IBM Tivoli Directory Server

Problem Determination Guide

Version 6.3

GC27-2752-00

IBM Tivoli Directory Server

Problem Determination Guide

Version 6.3

GC27-2752-00

Note

Before using this information and the product it supports, read the general information under Appendix C, "Notices," on page 139.

This edition applies to version 6, release 3, of IBM Tivoli Directory Server and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

About this book

IBM® Tivoli® Directory Server is the IBM implementation of Lightweight Directory Access Protocol for supported Windows®, AIX®, Linux® (System x®, System i®, System p®, and System z®), Solaris, and Hewlett-Packard UNIX® (HP-UX) (Itanium®) operating systems.

IBM Tivoli Directory Server Version 6.3 Problem Determination Guide contains information about possible limitations, problems, and corrective actions that can be attempted before contacting IBM Software Support. This guide also includes information about tools you can use for determining problems with IBM Tivoli Directory Server 6.3.

Intended audience for this book

This book is intended for system administrators and directory server administrators who are responsible for maintaining and troubleshooting IBM Tivoli Directory Server.

Publications

This section lists publications in the IBM Tivoli Directory Server version 6.3 library and related documents. The section also describes how to access Tivoli publications online and how to order Tivoli publications.

IBM Tivoli Directory Server version 6.3 library

The following documents are available in the IBM Tivoli Directory Server version 6.3 library:

- *IBM Tivoli Directory Server Version 6.3 What is New for This Release*, GC27-2746-00

  Provides information about the new features in the IBM Tivoli Directory Server Version 6.3 release.
- *IBM Tivoli Directory Server Version 6.3 Quick Start Guide*, GI11-9351-00

  Provides help for getting started with IBM Tivoli Directory Server 6.3. Includes a short product description and architecture diagram, as well as a pointer to the product Information Center and installation instructions.
- *IBM Tivoli Directory Server Version 6.3 System Requirements*, SC27-2755-00

  Contains the minimum hardware and software requirements for installing and using IBM Tivoli Directory Server 6.3 and its related software. Also lists the supported versions of corequisite products such as DB2® and GSKit.
- *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide*, SC27-2747-00

  Contains complete information for installing, configuring, and uninstalling IBM Tivoli Directory Server. Includes information about upgrading from a previous version of IBM Tivoli Directory Server.
- *IBM Tivoli Directory Server Version 6.3 Administration Guide*, SC27-2749-00

  Contains instructions for performing administrator tasks through the Web Administration Tool and the command line.
- *IBM Tivoli Directory Server Version 6.3 Command Reference*, SC27-2753-00

Describes the syntax and usage of the command-line utilities included with IBM Tivoli Directory Server.

- *IBM Tivoli Directory Server Version 6.3 Server Plug-ins Reference*, SC27-2750-00

  Contains information about writing server plug-ins.

- *IBM Tivoli Directory Server Version 6.3 Programming Reference*, SC27-2754-00

  Contains information about writing Lightweight Directory Access Protocol (LDAP) client applications in C and Java™.

- *IBM Tivoli Directory Server Version 6.3 Performance Tuning and Capacity Planning Guide*, SC27-2748-00

  Contains information about tuning the directory server for better performance. Describes disk requirements and other hardware needs for directories of different sizes and with various read and write rates. Describes known working scenarios for each of these levels of directory and the disk and memory used; also suggests rough rules of thumb.

- *IBM Tivoli Directory Server Version 6.3 Problem Determination Guide*, GC27-2752-00

  Contains information about possible problems and corrective actions that can be taken before contacting IBM Software Support.

- *IBM Tivoli Directory Server Version 6.3 Messages Guide*, GC27-2751-00

  Contains a list of all informational, warning, and error messages associated with IBM Tivoli Directory Server 6.3.

- *IBM Tivoli Directory Server Version 6.3 White Pages*, SC27-2756-00

  Describes the Directory White Pages application, which is provided with IBM Tivoli Directory Server 6.3. Contains information about installing, configuring, and using the application for both administrators and users.

Related publications

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

- *Java Naming and Directory Interface™ 1.2.1 Specification* on the Sun Microsystems Web site at http://java.sun.com/products/jndi/1.2/javadoc/index.html.

  IBM Tivoli Directory Server Version 6.1 and later versions uses the Java Naming and Directory Interface (JNDI) client from Sun Microsystems. See this document for information about the JNDI client.

- 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://publib.boulder.ibm.com/tividd/td/link/tdprodlist.html

- The DB2 documentation library is located at http://www.ibm.com/software/data/db2/library/.

Accessing terminology online

The IBM Terminology Web site consolidates the terminology from IBM product libraries in one convenient location. You can access the Terminology Web site at the following Web address:

http://www.ibm.com/software/globalization/terminology

Accessing publications online

IBM posts publications for this and all other Tivoli products, as they become available and whenever they are updated, to the Tivoli Information Center Web site at http://publib.boulder.ibm.com/tividd/td/link/tdprodlist.html.

In the Tivoli Information Center window, click **Tivoli product manuals**. Click the letter that matches the first letter of your product name to access your product library. For example, click **M** to access the IBM Tivoli Monitoring library or click **O** to access the IBM Tivoli OMEGAMON® library.

IBM posts publications for this and all other Tivoli products, as they become available and whenever they are updated, to the Tivoli Documentation Central Web site at http://www.ibm.com/tivoli/documentation.

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

Ordering publications

You can order many Tivoli publications online at: http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

You can also order by telephone by calling one of these numbers:

- In the United States: 800-879-2755
- In Canada: 800-426-4968

In other countries, contact your software account representative to order Tivoli publications. To locate the telephone number of your local representative, perform the following steps:

1. Go to http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.
2. Select your country from the list and click **Go**.
3. Click **About this site** in the main panel to see an information page that includes the telephone number of your local representative.

Accessibility

Accessibility features help users with a physical disability, such as restricted mobility or limited vision, to use software products successfully. With this product, you can use assistive technologies to hear and navigate the interface. You can also use the keyboard instead of the mouse to operate all features of the graphical user interface.

Visit the IBM Accessibility Center at http://www.ibm.com/alphaworks/topics/accessibility/ for more information about the IBM commitment to accessibility.

For additional information, see the Accessibility Appendix in the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide*.

Tivoli technical training

For Tivoli technical training information, refer to the IBM Tivoli Education Web site at: http://www.ibm.com/software/tivoli/education.

Tivoli user groups

Tivoli user groups are independent, user-run membership organizations that provide Tivoli users with information to assist them in the implementation of Tivoli Software solutions. Through these groups, members can share information and learn from the knowledge and experience of other Tivoli users. Tivoli user groups include the following members and groups:

- 23,000+ members
- 144+ groups

Access the link for the Tivoli Users Group at www.tivoli-ug.org.

Support information

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides the following ways for you to obtain the support you need:

Online

Access the Tivoli Software Support site at http://www.ibm.com/software/sysmgmt/products/support/index.html?ibmprd=tivman. Access the IBM Software Support site at http://www.ibm.com/software/support/probsub.html .

IBM Support Assistant

The IBM Support Assistant is a free local software serviceability workbench that helps you resolve questions and problems with IBM software products. The Support Assistant provides quick access to support-related information and serviceability tools for problem determination. To install the Support Assistant software, go to http://www.ibm.com/software/support/isa.

For more information about resolving problems, see Appendix B, "Support information," on page 135.

Conventions used in this book

This book uses several conventions for special terms and actions, operating system-dependent commands and paths, and margin graphics.

Typeface conventions

This book uses the following typeface conventions:

Bold

- Lowercase commands and mixed case commands that are otherwise difficult to distinguish from surrounding text
- Interface controls (check boxes, push buttons, radio buttons, spin buttons, fields, folders, icons, list boxes, items inside list boxes, multicolumn lists, containers, menu choices, menu names, tabs, property sheets), labels (such as **Tip:**, and **Operating system considerations**:)
- Keywords and parameters in text

Italic

- Citations (examples: titles of books, diskettes, CDs, and DVDs)
- Words defined in text (example: a nonswitched line is called a *point-to-point line*)

- Emphasis of words and letters (words as words example: "Use the word *that* to introduce a restrictive clause."; letters as letters example: "The LUN address must start with the letter *L*.")
- New terms in text (except in a definition list): a *view* is a frame in a workspace that contains data.
- Variables and values you must provide: ... where *myname* represents....

