

This file sets a baseline for the log-strategy for the server (*ibmdisrv*).

Log options configured in the Logging tab in the Configuration Editor are written into the Config file, and are supplementary to or supersede the following:

```
# This file controls the logging strategy for the server (ibmdisrv) when started
# from the command line.
# Look at executetask.properties for the logging strategy of the server when started
# from the Configuration Editor (ibmditk).
# Look at ce-log4j.properties for the logging behavior of the Configuration Editor (ibmditk).
#
# You will normally configure the logging strategy of the server by adding appenders
# using the Configuration Editor (ibmditk). This file only defines the baseline
# that is independent of the configuration files you are using.
#
# See the IDI documentation for more information on the contents of this file.
#
```



```

log4j.rootCategory=INFO, Default

# This is the default logger, you will see that it logs to ibmdi.log
log4j.appender.Default=org.apache.log4j.FileAppender
log4j.appender.Default.file=logs/ibmdi.log
log4j.appender.Default.layout=org.apache.log4j.PatternLayout
log4j.appender.Default.layout.ConversionPattern=%d{ISO8601} %-5p [%c] - %m%n
log4j.appender.Default.append=false

#Example settings for changing the default logger

#####ROLLING FILE SIZE APPENDER
##RollingFileAppender rolls over log files when they reach a certain size specified by the
##MaxFileSize parameter

#log4j.appender.Default=org.apache.log4j.RollingFileAppender
#log4j.appender.Default.File=logs/ibmdi.log
#log4j.appender.Default.Append=true
#log4j.appender.Default.MaxFileSize=10MB
#log4j.appender.Default.MaxBackupIndex=10
#log4j.appender.Default.layout=org.apache.log4j.PatternLayout
#log4j.appender.Default.layout.ConversionPattern=%d{ISO8601} %-5p [%c] - %m%n

#####DAILY OUTPUT LOG4J SETTINGS
## With the DailyRollingFileAppender the underlying file is rolled over at a user chosen frequency.
##The rolling schedule is specified by the DatePattern option

#log4j.appender.Default=org.apache.log4j.DailyRollingFileAppender
#log4j.appender.Default.file=logs/ibmdi.log
#log4j.appender.Default.DatePattern='.'yyyy-MM-dd
#log4j.appender.Default.layout=org.apache.log4j.PatternLayout
#log4j.appender.Default.layout.ConversionPattern=%d{ISO8601} %-5p [%c] - %m%n

# You may change the logging category of these subsystems to DEBUG
# if you want to investigate particular problems. This may
# generate a lot of output.
# ...com.ibm.di.config describes the loading of the configuration file (.xml),
# and how the internal configuration structure is built.
# ...com.ibm.di.loader gives information about jar files, and where classes are found.
# It also loads idi.inf files, which provides Connectors/Parsers/EH information
# for the Configuration Editor.
log4j.logger.com.ibm.di.config=WARN
log4j.logger.com.ibm.di.loader=WARN

# Uncomment the lines below to activate them

# Here is an example on how to make a logger that logs to the console
#log4j.appender.CONSOLE=org.apache.log4j.ConsoleAppender
#log4j.appender.CONSOLE.layout=org.apache.log4j.PatternLayout
#log4j.appender.CONSOLE.layout.ConversionPattern=%d [%t] %-5p - %m%n

# Here is an example that logs to myFile.log
#log4j.appender.fileLOG=org.apache.log4j.FileAppender
#log4j.appender.fileLOG.file=myFile.log
#log4j.appender.fileLOG.layout=org.apache.log4j.PatternLayout
#log4j.appender.fileLOG.layout.ConversionPattern=%d{ISO8601} %-5p [%c] - %m%n
#log4j.appender.fileLOG.append=false

# Finally, make use of the loggers defined above:
# Tell AssemblyLines myAL to log using CONSOLE logger defined above.

# log4j.logger.AssemblyLine.AssemblyLines/myAL=INFO, CONSOLE

# Or you could log to myFile.log

# log4j.logger.AssemblyLine.AssemblyLines/myAL=INFO, fileLOG

# Finally, make EventHandler myEH log to myFile.log
# log4j.logger.EventHandler.EventHandlers/myEH=INFO, fileLOG

```

jlog.properties

This file configures the JLOG-based Chapter 15, “Tracing and FFDC,” on page 179 of the TDI server. These values can be modified dynamically (during Server execution) using the LogCmd script if the property jlog.noLogCmd was set to **false** when the Server started.