`Monospace`

- Examples and code examples
- File names, programming keywords, and other elements that are difficult to distinguish from surrounding text
- Message text and prompts addressed to the user
- Text that the user must type
- Values for arguments or command options

Operating system-dependent variables and paths

This book uses the UNIX convention for specifying environment variables and for directory notation.

When using the Windows command line, replace **$***variable* with **%** *variable***%** for environment variables and replace each forward slash (**/**) with a backslash (\) in directory paths. The names of environment variables are not always the same in the Windows and UNIX environments. For example, %TEMP% in Windows environments is equivalent to $TMPDIR in UNIX environments.

Note: If you are using the bash shell on a Windows system, you can use the UNIX conventions.

Chapter 1. Introduction to problem determination

Problem determination, or troubleshooting, is the process of determining why a product is malfunctioning or not functioning as you expect it to. This chapter introduces problem determination as it relates to IBM Tivoli Directory Server Version 6.3.

IBM Tivoli Directory Server overview

IBM Tivoli Directory Server is the IBM implementation of Lightweight Directory Access Protocol (LDAP) for supported Windows, AIX, Linux System x, Linux System i, Linux System p, Linux System z, Solaris, and HP-UX (Itanium) operating systems. IBM Tivoli Directory Server provides a specialized directory in which to store, organize, and retrieve information about objects.

IBM Tivoli Directory Server provides diagnostic tools that can be used to collect information and determine the exact cause of problems that occur. In addition, this guide provides scenarios and workarounds dealing with such topics as installation, configuration, and replication to help you fix problems you might encounter.

Built-in troubleshooting features

IBM Tivoli Directory Server contains several tools in addition to the operating system tools to help you determine the source of problems you encounter:

Core file generation

Core files, generated by the operating system, collect the contents of a program's memory space at the time the program ended. A core file helps IBM Software Support diagnose your problem.

You must have core file generation enabled in order for core file information to be generated. See "Generating core files" on page 11 for more information about core files and for instructions for enabling core file generation.

Error logs

Error logs record error messages that occur during directory server processing. IBM Tivoli Directory Server detects and saves these errors in a text file. See Chapter 2, "Logging utilities," on page 5 for more information.

Server audit logs

Server audit logs record suspicious patterns of activity in order to detect security violations. If security is violated, the Server audit log can be used to determine how and when the problem occurred. IBM Tivoli Directory Server detects and saves these errors in a text file. See Chapter 2, "Logging utilities," on page 5 for more information.

Tools for troubleshooting IBM Tivoli Directory Server

IBM Support Assistant Lite

IBM Support Assistant (ISA) Lite is a software support solution that helps to quickly collect diagnostic files (such as logs and configuration files, schema files, and traces and core files):

- Customized to automate product specific data collection.

- Collects the data files that IBM Support analysts need to identify, diagnose, and recover from occasional operational problems with IBM products.
- Collects files automatically and package them for sending to IBM (with consent) or for your own analysis.

To get an overview and to know about the features of IBM Support Assistant Lite, see http://publib.boulder.ibm.com/infocenter/ieduasst/v1r1m0/index.jsp?topic=/com.ibm.iea.selfassist/isalite/1.3/Overview.html. To know more about IBM Supports Assistant Lite and the best practices, refer to http://www-01.ibm.com/support/docview.wss?&uid=swg27017356. To download IBM Support Assistant Lite, visit http://www-01.ibm.com/software/support/isa/download.html.

Using the Messages Guide to resolve errors

The *IBM Tivoli Directory Server Version 6.3 Messages Guide* contains a list of messages you might encounter in the IBM Tivoli Directory Server logs, graphical user interfaces, and the command line. Use the unique message ID associated with a message to locate detailed explanations and suggested operator responses in the *IBM Tivoli Directory Server Version 6.3 Messages Guide*.

For example, you encounter the following error message in the Server log:

```
Sep 13 14:31:04 2006  GLPL2D014E Suffix entry has not been created for entry
    cn=Robert Dean, ou=In Flight Systems, ou=Austin, o=sample.
```

You can search for "GLPL2D014E" in the *IBM Tivoli Directory Server Version 6.3 Messages Guide* for information about why the error occurred and how to resolve it.

DB2 error log messages, lost and found log messages, admin audit log messages, and server audit log messages are not contained in the *IBM Tivoli Directory Server Version 6.3 Messages Guide*.

Troubleshooting topics

In addition to information about built-in troubleshooting tools, this guide contains further troubleshooting information about the following topics:

- Installation and uninstallation: See Chapter 4, “Troubleshooting installation and uninstallation,” on page 21 for more information.
- Migration: See Chapter 5, “Troubleshooting migration,” on page 31 for more information.
- Instance Creation: See Chapter 6, “Troubleshooting instance creation and configuration,” on page 35 for more information.
- Configuration: See Chapter 6, “Troubleshooting instance creation and configuration,” on page 35 for more information.
- DB2: See Chapter 7, “Troubleshooting DB2,” on page 43 for more information.
- Web Administration Tool and application server: See Chapter 8, “Troubleshooting the Web Administration Tool and the application server,” on page 47 for more information.
- Replication: See Chapter 9, “Troubleshooting replication,” on page 55 for more information.
- Performance: See Chapter 10, “Troubleshooting performance,” on page 75 for more information.

- Scenarios: See Chapter 11, “Troubleshooting scenarios,” on page 81 for more information.
- General troubleshooting: See Chapter 13, “Known limitations and general troubleshooting,” on page 97 for more information.

Chapter 2. Logging utilities

IBM Tivoli Directory Server Version 6.3 provides several logs that can be viewed either through the Web Administration Tool or the system command line. See the *IBM Tivoli Directory Server Version 6.3 Administration Guide* for information about viewing the logs. See "Using the Messages Guide to resolve errors" on page 2 for information about resolving error messages that you find in the logs.

By default, all the logs listed in this section are in the *directory_server_instance_name*/logs (or *directory_server_instance_name*\logs on Windows) directory. The file names shown are the defaults, but you can change both the paths and the file names for the logs. See the *IBM Tivoli Directory Server Version 6.3 Administration Guide* for information. The IBM Tivoli Directory Server logs are:

Administration server log (ibmdiradm.log)

An administration server is a limited LDAP server that accepts searches and extended operations to stop, start, and restart the LDAP server. The administration server log allows you to view status and errors encountered by the administration server.

A sample of the log looks like this:

```
05/06/2010 02:05:57 PM GLPADM056I Admin server starting.
05/06/2010 02:05:58 PM GLPCOM025I The audit plugin is successfully loaded from
  libldapaudit.so.
05/06/2010 02:05:58 PM GLPCOM022I The database plugin is successfully loaded from
  libback-config.so.
05/06/2010 02:05:58 PM GLPADM060I The admin server backup and restore server
  configuration entry is not enabled.
05/06/2010 02:05:58 PM GLPCOM024I The extended Operation plugin is successfully
  loaded from libloga.so.
05/06/2010 02:05:58 PM GLPCOM003I Non-SSL port initialized to 3546.
05/06/2010 02:05:58 PM GLPADM028I Admin server audit logging is started.
05/06/2010 02:05:58 PM GLPADM004I 6.3.0.0       ibmdiradm started
05/06/2010 02:05:58 PM GLPSRV048I Started 5 worker threads to handle client requests.
```

Administration server audit log (adminaudit.log)

Administration server audit logging is used to improve the security of the administration server. The directory administrator and administrative group members can use the records stored in the audit log to check for suspicious patterns of activity in an attempt to detect security violations. If security is violated, the audit log can be used to determine how and when the problem occurred and perhaps the amount of damage done.

Since the Administration server is integrated with the directory server's code base, to fine grain the auditing configuration, in addition to ibm-audit, auditing is extended to include audit configuration attributes such as ibm-auditbind, ibm-auditunbind, ibm-auditExtOp, ibm-auditSearch, ibm-auditVersion, and ibm-slapdLog. For the audit configuration changes to take effect, the Administration server must receive the dynamic update configuration request or you must restart the Administration server.

Note: If any additional "MAY" attributes are specified, the server will ignore the values and no error messages will be written.

A sample of the log looks like this:

```
2010-01-15-19:59:17.130-06:00GLPADM028I Admin Server audit logging
is started.
AuditV3--2010-01-16-22:04:50.93986-06:00--V3 Bind--bindDN: CN=ROOT
  --client: 127.0.0.1:3665--connectionID: 0--received:
 2010-01-16-22:04:50.93986-06:00--Success
AuditV3--2010-01-16-22:04:50.93986-06:00--V3 Search--bindDN: CN=ROOT
  --client: 127.0.0.1:3665--connectionID: 0--received:
 2010-01-16-22:04:50.93986-06:00--Success
AuditV3--2010-01-16-22:04:50.93986-06:00--V3 Unbind--bindDN: CN=ROOT
  --client: 127.0.0.1:3665--connectionID: 0--received:
  2010-01-16-22:04:50.93986-06:00--Success
AuditV3--2010-01-16-22:08:09.94185-06:00--V3 Bind--bindDN: CN=ROOT
  --client: 127.0.0.1:3678--connectionID: 1--received:
  2010-01-16-22:08:09.94185-06:00--Invalid credentials
AuditV3--2010-01-16-22:08:09.94185-06:00--V3 Unbind--bindDN: --client:
  127.0.0.1:3678--connectionID: 1--received:
  2010-01-16-22:08:09.94185-06:00--Success
```

Server audit log (audit.log)

Audit logging is used to improve the security of the directory server. The primary directory administrator and administrative group members with AuditAdmin and ServerConfigGroupMember roles can use the activities stored in the Server audit log to check for suspicious patterns of activity in an attempt to detect security violations. If security is violated, the Server audit log can be used to determine how and when the problem occurred and perhaps the amount of damage done. This information is very useful, both for recovery from the violation and, possibly, in the development of better security measures to prevent future problems.

The Server audit log records the DN's of the Administrative Group members and their assigned roles each time the server starts and anytime their roles change. The format of the record is displayed. Records to be logged after server starts is as follows:

```
<date>-<time>--<message ID> Administrative roles assigned to <user DN>
are: <role> <role> ...
```

See the section *"Creating the administrative group" in IBM Tivoli Directory Server Version 6.3 Administration Guide* to know more about administrative roles and permissions required to access various objects.

The following is a sample of the Server audit log:

```
2010-01-16-17:38:15.484-06:00--GLPSRV023I Audit logging started.
The audit configuration options are:
 ibm-slapdLog = C:\idsslapd-ldaptest\logs\audit.log,
  ibm-auditVersion = true,ibm-audit = true,
 ibm-auditFailedOPonly = true,ibm-auditBind = true,
  ibm-auditUnbind = true,ibm-auditSearch = true,
  ibm-auditAdd = true,ibm-auditModify = true,
  ibm-auditDelete = true,ibm-auditModifyDN = true,
  ibm-auditExtOPEvent = true,ibm-auditExtOp = true,
  ibm-auditAttributesOnGroupEvalOp = true,ibm-auditCompare = true,
  ibm-auditGroupsOnGroupControl = true.
2010-01-16-17:38:15.656-06:00--GLPSRV009I IBM Tivoli Directory (SSL),
 Version 6.3 Server started.
AuditV3--2009-01-16-17:39:28.468-06:00--V3 anonymous Search--bindDN:
<*CN=NULLDN*>--client: 127.0.0.1:3792--connectionID: 1
--received: 2009-01-16-17:39:28.453-06:00-- No such object
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: cn=monitor
scope: wholeSubtree
derefAliases: neverDerefAliases
typesOnly: false
filter: (objectclass=*)
```

Bulkload error log (bulkload.log)

The **idsbulkload** (or **bulkload**) command is used to load entries. The bulkload log allows you to view status and errors related to bulkload.

For example, the command `bulkload -I ldapdb2 -i bad.ldif` was used to load entries for instance ldapdb2 from an invalid LDIF file named bad.ldif, which contained the following lines:

```
dn: cn=abc,o=sample
objectclass:person
cn:caaa
sn:abc
```

The following bulkload error log resulted:

```
04/05/09 09:31:19 GLPCTL113I Largest core file size creation limit for
         the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/09 09:31:19 GLPCTL114I Largest file size creation limit for
         the process (in bytes): '-1'(Soft limit) and'-1'(Hard limit).
04/05/09 09:31:19 GLPCTL115I Maximum data segment limit for
         the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/09 09:31:19 GLPCTL116I Maximum physical memory limit for
         the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/09 09:31:19 GLPBLK072I Bulkload started.
04/05/09 09:31:19 GLPBLK050I Extracting parent DNs ...
04/05/09 09:31:19 GLPBLK116E Invalid line detected: 3
04/05/09 09:31:19 GLPBLK044I 1 errors detected during parsing phase.
04/05/09 09:31:20 GLPBLK073I Bulkload completed.
```

Tools log (idstools.log)

The tools log contains status and error messages related to the configuration tools, such as **idscfgdb**, **idsucfgdb**, **idscfgchglog**, **idsucfgchglog**, **idscfgsuf**, **idsucfgsuf**, **idsdnpw**, **idsxcfg**, **idsxinst**, **idscfgsch**, and **idsucfgsch**.

The following is a sample of the tools log:

```
Jan 09 16:41:02 2006 GLPDPW009I Setting the directory server administrator DN.
Jan 09 16:41:02 2006 GLPDPW010I Set the directory server administrator DN.
Jan 09 16:41:02 2006 GLPDPW006I Setting the directory server administrator
 password.
Jan 09 16:41:11 2006 GLPDPW007I Set the directory server administrator
 password.
Jan 09 16:41:17 2006 GLPCDB035I Adding database 'ldaptest' to directory server
  instance: 'ldaptest'.
Jan 09 16:41:18 2006 GLPCTL017I Cataloging database instance node: 'ldaptest'.
Jan 09 16:41:19 2006 GLPCTL018I Cataloged database instance node: 'ldaptest'.
Jan 09 16:41:19 2006 GLPCTL008I Starting database manager for database
  instance: 'ldaptest'.
Jan 09 16:41:22 2006 GLPCTL009I Started database manager for database
  instance: 'ldaptest'.
Jan 09 16:41:22 2006 GLPCTL026I Creating database: 'ldaptest'.
Jan 09 16:43:11 2006 GLPCTL027I Created database: 'ldaptest'.
Jan 09 16:43:11 2006 GLPCTL034I Updating the database: 'ldaptest'
Jan 09 16:43:19 2006 GLPCTL035I Updated the database: 'ldaptest'
Jan 09 16:43:19 2006 GLPCTL020I Updating the database manager: 'ldaptest'.
Jan 09 16:43:22 2006 GLPCTL021I Updated the database manager: 'ldaptest'.
Jan 09 16:43:23 2006 GLPCTL023I Enabling multi-page file allocation:
  'ldaptest'
Jan 09 16:43:37 2006 GLPCTL024I Enabled multi-page file allocation:
  'ldaptest'
Jan 09 16:43:38 2006 GLPCDB005I Configuring database 'ldaptest' for
 directory server instance: 'ldaptest'.
Jan 09 16:43:39 2006 GLPCDB006I Configured database 'ldaptest' for
 directory server instance: 'ldaptest'.
Jan 09 16:43:39 2006 GLPCDB003I Added database 'ldaptest' to directory
 server instance: 'ldaptest'.
```

DB2 log (db2cli.log)

Database errors that occur as a result of LDAP operations are recorded in the DB2 log.

The following is a sample of the DB2 log:

```
2006-09-13-19:18:29.native retcode = -1031; state = "58031";
       message = "SQL1031N"
 The database directory cannot be found on the indicated file system.

SQLSTATE=58031

"
2006-09-13-19:18:29.native retcode = -1018; state = "E8";
       message = "SQL1018N"
 The node name "idsinode" specified in the CATALOG NODE command
already exists.

"
2006-09-13-19:18:30.native retcode = -1026; state = "C8";
       message = "SQL1026N"
 The database manager is already active.
```

Lost and found log (lostandfound.log)

The lost and found log archives entries that were replaced due to replication conflict resolution. The log of these entries allows you to recover the data in the replaced entries if necessary.