Note: You would normally use Log4J to trace execution flow in your solution; the JLOG-based tracing and FFDC is meant to aid IBM Support should you have problems with IBM Tivoli Directory Integrator.

```
#####
# This file controls the tracing and First Failure Data Capture (FFDC) strategy for ITDI 7.0
# See the IDI documentation for more information on the contents of this file.
#####

#-----
# Enable the JLOG's command server
#
# If the jlog.noLogCmd is set to false, then the JLOG LogManager will listen on the
# default port (9992) for JLOG log commands.
# Setting this property to false will enable you to modify the JLOG properties dynamically using the
# logcmd scripts. The logcmd scripts are placed under ITDI_HOME directory.
# The default value is set to true.
#-----
jlog.noLogCmd=true

#-----
# Set listen port for JLOG's command server
#
# If you want LogManager to listen on different port than the default one (9992) you should
# uncomment the property jlog.logCmdPort and set it to the desired port. If not uncommented
# the LogManager will listen on the default port - 9992.
#-----
jlog.logCmdPort=9992

#-----
# Configure Jlog FileHandler for tracing into a file.
#
# By default the FileHandler is not attached to the Jlog Logger.
# Uncomment the properties with the prefix jlog.filehandler below to configure a FileHandler.
# After uncommenting this you need to add the filehandler to the logger's listeners names as shown
# below
# e.g: jlog.logger.listenerNames=jlog.snapmemory jlog.snaphandler jlog.filehandler
#-----
jlog.filehandler.className=com.ibm.log.FileHandler
jlog.filehandler.description=JLOG File Handler for Logging and Tracing
jlog.filehandler.encoding=UTF8
jlog.filehandler.maxFiles=10
jlog.filehandler.maxFileSize=2048
jlog.filehandler.appending=true
jlog.filehandler.fileDir=logs/
jlog.filehandler.trace.fileName=trace.log
#-----

#-----
# create a level filter.
# The level filter is used to define the level at which JFFDC action will be triggered.
# For JFFDC to be meaningful this should be set to either FATAL or ERROR (case-insensitive).
# NOTE: Setting the trigger level to other levels such as DEBUG_MIN will trigger unwanted JFFDC
# action causing a performance drop.
#-----
jlog.levelflt.className=com.ibm.log.LevelFilter
jlog.levelflt.level=FATAL

#-----
# Configure the SnapMemoryHandler for tracing into a memory buffer.
# The SnapMemoryHandler traces into a memory buffer and dumps the contents of the memory to a file on
# trigger of a event (as defined by the level filter above) and writes the content to the specified
# file
# Properties:
# jlog.snapmemory.queueCapacity : Sets the nnumber of LogEvents that can be buffered in the memory
# jlog.snapmemory.snapFile : name of the file to which the contents of the memory will be dumped
# jlog.snapmemory.baseDir : The directory where the snapFile is placed.
# daily subdirectories will be created under this base directory, as:
# [baseDir]/[YYYY-MM-DD]/
# Note: MS-DOS style path names need to be escaped with backslashes
# eg: c:\CTGI\FFDC
# jlog.snapmemory.userSnapFile : The name of the file to which the user initiated (from logcmd) dumps
# will be written to.
# jlog.snapmemory.userSnapDir : The directory where the userSnapfile is placed.
# jlog.snapmemory.msgIds : The list of TMS IDs
# jlog.snapmemory.msgIDRepeatTime : The minimum time, in milliseconds, after passing a log event with a
# given TMS message id, before another log event with the same id can
# be passed.
#-----
jlog.snapmemory.className=com.tivoli.log.SnapMemoryHandler
```

```
jlog.snapmemory.description=Memory handler used to trace to memory
jlog.snapmemory.queueCapacity=10000
jlog.snapmemory.dumpEvents=true
jlog.snapmemory.snapFile=trace.log
jlog.snapmemory.baseDir=CTGDI/FFDC/
jlog.snapmemory.userSnapFile=userTrace.log
jlog.snapmemory.userSnapDir=CTGDI/FFDC/user/
jlog.snapmemory.triggerFilter=jlog.levelflt
jlog.snapmemory.msgIds=*E
jlog.snapmemory.msgIDRepeatTime=10000

#-----
# Configure the JLogSnapHandler taking a snapshot of the SnapMemoryHanlders buffer
# The JLogSnapHanlder takes a snapshot of the associated SnapMemoryBuffer.
#-----
jlog.snaphandler.className=com.tivoli.log.JLogSnapHandler
jlog.snaphandler.description=snaphandler to dump the memory trace
jlog.snaphandler.baseDir=CTGDI/FFDC/
jlog.snaphandler.snapMemoryHandler=jlog.snapmemory
jlog.snaphandler.triggerFilter=jlog.levelflt

#-----
# Configure the PDLogger (Problem Determination) Object and attach the Listeners to it.
# jlog.logger.level can be FATAL | ERROR | WARNING | INFO | DEBUG_MIN | DEBUG_MID | DEBUG_MAX
# The heirarchy of the log levels is from the most severe (FATAL) to the least severe (DEBUG_MAX)
# The value for this property is case-insensitive
#-----
jlog.logger.level=FATAL
#jlog.logger.listenerNames=jlog.snapmemory jlog.snaphandler
jlog.logger.listenerNames=jlog.filehandler.trace
jlog.logger.className=com.ibm.log.PDLogger

#-----
# Configure the PDLogger for the Config Editor and attach the Listeners to it.
# By default, no listeners are attached
#-----
jlog.logger.config-editor.level=FATAL
jlog.logger.config-editor.listenerNames=
```

derby.properties

This file contains some defaults for Derby in networked mode. Most TDI-related Derby parameters are not maintained here but in `global.properties` and `solution.properties`. More information about these parameters can be obtained from the Derby documentation.

```
# This is a sample properties file provided to show the proper format.
# We're also setting one property which make sure that
# Derby adds to the error log instead of overwriting it.
# This mode is useful for development.
derby.drda.logConnections=true
derby.drda.maxThreads=0
derby.drda.portNumber=1527
derby.drda.traceAll=true
derby.drda.timeSlice=0
derby.drda.traceDirectory=/trace
```

global.properties

This file is read by `ibmditk` (the CE) and `ibmdisrv` (the server) on startup. This file is read and applied before a file called `solution.properties` from your Solution Directory is read and applied.

Note:

The rendition here, due to extremely long line lengths, may not be complete. Refer to an actual `global.properties` file instead.

```
##
## This file is read by ibmditk/ibmdisrv on startup
##
## Enter <name>=<value> to set system properties.
## Enter !include <file | url> to include other files
##

com.ibm.di.securityTransformation=DES/ECB/NoPadding
```

```

##
## Modify the line below to add your own jar/zip files.
## The property may specify several directories or jar files, separated by the Java Property "path.separator",
## which is ":" on Linux and ";" on Windows
## Directories will be searched recursively by the TDI Loader for jar files containing classes and resources.
## Only files with a ".zip" or ".jar" extension are searched.
# com.ibm.di.loader.userjars=c:\myjars

##
## Modify the line below to enable the config autoload feature. When this property is defined, the "ibmdisrv -d" command
## line will look for *.xml files in the directory specified by this property and start each one.
##
# com.ibm.di.server.autoload=autoload.tdi

##
## SYSTEM STORE
##

## Location of the database (embedded mode) - Cloudscape 10
#com.ibm.di.store.database=TDISysStore
#com.ibm.di.store.jdbc.driver=org.apache.derby.jdbc.EmbeddedDriver
#com.ibm.di.store.jdbc.urlprefix=jdbc:derby:
#com.ibm.di.store.jdbc.user=APP
#{protect}-com.ibm.di.store.jdbc.password=APP

## Location of the database to connect (networked mode) - Cloudscape 10 - DerbyClient driver
com.ibm.di.store.database=jdbc:derby://localhost:1527/C:\Program Files\IBM\TDI\V7.0\TDISysStore;create=true
com.ibm.di.store.jdbc.driver=org.apache.derby.jdbc.ClientDriver
com.ibm.di.store.jdbc.urlprefix=jdbc:derby://localhost:1527/
com.ibm.di.store.jdbc.user=APP
{protect}-com.ibm.di.store.jdbc.password=APP

#
## Derby (Cloudscape) properties required for enabling authentication
#
derby.drda.startNetworkServer=true
derby.connection.requireAuthentication=true
derby.authentication.provider=BUILTIN
derby.database.defaultAccessMode=fullAccess

#
## Details for starting Cloudscape in network mode.
## Note: If the com.ibm.di.store.hostname is set to localhost then remote connections will not be allowed.
## If it is set to the IP address of the local machine - then remote clients can access this Cloudscape
## instance by mentioning the IP address. The network server can only be started for the local machine.
#
#com.ibm.di.store.start.mode=automatic
com.ibm.di.store.hostname=localhost
com.ibm.di.store.port=1527
com.ibm.di.store.sysibm=true

# the varchar(length) for the ID columns used in system store and pes connector tables
com.ibm.di.store.varchar.length=512

## create statements for system store tables (Cloudscape 5.1)
#com.ibm.di.store.create.delta.systable=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SEQUENCEID int, VERSION int)
#com.ibm.di.store.create.delta.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SEQUENCEID int, ENTRY long varbinary )
#com.ibm.di.store.create.property.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, ENTRY long varbinary )
#com.ibm.di.store.create.sandbox.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, ENTRY long varbinary )

## create statements for system store tables (Cloudscape 10)
com.ibm.di.store.create.delta.systable=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SEQUENCEID int, VERSION int);ALTER TABLE {0} ADD CONSTRAINT
com.ibm.di.store.create.delta.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SEQUENCEID int, ENTRY BLOB );ALTER TABLE {0} ADD CONSTRAINT
com.ibm.di.store.create.property.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, ENTRY BLOB );ALTER TABLE {0} ADD CONSTRAINT IDI_PS_{UNI
com.ibm.di.store.create.sandbox.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, ENTRY BLOB )
com.ibm.di.store.create.recal.conops=CREATE TABLE {0} (METHOD varchar(VARCHAR_LENGTH), RESULT BLOB, ERROR BLOB)

## create statements for system store tables DB2 on z/OS
#com.ibm.di.store.create.delta.systable=CREATE TABLESPACE TS1DSYS LOCKSIZE ROW BUFFERPOOL BP32K;CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SE
#com.ibm.di.store.create.delta.store=CREATE TABLESPACE TS1DST LOCKSIZE ROW BUFFERPOOL BP32K;CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SE
#com.ibm.di.store.create.property.store=CREATE TABLESPACE PS3DST LOCKSIZE ROW BUFFERPOOL BP32K;CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, SE
#com.ibm.di.store.create.sandbox.store=CREATE TABLE {0} (ID VARCHAR(VARCHAR_LENGTH) NOT NULL, ENTRY BLOB)
#com.ibm.di.store.create.recal.conops=CREATE TABLESPACE IM{UNIQUE} LOCKSIZE ROW BUFFERPOOL BP32K;CREATE TABLE {0} (METHOD VARCHAR(VARCHAR_LENGTH), RE

# Set a customized SQL statement for creation of the Tombstone Manager table. Keep the same table and field names.
#com.ibm.di.store.create.tombstones=CREATE TABLE IDI_TOMBSTONE ( ID INT GENERATED ALWAYS AS IDENTITY, COMPONENT_TYPE_ID INT, EVENT_TYPE_ID INT, START