The information logged for each replaced entry includes:

- The distinguished name (DN) of the entry that is archived as a result of conflict resolution
- The type of operation that results in the conflict; for example, add or delete.
- The time the entry was created
- The time the entry was last modified
- The TCP/IP address of the supplier whose update caused the conflict
- The LDAP Data Interchange Format (LDIF) representation of the entry associated with the failed update, including all the operational attributes such as ibm-entryUUID.

The following is a sample of the lost and found log:

```
#Entry DN: cn=t6,o=ut1,c=us
#Operation type:Add
#Corrective action:Replace
#Entry createTimestamp: 20061106211242.000000Z
#Entry modifyTimestamp: 20061030202533.000000Z
#Supplier address: 9.53.21.187
dn: cn=t6,o=ut1,c=us
objectclass: person
objectclass: top
sn: aa
cn: aa
cn: t6
description: this should not be here
ibm-entryuuid: 0c4559de-0a76-4c91-96e4-5ae81d405466
```

Server log (ibmslapd.log)

The server log contains status and error messages related to the server.

The following is a sample of the server log with no errors:

```
Sep 13 14:31:04 2006  GLPL2D014E Suffix entry has not been created for
 entry cn=Robert Dean, ou=In Flight Systems, ou=Austin, o=sample.
Sep 13 14:31:04 2006  GLPRDB002W ldif2db: 0 entries have been successfully
```

```
 added out of 50 attempted.
Sep 13 14:39:41 2006  GLPCOM024I The extended Operation plugin is
 successfully loaded from libevent.dll.
Sep 13 14:39:41 2006  GLPCOM024I The extended Operation plugin is
 successfully loaded from libtranext.dll.
```

Installation and uninstallation logs

In addition, there are logs created during installation and uninstallation. The InstallShield GUI installation and uninstallation logs are: ldapinst.log, ldapuninst.log and ldaplp_inst.log (for language packs). For more information about these logs, see Chapter 4, “Troubleshooting installation and uninstallation,” on page 21.

Backup status file (dbback.dat)

The Administration Server reads entry from the server configuration file that contains backup and restore configuration details if the directory server is with RDBM back-end. If the server backup entries are not present or not enabled in the directory server instance's configuration file, then the Administration Server will log a message such as "The admin server backup and restore server configuration entry is not enabled." in the ibmdiradm.log file during the directory server startup. If the backuprestore LDAP extended operation is initiated at this stage when the backup entries are not present or not enabled, it will result in a "Protocol error" and the Administration Server will log a message such as "Unsupported extended operation request OID '1.3.18.0.2.12.81'" in the ibmdiradm.log file.

Note: If the directory server is a Proxy Server, then the backup configuration entry will not be read and the LDAP extended operation for backups and restores will not be registered.

If the entry is present and enabled, the Administration Server will check the backup location from the configuration for the date and time of current backup. Monitor searches can be used to fetch latest snapshot of the data pertaining backup/restore. The file dbback.dat is the prime source for monitor searches to fetch their data from. The dbback.dat file is created at a backup location that you specify when configuring backup, for example <backup_location>/BACKUP_FILES.

The dbback.dat file records information like "is backup configured", "database backup location", "date and time of the last backup", "is online backup configured for database and changelog", and other backup related information. This information can be very handy in troubleshooting issues. For example, if restore fails one of the reasons for failure could be that no backup image is available at the configured backup locations. This can be deduced by performing monitor searches or analyzing dbback.dat to fetch the backup information. The timestamp for the last backup is NONE in this case.

If no backup is available at the configured locations, the timestamp for the last backup will be NONE and restore requests will fail.

Note: User must not edit the contents of the dbback.dat file manually.

Chapter 3. Other diagnostic tools

Several diagnostic tools are built into IBM Tivoli Directory Server and operating systems to help users and IBM Software Support determine why a problem is occurring. This chapter describes these tools and explains how to configure and gather information from them.

Generating core files

A core file contains the contents of a program's memory space at the time the program ended. You can send core files to IBM Software Support. The information in the core file helps IBM Software Support determine the source of a server error.

To produce a core file, you must enable core file generation. After you have enabled core file generation, core files are created automatically when an error occurs. The following sections show you how to enable core file generation for your operating system.

For Windows operating systems (Dr. Watson debugger)

Windows uses a tool called Dr. Watson to generate a text file called Drwtsn32.log, which is the Windows equivalent of a core file. This file is generated whenever an error is detected.

If a program error occurs, Dr. Watson will start automatically. If you want to start Dr. Watson manually using the GUI, do the following:

1. Click **Start**.
2. Click **Run**.
3. Type **drwtsn32**.

To start Dr. Watson from a command prompt, change to the root directory, and then type **drwtsn32**.

Dr. Watson (Drwtsn32.exe) is installed in your system folder when you set up Windows. The default options are set the first time Dr. Watson runs, which can be either when a program error occurs or when you start Dr. Watson yourself. To find the location of the Dr. Watson log file, run **drwtsn32**; the **Log File Path** field will specify the path. To determine if the crash dump file will be generated, run drwtsn32 and check the status of the **Create Crash Dump File** check box.

For Linux operating systems

To enable core file generation, run the following command and then start the server from the same command line:

```
ulimit -c unlimited
ulimit -H -c unlimited
```

The ulimit for core files might be set to zero. Be sure to run these commands so that the core file size is not limited.

For AIX operating systems

To enable core file generation, run the following command and then start the server from the same command line. Be sure that the limit for the core file size is set to unlimited:

```
ulimit -c unlimited
```

For Solaris operating systems

To enable core file generation, run the following command and then start the server from the same command line:

```
coreadm -e proc-setid
```

If the application terminates unexpectedly, a core file named 'core' will be in the working directory of the process. This is true unless the global core file pattern or init core file pattern is set to a different setting. To set the file pattern to 'core' issue the following command:

```
coreadm -i core
```

To be sure that a core file is really being generated, start the **ibmslapd** process and then issue the following command :

```
"kill -6 <slapd process id>"
```

You should see a core file generated.

The ulimit for core files might be set to zero, so be sure to run the following commands so that the core file size is not limited:

```
ulimit -c unlimited
ulimit -H -c unlimited
```

To determine the current coreadm settings, run **coreadm** as root. Output such as the following will be generated:

```
global core file pattern: <setting>
init core file pattern: <setting>
global core dumps: <setting>
per-process core dumps: <setting>
global setid core dumps: <setting>
per-process setid core dumps: <setting>
global core dump logging: <setting>
```

For example:

```
global core file pattern:
init core file pattern: core
global core dumps: disabled
per-process core dumps: disabled
global setid core dumps: disabled
per-process setid core dumps: enabled
global core dump logging: disabled
```

You can disable core file generation using the following command:

```
coreadm  -d proc-setid
```

Server debug mode

If the error logs do not provide enough information to resolve a problem, you can run IBM Tivoli Directory Server in a special debug mode that generates very detailed information. You must run the server command **idsslapd** from a command prompt to enable debug output. The syntax is as follows:

```
ldtrc on
idsslapd -I <instance_name> -h <debug_mask>
```

where the specified *debug_mask* value determines which categories of debug output are generated.

Note: Running the server with the debug output option has a noticeable negative impact on performance.

After running the `ldtrc on` command, you can also use the **-d** *debug_mask* with any of the server commands except for **idsxinst** and **idsxcfg**.

You can also use the LDAP_DEBUG environment variable to specify the debug level. Set this environment variable with the value you would use for *debug_mask*.

If the LDAP_DEBUG environment variable is set and you use the **-d** option with a different debug mask, the debug mask specified with the **-d** option overrides the debug mask specified in the environment variable.