```

```

# the ibmsnap_commitseq column name used by the RDBMS changelog connector
com.ibm.di.conn.rdbmschlog.cdcolname=ibmsnap_commitseq

## server authentication
javax.net.ssl.trustStore=serverapi/testadmin.jks
{protect}-javax.net.ssl.trustStorePassword=administrator
javax.net.ssl.trustStoreType=jks

## client authentication
javax.net.ssl.keyStore=serverapi/testadmin.jks
{protect}-javax.net.ssl.keyStorePassword=administrator
javax.net.ssl.keyStoreType=jks

##PKCS11 options
##Set the value of following properties to use PKCS11 enabled devices to store TDI servers private key / certificate.
com.ibm.di.pkcs11cfg=etc/pkcs11.cfg
com.ibm.di.server.pkcs11=false
com.ibm.di.server.pkcs11.library=
com.ibm.di.server.pkcs11.slot=
{protect}-com.ibm.di.server.pkcs11.password=PASSWORD

## Turns on java debug
# javax.net.debug=true

## java interpreter override
# com.ibm.di.javacmd=
# com.ibm.di.installdir=

## Limits the number of threads IDI uses
## Must be set higher than 3 to have any effect

# com.ibm.di.server.maxThreadsRunning=500

com.ibm.di.server.securemode=false

## Following properties modified in TDI 7.0 .Added property for
## keystore password and keypassword
## com.ibm.di.server.keystore
## com.ibm.di.server.key.alias

api.keystore=testserver.jks
api.key.alias=server
{protect}-api.keystore.password=server
{protect}-api.key.password=

## Encryption properties added in TDI 7.0
com.ibm.di.server.encryption.keystore = testserver.jks
com.ibm.di.server.encryption.key.alias = server
com.ibm.di.server.encryption.keystoretype = jks
com.ibm.di.server.encryption.transformation = RSA

## Server API properties
## -----

api.on=true
api.audit.on=false
api.user.registry=serverapi/registry.txt
api.user.registry.encryption.on=false

api.remote.on=true
api.remote.ssl.on=true
api.remote.ssl.client.auth.on=true
api.remote.naming.port=1099
api.truststore=testserver.jks
{protect}-api.truststore.pass=server

## Specifies a list of IP addresses to accept non SSL connections from (host names are not accepted).
## Use space, comma or semicolon as delimiter between IP addresses. This property is only taken into account
## when api.remote.ssl.on is set to false.
## api.remote.nonssl.hosts=

api.jmx.on=false
api.jmx.remote.on=false

## The configuration files placed in this folder can be edited through the Server API.
## Configuration files placed in other folders cannot be edited through the Server API.
api.config.folder=configs

## Timeout in minutes for configuration locks. A value of 0 means no timeout.
api.config.lock.timeout=0

```

```

## Timeout in minutes for loading a configuration.
api.config.load.timeout=2

## Specifies if the Server API methods for custom method invocation (Session.invokeCustom(...)) are allowed to be used.
## When api.custom.method.invoke.on is set to false and the Server API methods for custom method invocation are used,
## then an exception will be thrown.
## Only classes listed in api.custom.method.invoke.allowed.classes are allowed to be directly invoked.
## The default value is false.
api.custom.method.invoke.on=false

## Specifies the list of classes which can be directly invoked by the Server API methods for custom
## method invocation (Session.invokeCustom(...)).
## This property is only taken into account if api.custom.method.invoke.on is set to true.
## The classes in this list must be separated by a space, a comma or a semicolon.
## Example:
## api.custom.method.invoke.allowed.classes=com.ibm.MyClass,com.ibm.MyOtherClass
## In the above example only methods from the com.ibm.MyClass and com.ibm.MyOtherClass classes are
## allowed to be directly invoked.
api.custom.method.invoke.allowed.classes=

## Specifies a list of Server notification types, which will be suppressed.
## Notifications of suppressed types will not be propagated by the notifications framework.
## The notification types in the list are separated by spaces. Wildcards may be included.
## Example:
## api.notification.suppress=di.al.* di.ci.start
## The above example will suppress all Assembly Line related notifications as well as
## notifications for starting a configuration instance.
## If the property is missing or is empty, no notifications will be suppressed.
api.notification.suppress=di.server.api.authenticate di.server.api.authorize.*

## api.custom.authentication points to a JavaScript text file that contains custom authentication code.
## For example: api.custom.authentication=ldap_auth.js.
## To enable the built-in LDAP Authentication mechanism, set this property to "[ldap]".
## To enable the built-in JAAS Authentication mechanism, set this property to "[jaas]".
## For example: api.custom.authentication=[ldap]

##api.custom.authentication=[ldap]

## LDAP Authentecation properties
## -----

## If this parameter is set to "true" and the LDAP Authentecation initialization fails, the whole Server API will not be started.
## If this parameter is missing or is set to "false" any LDAP Authentication initialization errors will be logged and the Server API will be started.
api.custom.authentication.ldap.critical=false

## LDAP Server hostname.
api.custom.authentication.ldap.hostname=

## LDAP server port number. For example, 389 for non-SSL or 636 for SSL.
api.custom.authentication.ldap.port=

## Specifies whether SSL is used to communicate with the LDAP Server.
## When set to "true" SSL will be used, otherwise SSL will not be used.
api.custom.authentication.ldap.ssl=

## Specifies the LDAP directory location where user searches will be preformed.
## When this property is not specified user searches will not be performed.
api.custom.authentication.ldap.searchbase=

## Specifies the user id attribute to be used in searches.
## When this property is not specified user searches will not be performed.
api.custom.authentication.ldap.userattribute=

## Specifies an LDAP Server administrator distinguished name that will be used for user searches.
## When this property is not specified anonymous bind will be used for user searches.
api.custom.authentication.ldap.admindn=

## Password for the LDAP Server administrator distinguished name.
{protect}-api.custom.authentication.ldap.adminpassword=

## This property specifies whether LDAP Group authentication is turned on.
## If it is set to 'true', the group membership of the authenticating user will be resolved and will be taken into account during authorization.
## If it is missing, the default value 'false' is used.
api.custom.authentication.ldap.groupsupport=false

## Specifies the name of the attribute of a user in LDAP that contains a list of the groups of which the user is a member.
## It is taken into account only if 'api.custom.authentication.ldap.groupsupport' is set to true.
api.custom.authentication.ldap.usermembershipattribute=

## Specifies how groups are named in the membership attribute of a user.
## For example, if the user's membership attribute contains values, which correspond to the 'objectSID' attributes of groups, set this property to 'objectSID'.
## If the user's membership attribute contains distinguished names of groups, then set this property to 'dn'.
## The property is required in case 'api.custom.authentication.ldap.groupsupport' is set to true.