Table 1. Debug categories

Hex	Decimal	Value	Description
0x0001	1	LDAP_DEBUG_TRACE	Entry and exit from routines
0x0002	2	LDAP_DEBUG_PACKETS	Packet activity
0x0004	4	LDAP_DEBUG_ARGS	Data arguments from requests
0x0008	8	LDAP_DEBUG_CONNS	Connection activity
0x0010	16	LDAP_DEBUG_BER	Encoding and decoding of data
0x0020	32	LDAP_DEBUG_FILTER	Search filters
0x0040	64	LDAP_DEBUG_MESSAGE	Messaging subsystem activities and events
0x0080	128	LDAP_DEBUG_ACL	Access Control List activities
0x0100	256	LDAP_DEBUG_STATS	Operational statistics
0x0200	512	LDAP_DEBUG_THREAD	Threading statistics
0x0400	1024	LDAP_DEBUG_REPL	Replication statistics
0x0800	2048	LDAP_DEBUG_PARSE	Parsing activities
0x1000	4096	LDAP_DEBUG_PERFORMANCE	Relational backend performance statistics
0x1000	8192	LDAP_DEBUG_RDBM	Relational backend activities (RDBM)
0x4000	16384	LDAP_DEBUG_REFERRAL	Referral activities
0x8000	32768	LDAP_DEBUG_ERROR	Error conditions
0xffff	65535	LDAP_DEBUG_ANY	All levels of debug

For example, specifying a bit mask value of 65535 turns on full debug output and generates the most complete information.

To turn off the environment variable, use the `unset LDAP_DEBUG` command.

When you are finished, type the following command at a command prompt:

```
ldtrc off
```

Note: If you set the debug output option but tracing is off, no debug output is generated.

The generated debug output is displayed to standard error. To place the output in a file, you can do one of the following:

- Set the LDAP_DEBUG_FILE environment variable.
- On server commands (but not the **idsslapd** command), you can use the **-b** option to specify a file. If the LDAP_DEBUG_FILE environment variable is set and you use the **-b** option and specify a different file, the file you specify overrides the file specified in the environment variable.

Contact IBM Software Support for assistance with interpreting the debug output and resolving the problem.

Note: The **idsldaptrace** tracing utility can be used to dynamically activate or deactivate tracing of the directory server. See the *IBM Tivoli Directory Server Version 6.3 Command Reference* for information about the **idsldaptrace** utility.

Tracing and debugging LDAP client APIs

Before you enable tracing for LDAP client APIs, you must first stop the LDAP client application. To enable tracing consider the following steps:

1. Set the appropriate debug level using the LDAP_DEBUG environment variable.

   On AIX, Linux, Solaris, and HP-UX (Itanium) platforms
   $export LDAP_DEBUG=<debug_level>

   On Windows platform
   c:\>set LDAP_DEBUG=<debug_level>

   The different debug levels for various categories are provided in the below table.

Table 2. Debug levels

Decimal	Value	Description
1	LDAP_DEBUG_TRACE	Entry and exit from routines
2	LDAP_DEBUG_PACKETS	Packet activity
4	LDAP_DEBUG_ARGS	Data arguments from requests
8	LDAP_DEBUG_CONNS	Connection activity
16	LDAP_DEBUG_BER	Encoding and decoding of data
32	LDAP_DEBUG_FILTER	Search filters
64	LDAP_DEBUG_MESSAGE	Messaging subsystem activities and events
128	LDAP_DEBUG_ACL	Access Control List activities
256	LDAP_DEBUG_STATS	Operational statistics
512	LDAP_DEBUG_THREAD	Threading statistics
1024	LDAP_DEBUG_REPL	Replication statistics
2048	LDAP_DEBUG_PARSE	Parsing activities
4096	LDAP_DEBUG_PERFORMANCE	Relational backend performance statistics
8192	LDAP_DEBUG_RDBM	Relational backend activities (RDBM)
16384	LDAP_DEBUG_REFERRAL	Referral activities

Table 2. Debug levels (continued)

Decimal	Value	Description
32768	LDAP_DEBUG_ERROR	Error conditions
65535	LDAP_DEBUG_ANY	All levels of debug

For example, specifying a bit mask value of 65535 turns on full debug output and generates the most complete information. To know more about debug levels , see Server debug mode.

2. Set the debug file name using the LDAP_DEBUG_FILE environment variable.

   On AIX, Linux, Solaris, and HP-UX (Itanium) platforms
   $export LDAP_DEBUG_FILE=<filename>

   ON Windows platform
   c:\>set LDAP_DEBUG_FILE=<filename>

   Note: Ensure that your client application has write access to this file.

3. Run the application from the same terminal where you have environment set. Recreate the problem that you want to debug.
4. The debug information will be captured in the file pointed by the LDAP_DEBUG_FILE environment variable. You can now debug the problem using the information captured in the file, or send this file to the IBM Support team for further analysis.

Collecting an ASCII server trace at startup

Collecting an ASCII server trace helps to determine and debug issues involving a failed directory server startup or to trace a specific operation at directory server startup. To collect an ASCII server trace, perform the following steps

1. Stop the directory server instance, if running. Issue the command of the following format:

   ```
   ibmslapd -I <instance_name> -k
   ```

2. Determine if tracing is enabled or not. Issue the following command:

   ```
   ldtrc info
   ```

3. Enable tracing if it is currently disabled (in "off" mode). Issue the following command:

   ```
   ldtrc on
   ```

4. Start the directory server in DEBUG mode and redirect the output to a file.

   On AIX, Linux, and Solaris platforms
   Issue the command of the following format:

   ```
   ibmslapd –I <instance_name> –n -h 65535 2>&1 | tee /tmp/slapd_trace.out
   ```

   On Windows platform
   Issue the command of the following format:

   ```
   (ibmslapd –I <instance_name> -n  -h 65535 2>&1) > C:\slapd_trace.out
   ```

5. Recreate the problem. Once the error or the condition you want to trace occurs and the screen no longer has messages written out, press Ctrl + C to stop the process. Now, you can analyze the trace file.
6. Disable tracing. Issue the following command:

   ```
   ldtrc off
   ```

Collecting a binary server trace at startup

To debug issues involving a failed directory server startup or to trace a specific operation at directory server startup, collecting a binary server trace is important. To collect a binary server trace at startup, do the following:

1. Stop the directory server instance, if running. Issue the command of the following format:

   ```
   ibmslapd -I <instance_name> -k
   ```

2. Determine if tracing is enabled or not. Issue the following command:

   ```
   ldtrc info
   ```

3. If trace is currently enabled, disable the trace. Issue the following command:

   ```
   ldtrc off
   ```

4. Enable binary tracing. Issue the following command:

   ```
   ldtrc on -l 50000000
   ```

   In the command, the value for the buffer size is set to 50 million bytes. This will store the latest 50 million bytes of trace record data in the shared memory (it flushes the oldest data once the 50 MB value is reached). If for some reason the command fails due to what may appear to be not enough shared memory resources, you can scale the number down but less than 20 million might not provide the desired information.

5. Start the directory server instance. Issue the command of the following format:

   ```
   ibmslapd -I <instance_name> -n
   ```

6. Point the environment variable TRCTFIDIR to the <TDS_INSTALL_HOME> directory. To do this, use the following command:

   - On AIX, Linux, and Solaris platforms: **$export TRCTFIDIR=<TDS_INSTALL_HOME>/etc**.
   - On Windows platform: **C:\> set TRCTFIDIR=<TDS_INSTALL_HOME>\etc**.

   where, <TDS_INSTALL_HOME> on different operating system is as follows:

   - On AIX and Solaris: /opt/IBM/ldap/V6.3/etc
   - On Linux: /opt/ibm/ldap/V6.3/etc
   - On Windows: <Install_Drive>:\Program Files\IBM\LDAP\V6.3\etc

7. Recreate the problem to produce the error or condition that you want to trace.
8. Collect the trace records. Once the error or the condition you want to trace occurs, issue the following command:

   ```
   ldtrc dump trace.raw
   ```

   where, trace.raw is the path name and file name that will be used to capture the records in shared memory.

9. Change to the <TDS_INSTALL_HOME>/etc directory and then collect the format and flow of the binary trace. Issue the following commands:

   ```
   ldtrc fmt trace.raw trace.fmt
   ldtrc flw trace.raw trace.flw
   ```

   Send the trace.fmt and trace.flw files to support.