```

```

api.custom.authentication.ldap.usermembershipattributecontent=

## Specifies the name of a group's attribute in LDAP, which corresponds to the way the group is named in the TDI User Registry.
## For example, if LDAP groups are addressed in the TDI registry by their common name, then set this property to 'cn'.
## If the User Registry contains the distinguished names of the groups, then set this property to 'dn'.
api.custom.authentication.ldap.groupnameattribute=

## Represents the LDAP directory context, where groups will be searched.
## It is required only when LDAP group support is enabled
api.custom.authentication.ldap.groupsearchbase=

## Optional property, which represents a list of space-separated attribute names. Specifies attributes which have non-string syntax.
api.custom.authentication.ldap.binaryattributes=

## JAAS Authentication properties
## -----
java.security.auth.login.config=

## Enabling/Disabling FIPS Mode in TDI
##-----
## If the below property is set to true then TDI will be enforced to run in FIPS Compliant Mode.
## The default value is false, i.e. TDI will not run in FIPS Mode by default.
com.ibm.di.server.fipsmode.on=false

## Specify the unique ID for the TDI Server
## -----
## This property helps a client connecting to the TDI server to identify different servers
## running on the same IP and the same port in different time. (Default is DEFAULT_ID)
com.ibm.di.server.id=DEFAULT_ID

## Tombstone Manager properties
## -----

com.ibm.di.tm.on=false
com.ibm.di.tm.autodel.age=0
com.ibm.di.tm.autodel.records.trigger.on=10000
com.ibm.di.tm.autodel.records.max=5000
com.ibm.di.tm.create.all=false

## -----
## Help system properties
## -----

## Name of help server, comment out if you want local help system
## The Tivoli library is at the following URL:
## http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp?topic=/com.ibm.iea.tdi/tdi/TDIv70_Task.html
com.ibm.di.helpHost=publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp?topic=

## Port for help system
com.ibm.di.helpPort=80

## -----
## AssemblyLinePool: Connector pooling defaults
## -----
##
## Note! These settings are only used when an AssemblyLine uses
## an AssemblyLinePool in combination with a Server mode connector.

## The number of seconds before a pooled connector times (e.g. is closed and no longer reused)
## Less than zero means disable connector pooling
## Zero means never timeout
## Greater than zero sets the number of seconds before a connector is closed
com.ibm.di.server.connectorpooltimeout=42

## Comma separated list of connector interfaces that we never pool
com.ibm.di.server.connectorpoolexclude=com.ibm.di.connector.FileConnector,com.ibm.di.connector.ScriptConnector

## Properties for Windows IPv6 communications.
## Uncomment these properties for Windows IPv6 communication only.
## These properties will not affect IPv4 communication or IPv6 communication on Unices.
#java.net.preferIPv4Stack=false
#java.net.preferIPv6Addresses=true

## -----
## Performance settings
## -----
##
## Enable/Disable performance logging
com.ibm.di.server.perfStats=false

### -----