10. Disable tracing. Issue the following command:

    ```
    ldtrc off
    ```

Collecting performance records dynamically

The performance profile information in trace is intended to help users diagnose performance problems. By using the independent trace facility, performance profiling is accomplished with minimum impact on server performance. The independent trace facility profiles operation performance that consists of timestamps at key points traversed during an operation execution for a running server instance. The timestamps are profiled during different stages such as the following:

- RDBM search processing
- RDBM bind processing
- RDBM compare processing
- RDBM write processing

To activate tracing of performance records dynamically, do the following:

1. Activate tracing for performance records. To do this, issue the following command:

   ```
   ldaptrace –h <hostname> -p <port number> -D <adminDN> -w <adminPW> -l on \
       –t start -- -perf
   ```

2. Dump the trace to a binary trace file. To do this, issue the following command:

   ```
   ldtrc dump trace.bin
   ```

3. Format the trace. To do this, issue the following command:

   ```
   ldtrc fmt trace.bin trace.txt
   ```

After formatting the trace you can analyze the trace and diagnose performance problems. To turn off tracing, issue the following command:

```
ldtrc off
```

To know more about "Performance profiling", see the *IBM Tivoli Directory Server Version 6.3 Administration Guide*.

Collecting a dynamic ASCII server trace

Collecting a dynamic ASCII server trace helps in debugging issues related to a specific operation of a server. You can collect a dynamic server trace only if the Tivoli Directory Server instance that you wish to debug is running. To collect a dynamic server trace, do the following:

1. Verify the ports used by your directory server instance. Issue the following command:

   ```
   idsilist -I <instance_name> -a
   ```

2. Start the dynamic ASCII server trace for your directory server instance.

   On AIX, Linux, and Solaris platforms

   Issue the command of the following format:

   ```
   idsldaptrace -p <port> -a <admin_port> –D <adminDN> –w <adminPW> \
       –h <hostname> –l on -t start –m 65535 –o /tmp/ibmslapd.dbg
   ```

   On Windows platforms

   Issue the command of the following format:

   ```
   idsldaptrace -p <port> -a <admin_port> –D <adminDN> –w <adminPW> \
      –h <hostname> –l on -t start –m 65535 –o C:\Temp\ibmslapd.dbg
   ```

3. Recreate the problem and issue the specific operation that is failing.
4. Disable the dynamic ASCII server trace. Issue the command of the following format:

```
idsldaptrace -p <port> -a <adminPort> –D <adminDN> –w <adminPW> \
         –h <hostname> –l off -t stop
```

Note: You can issue **idsldaptrace -?** to see the usage information for the command.

Collecting trace information

Enabling tracing is a multistep process that involves starting the trace facility which allows tracing of directory server and other commands. Once the trace facility is enabled, an administrator can request for a specific processes, like directory server, or commands, like ldif2db , to print trace information. Trace information can be sent to the command line or to a file. To enable tracing perform the following steps:

1. Enable the trace facility. From the command line, issue the following command:

   ```
   ldtrc on
   ```

   OR,

   ```
   idsldaptrace -p <adminSereverPort> -h <host_name> -D cn=<adminDN> \
     -w <adminPW> -l on
   ```

   Note: You can use the idsldaptrace command from any system that has the directory server installed. The Admininstration Server must be running for this command to work.

2. Enable the tracing for a specific process or a command. Select a debug level for the trace. For example, specifying a bit mask value of 65535 turns on full debug output and generates the most complete information. To know more about debug levels , see Server debug mode. You can use one of the following options to set the debug level based on the process or command that you want to trace..

   - Set the LDAP_DEBUG environment variable to specify the debug level. Set this environment variable with a value that you want to use for debug_mask. If the LDAP_DEBUG environment variable is set and you use the -d option with a different debug mask, the debug mask specified with the -d option overrides the debug mask specified in the environment variable.

     On AIX, Linux, and Solaris platforms: **$export LDAP_DEBUG=<debug_level>**.

     On Windows platform: **C:\> set LDAP_DEBUG=<debug_level>**. To disable the environment variable, use the **unset LDAP_DEBUG** command.

   - For a directory server instance, you can enable tracing at the server startup by setting the attributes in the server configuration file. To do this, set the ibm-slapdStartupTraceEnabled attribute to TRUE in the server configuration file. There are configuration options for setting the level using the ibm-slapdTraceMessageLevel attribute and routing the output to a file by specifying a file name as value for the ibm-slapdTraceMessageLog attribute. The following example shows the ibm-slapdStartupTraceEnabled attribute set to true in the cn=Configuration entry:

     ```
     idsldapmodify -p <port> -D cn=<adminDN> -w <adminPW>
     dn: cn=Configuration
     changetype: modify
     replace: ibm-slapdStartupTraceEnabled
     ibm-slapdStartupTraceEnabled: TRUE
     -
     replace: ibm-slapdTraceMessageLevel
     ibm-slapdTraceMessageLevel: 0xFFFF
     -
     replace: ibm-slapdTraceMessageLog
     ```

```
ibm-slapdTraceMessageLog: /var/ibmslapd.trace.log

Operation 0 modifying entry cn=Configuration
```

Restart the directory server instance for the changes to take effect.

Note: To disable tracing modify the value of the ibm-slapdStartupTraceEnabled attribute to False using the idsldapmodify command.

- You can also dynamically enable tracing after a directory server instance has started using the idsldaptrace command.

  To start tracing IBM Tivoli Directory Server, issue the idsldaptrace command of the following format:

  ```
  idsldaptrace -h <host_name> -D cn=<adminDN> -w <adminPW> -p <port> \
   -m <debug_level> -o <output_file> -t start
  ```

  To stop tracing of IBM Tivoli Directory Server, issue the idsldaptrace command of the following format:

  ```
  idsldaptrace -h <host_name> -D cn=<adminDN> -w <adminPW> -p <port> -t stop
  ```

3. When you are finished with tracing you must disable tracing. You can use one of the following options to stop tracing depending on the method that you have used to enable tracing.
   - To stop tracing, issue the following command:

     ```
     ldtrc off
     ```

     OR,

     ```
     idsldaptrace -p <adminSereverPort> -h <host_name> -D cn=<adminDN> \
       -w <adminPW> -l off
     ```

     Note: You can use the ldaptrace command from any system that has the directory server installed. The Administration Server must be running for this command to work.

Alternatively, you can also use the GUI, Web administration tool, to start or stop tracing. To do this, select **Logs** under **Server administration** in the Web administration navigation area. On the expanded list, select **Start/Stop server trace** to enable or disable server tracing. The GUI has fields Trace debug levels and Trace debug file, where you can specify the debug level and output file to store trace information. If you use the GUI, it will take care of starting and stopping the trace facility. To know more about logging utilities, see *IBM Tivoli Directory Server Version 6.3 Administration Guide*.

Collecting IBM Tivoli Directory Server's log and configuration file

In order to diagnose any issue related to Tivoli Directory Server, it is important to collect the log and configuration file. The log and configuration file that you usually check to determine issues related to Tivoli Directory Server are located in the following path for 6.3 version.

On AIX, Linux, and Solaris platforms

- Configuration file - <instance_home>/idsslapd-<instance_name>/etc/ibmslapd.conf
- Administration Server log file - <instance_home>/idsslapd-<instance_name>/logs/ibmslapd.log

- DB2 error log - <instance_home>/idsslapd-<instance_name>/logs/db2cli.log
- Audit log file - <instance_home>/idsslapd-<instance_name>/logs/audit.log

You can issue the **idsilist -a** command at the command line to view the <*instance_home*> and directory server instance names, <*instance_name*>, on a given computer.

On Windows platforms

- Configuration file - <install_path>\idsslapd-<instance_name>\etc\ibmslapd.conf
- Administration Server log file - <install_path>\idsslapd-<instance_name>\logs\ibmslapd.log
- DB2 error log - <install_path>\idsslapd-<instance_name>\logs\db2cli.log
- Audit log file - <install_path>\idsslapd-<instance_name>\logs\audit.log

You can issue the **idsilist -a** command at the command line to view the <*install_path*> and directory server instance names, <*instance_name*>, on a given computer.

Chapter 4. Troubleshooting installation and uninstallation

There are many points during the installation of a product and its prerequisite software where problems might be encountered. This chapter explains how to troubleshoot problems during the installation process and perform recovery actions.

Product installation overview

When you install IBM Tivoli Directory Server, you can install the following components:

- Client SDK
- Java client
- Server
  - IBM Tivoli Directory Proxy Server
  - IBM Tivoli Directory Full (Backend Enabled) Server
- Web Administration Tool
- Embedded WebSphere® Application Server
- IBM DB2 (Enterprise Server Edition / Workgroup Server Edition)
- Global Security Kit (GSKit)

You can install these components using an InstallShield graphical user interface (GUI) or use operating-system-specific installation methods such as the command line or installation tools for the operating system.

Prerequisite software

If you are installing using the InstallShield GUI, prerequisite software is available for installation as part of the IBM Tivoli Directory Server overall installation process. If you are using the operating system utilities to install, installation might fail if you do not have the prerequisite software installed. Before you install, be sure to read the "System requirements and supported software versions" section in the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide*.

If installation does not complete, the first place you can look for information is the ldapinst.log file. If the installation destination directory (*install_directory*) was created, this log is in the following location:

- On Windows, in *install_directory*\var. For example, if you installed in the default location, the ldapinst.log file and other install log files are in c:\Program Files\IBM\LDAP\V6.3\var.
- On AIX, Linux, and Solaris systems, in *install_directory*\var. Other install logs are in the /var/idsldap/V6.3 directory.

If *install_directory* was not created before the installation failed, the log might be in a temporary directory. To find it, search for "ldapinst.log". Review this log for any messages about why the installation failed.

If you are installing language packs using the InstallShield GUI, the installation log is in the *install_directory*\LangPack\ldaplp_inst.log file on Windows systems or in *install_directory*/LangPack/ldaplp_inst.log on AIX, Linux, and Solaris systems.

Because some of the LDAP features require corequisite products, it is possible that a failure in a corequisite installation caused the IBM Tivoli Directory Server installation to fail. For example, when IBM Tivoli Directory Full Server is being installed but if the DB2 installation fails, then IBM Tivoli Directory Full Server cannot be installed.

Failure when installing IBM Tivoli Directory Server with corequisite software

If a failure occurs while you are installing prerequisite software, you will see different results depending on the software you are installing when the failure occurs and other related components you are installing. For example:

- If you are installing IBM Tivoli Directory Server and the installation fails, an "Installation cannot continue" message is displayed and the installation exits. This can be because the corequisite software could not be located, for example IBM Tivoli Directory Full Server cannot be installed without DB2. This is considered to be a critical failure and a reason for exiting the installation completely. The reason for not able to locate corequisite software can vary, but the result is the same. Examples of reasons for the failure are:
  - If installing from downloaded and uncompressed .zip, .tar, or .iso files:
    - If you downloaded .zip files on your computer, uncompress the files to your computer into a path that has no spaces in the name. Uncompress all .zip files in the same directory.
    - If you downloaded .tar file on your computer, uncompress all .tar files in the same directory.
    - The .iso file versions of the product are used to burn installation DVDs that can then be used in the installation process. The .iso files are images that must be processed through a DVD burner program to create DVDs. When you create the DVDs, be sure that you do not make data DVDs of the .iso files. Select the option that unencapsulates the data from the .iso files and burns the files on the DVD.

  In these cases the installation exits prematurely, and the installation log is stored in a temporary location and not in the installation path. On Windows systems, the installation log file is usually stored in the *C:\Documents and Settings\Administrator\Local Settings\Temp* directory. On AIX, Linux, and Solaris systems, the installation log file is stored in the */tmp* directory. In order for install to find the image correctly, users must download and uncompress or untar all the images into the same directory.

To know more about installing IBM Tivoli Directory Server version 6.3, see *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide*.

The idsldap user and group

During installation of a server, the idsldap user and group are created if they do not already exist. If your AIX, Linux, or Solaris environment requires that you have more control over this user and group, you can create them before you install. The requirements are:

- The idsldap user must be a member of the idsldap group.
- The root user must be a member of the idsldap group.
- The idsldap user must have a home directory.
- The default shell for the idsldap user must be the Korn shell.
- The idsldap user can have a password, but is not required to.
- The idsldap user can be the owner of the director server instance.

If you do not want the installer to automatically create the idsldap user and group, you can use the following commands to create them and set them up correctly:

On AIX systems:

Use the following commands.

To create the idsldap group:

```
mkgroup idsldap
```

To create user ID idsldap, which is a member of group idsldap, and set the korn shell as the default shell:

```
mkuser pgrp=idsldap home=/home/idsldap shell=/bin/ksh idsldap
```

To set the password for user idsldap:

```
passwd idsldap
```

To modify the root user ID so that root is a member of the group idsldap:

```
/usr/bin/chgrpmem -m + root idsldap
```

On Linux systems:

Use the following commands.

To create the idsldap group:

```
groupadd idsldap
```

To create user ID idsldap, which is a member of group idsldap, and set the korn shell as the default shell:

```
useradd -g idsldap -d /home/idsldap -m -s /bin/ksh idsldap
```

To set the password for user idsldap:

```
passwd idsldap
```

To modify the root user ID so that root is a member of the group idsldap:

```
usermod -G idsldap,rootgroups root
```

where *rootgroups* can be obtained by using the command: `groups root`

On Solaris systems:

Use the following commands.

To create the idsldap group:

```
groupadd idsldap
```

To create user ID idsldap, which is a member of group idsldap, and set the korn shell as the default shell:

```
useradd -g idsldap -d /export/home/idsldap -m -s /bin/ksh idsldap
```

To set the password for user idsldap:

```
passwd idsldap
```

To modify the root user ID so that root is a member of the group idsldap, use the AdminTool or another appropriate tool.

Be sure that all these requirements are met before you install. IBM Tivoli Directory Proxy Server does not install correctly if the idsldap user exists but does not meet the requirements.

Installation logs

The following sections describe logs used during installation by the InstallShield GUI.

Logs for Embedded WebSphere Application Server

Logs used by the InstallShield GUI when installing Embedded WebSphere Application Server are:

On Windows platforms

- <*install_home*>\var\installApp.log
- <*install_home*>\var\installAppErr.log
- <*install_home*>\var\configApp.log
- <*install_home*>\var\configAppErr.log
- <*install_home*>\var\migrateApp.log
- <*install_home*>\var\migrateAppErr.log

The following logs are used when adding Embedded WebSphere Application Server Web Administration tool as a Windows service.

- addWebAdminSrv.log
- addWebAdminSrvErr.log

The following logs are used when starting Embedded WebSphere Application Server Web Administration tool as a Windows service.

- startWebAdminSrv.log
- startWebAdminSrvErr.log

On AIX, Linux, and Solaris platforms

- /var/idsldap/V6.3/installApp.log
- /var/idsldap/V6.3/installAppErr.log
- /var/idsldap/V6.3/configApp.log
- /var/idsldap/V6.3/configAppErr.log
- /var/idsldap/V6.3/migrateApp.log
- /var/idsldap/V6.3/migrateAppErr.log

where *install_home* is the location where you installed IBM Tivoli Directory Server.

DB2 logs on Windows

Logs used by the InstallShield GUI when installing and uninstalling DB2 on Windows are:

When installing

- <*install_home*>\var\DB2setup.log
- <*install_home*>\var\db2inst.log
- <*install_home*>\var\db2insterr.log
- <*install_home*>\var\db2wi.log

Note: Sometimes, the db2wi.log file is located in the temporary directory instead of <*install_home*>. The temporary directory is whatever the temp environment variable is set to, for example, \Documents and Settings\<userid>\Local Settings\temp.

When uninstalling

The directory is whatever the temp environment variable is set to, usually \Documents and Settings\<userid>\Local Settings\temp and then the files are:

- DB2remove.log
- db2uninst.log
- db2uninsterr.log
- DB2UninstTrc.log

DB2 logs on AIX, Linux, and Solaris systems

Logs used when installing DB2 on AIX, Linux, and Solaris systems are:

When installing using the InstallShield GUI

- /var/idsldap/V6.3/db2inst.log
- /var/idsldap/V6.3/db2insterr.log
- /var/idsldap/V6.3/DB2setup.log

When uninstalling

- /var/idsldap/V6.3/db2uninst.log
- /var/idsldap/V6.3/db2uninsterr.log

When installing using the db2_install utility

- /tmp/db2_install.rc.99999
- /tmp/db2_install.log.99999

idslink log on AIX, Linux, Solaris, and HP-UX (Itanium) systems

The **idslink** script should be run manually during InstallShield GUI and operating system utility installation of the client, IBM Tivoli Directory Proxy Server and IBM Tivoli Directory Full Server. The idslink.log and idslink.preview files are located in the /var/idsldap/V6.3/ directory.