```



```

### Used by Config Report
###-----
### set this is you want to override the local language for Config Reports
# com.ibm.di.admin.configreport.translation=en

###-----
## System Queue settings
###-----
## If set to "true" the System Queue is initialized on startup and can be used;
## otherwise the System Queue is not initialized and cannot be used.
systemqueue.on=true

## Specifies the fully qualified name of the class that will be used as a JMS Driver.
# systemqueue.jmsdriver.name=com.ibm.di.systemqueue.driver.IBMMQ
# systemqueue.jmsdriver.name=com.ibm.di.systemqueue.driver.JMSScriptDriver
systemqueue.jmsdriver.name=com.ibm.di.systemqueue.driver.IBMMQe

### MQe JMS driver initialization properties
## Specifies the location of the MQe initialization file.
## This file is used to initialize MQe on TDI server startup.
systemqueue.jmsdriver.param.mqe.file.ini=MQePWStore/pwstore_server.ini

### MQ JMS driver initialization properties
# systemqueue.jmsdriver.param.jms.broker=<host:port>
# systemqueue.jmsdriver.param.jms.serverChannel=<channel_name>
# systemqueue.jmsdriver.param.jms.qManager=<queuemanger_name>
# systemqueue.jmsdriver.param.jms.sslCipher=<cipherSuite_name>
# systemqueue.jmsdriver.param.jms.sslUseFlag=false

### JMS Javascript driver initialization properties
## Specifies the location of the script file
# systemqueue.jmsdriver.param.js.jsfile=driver.js

## This is the place to put any JMS provider specific properties needed by a JMS Driver,
## which connects to a 3rd party JMS system.
## All JMS Driver properties should begin with the 'systemqueue.jmsdriver.param.' prefix.
## All properties having this prefix are passes to the JMS Driver on initialization after
## removing the 'systemqueue.jmsdriver.param.' prefix from the property name.
# systemqueue.jmsdriver.param.user.param1=value1
# systemqueue.jmsdriver.param.user.param2=value2
# ...

## Credentials used for authenticating to the target JMS system
# {protect}-systemqueue.auth.username=<username>
# {protect}-systemqueue.auth.password=<password>

## -----
## Logging settings
## -----

## When false, all log calls made through the TDI Log class will be discarded.
com.ibm.di.logging.enabled=true

## -----
## IBM JavaScript Engine settings
## -----

## Set the type of platform - required by the IBM JS Engine when caching is used.
com.ibm.common.platform=com.ibm.common.platform.GenericPlatform

##
## Set this property to a directory to enable auto dumps of assemblylines that fails
##
# com.ibm.tdi.autodump.directory=<dump-directory>

```

Appendix C. Notices

This information was developed for products and services offered in the U.S.A. IBM might not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the information. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this information at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Department MU5A46
11301 Burnet Road

Austin, TX 78758
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Third-Party Statements

ICU License - ICU 1.8.1 and later

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1995-2006 International Business Machines Corporation and others

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

IBM	Tivoli	AIX	Lotus
Notes	pSeries®	DB2	WebSphere
S/390®	Domino	iNotes	Cloudscape

Java, JavaScript and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Microsoft, Windows NT and Windows are registered trademarks of Microsoft Corporation.

Intel is a trademark of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the U.S., other countries, or both.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

Other company, product, and service names may be trademarks or service marks of others.



Product Number: 5724-K74

Printed in USA

SC23-6560-00