GSKit logs on Windows operating systems

Logs used by the InstallShield GUI when installing and uninstalling GSKit on Windows systems are:

- <*install_home*>\var\gsksetup.log
- <*install_home*>\var\gskitinst.log
- <*install_home*>\var\gskitinsterr.log

Log files generated for native packages on AIX, Linux, Solaris, and HP-UX operating systems

On AIX, Linux, Solaris, and HP-UX (Itanium) platforms, two logs are generated for each native package that is installed. These logs give information about the native packages. The log files are created in the /var/idsldap/V6.3 directory. Users can refer to these logs to determine the reason why an install failed. These log files are of importance since InstallShield installs the native packages in the background.

Note: On HP-UX (Itanium) systems, the logs are created only for the client package (Client, Java client, GSKit) that is provided.

The various log files that are created during install are:

- baseServerErr.log, baseServer.log
- clientXXBitErr.log, clientXXBit.log

  Note: Here, XX can be either 64 or 32 depending on whether the hardware is 64-bit or 32-bit.
- clientBaseErr.log, clientBase.log
- engMsgErr.log, engMsg.log
- gsKitErr.log, gsKit.log
- javaClientErr.log, javaClient.log
- proxyErr.log, proxy.log
- serverErr.log, server.log
- srvBaseErr.log, srvBase.log
- webAdminErr.log, webAdmin.log

On AIX systems, additional logs are generated for SSL packages. The log files are created in the /var/idsldap/V6.3 directory.

- srvBaseMaxCrypto.log
- srvBaseMaxCryptoErr.log
- webAdminMaxCrypto.log
- webAdminMaxCryptoErr.log
- client64MaxCrypto.log
- client64MaxCryptoErr.log

On AIX, Linux, Solaris, and HP-UX (Itanium) platforms, when native packages are uninstalled log files are created in the /var/idsldap/V6.3/uninstall directory. The various log files that are created depending on the native packages that you install are:

- baseServer.log, baseServerErr.log
- baseSrv.log, baseSrvErr.log
- client64Bit.log, client64BitErr.log
- clientBase.log, clientBaseErr.log
- engMsg.log, engMsgErr.log
- entitle.log, entitleErr.log
- Gskit.log, GskitErr.log
- javaClient.log, javaClientErr.log
- proxy.log, proxyErr.log
- server.log, serverErr.log
- webAdmin.log, webAdminErr.log

On AIX systems, in addition to the above log files that are created for native packages during uninstall the following additional log files are created.

- baseSrvMaxCrypto.log, baseSrvMaxCryptoErr.log
- clientuninst64MaxCrypto.log, clientuninst64MaxCryptoErr.log
- webAdminMaxCryptouninst.log, webAdminMaxCryptouninstErr.log

Troubleshooting

If you are having problems installing IBM Directory Server, refer to the following sections for possible fixes.

InstallShield GUI installation

The following items relate to InstallShield GUI installation.

Installation failure due to lack of disk space

One reason for an installation failure is lack of disk space. IBM Tivoli Directory Server attempts to verify that there is enough space and generates messages if the required disk space is not found, but sometimes the InstallShield GUI cannot progress far enough to issue a message. Before installing, make sure you have the required free disk space available that is specified in the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide*. All platforms use temporary space. In addition, AIX, Linux, and Solaris platforms use the `/var` directory. When installation is first run, the JVM is installed to the installation directory, so be sure that your installation destination directory has enough space.

Recovering from a failed installation

The first step to recovering from a failed installation is to run the InstallShield Uninstall GUI to clean up any registry entries that might have been made by the installation process. If you do not run the InstallShield Uninstall GUI, the InstallShield GUI might fail the next time you try to use it to install IBM Tivoli Directory Server. See the following sections for information organized by operating system. See the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* for information about uninstalling using the InstallShield GUI.

When installing on AIX, Linux, and Solaris platforms, the InstallShield GUI uses the native packages (for example, AIX installp files, Solaris .pkg files, or Linux RPM files) to install IBM Tivoli Directory Server. Because of this, you will see these packages when you run the platform commands (such as `rpm -qa` on the Linux operating system) to query what is installed. Even though you can use the platform commands (such as `rpm -e`) to uninstall, you **must** use the InstallShield GUI to uninstall so that the InstallShield Registry is cleaned up.

Windows operating systems: To recover from a failed InstallShield GUI installation on Windows systems:

1. Correct any problems listed in the ldapinst.log file. See "Prerequisite software" on page 21 for more information about the ldapinst.log file.
2. Uninstall IBM Tivoli Directory Server using the InstallShield GUI. See the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* for more information about uninstalling IBM Tivoli Directory Server.
3. Remove the IBM Tivoli Directory Server installation directory. The default directory is C:\Program Files\IBM\LDAP\V6.3.
4. Use **regedit** to remove the LDAP entry in the registry: HKEY_LOCAL_MACHINE\SOFTWARE\IBM\IDSLDAP\6.3

AIX operating systems: To recover from a failed InstallShield GUI installation on AIX systems:

1. Correct any problems listed in the ldapinst.log file. See "Prerequisite software" on page 21 for more information about the ldapinst.log file.
2. Uninstall IBM Tivoli Directory Server using the InstallShield GUI. See the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* for more information about uninstalling IBM Tivoli Directory Server.
3. Type the following at a command prompt:

   ```
   lslpp -l |grep -i idsl
   ```

4. If any packages that were installed by IBM Tivoli Directory Server were left on the system, use **installp** to uninstall them, as follows:

```
installp -u packagename
```

See the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* for information about package names for IBM Tivoli Directory Server.

5. Remove the /opt/IBM/ldap/V6.3 directory.

Linux operating systems: To recover from a failed InstallShield GUI installation on Linux systems:

1. Correct any problems listed in the ldapinst.log file. See "Prerequisite software" on page 21 for more information about the ldapinst.log file.
2. Uninstall IBM Tivoli Directory Server using the InstallShield GUI. See the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* for more information about uninstalling IBM Tivoli Directory Server.
3. Type the following at a command prompt:

   ```
   rpm -qa | grep -i idsl
   ```

   If any packages that were installed by IBM Tivoli Directory Server were left on the system, use the **rpm** command to uninstall them. For example:

   ```
   rpm -ev  packagenames
   ```

   See the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* for information about package names for IBM Tivoli Directory Server.
4. If an **rpm** command hangs, try running the command with the **noscripts** option:

   ```
   rpm -ev  --noscripts packagenames
   ```
5. Remove the /opt/ibm/ldap/V6.3 directory.

Solaris operating systems: To recover from a failed InstallShield GUI installation on Solaris systems:

1. Correct any problems listed in the ldapinst.log file. See "Prerequisite software" on page 21 for more information about the ldapinst.log file.
2. Uninstall IBM Tivoli Directory Server using the InstallShield GUI. See the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* for more information about uninstalling IBM Tivoli Directory Server.
3. Type the following at a command prompt:

   ```
   pkginfo | grep -i idsl
   ```
4. If any packages that were installed by IBM Tivoli Directory Server were left on the system, use **pkgrm** to uninstall them:

   ```
   pkgrm packagenames
   ```

   See the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* for information about package names for IBM Tivoli Directory Server.

   Note: If you encounter problems removing these packages, try to remove the directories containing the packages from /var/sadm/pkg
5. Remove the /opt/IBM/ldap/V6.3 directory, and any other directories left from the installation, such as a language directory.

Missing files after server installation

After an InstallShield GUI installation on AIX, Linux, or Solaris systems, if there are files missing such as **idsxinst**, **idsicrt**, or **idsilist**, IBM Tivoli Directory Proxy Server feature might not have installed correctly. (You might notice this problem when instance creation begins because the Instance Administration Tool is not available.)

If you experience this situation:

1. Type `id idsldap` at a command prompt.
2. If the results do not show that the idsldap user is a member of the idsldap group, do one of the following:
   - Modify the idsldap user so that it belongs to the idsldap group.
   - Delete the idsldap user and the idsldap group and then do one of the following:
     – Recreate the idsldap user and group as described in the section The idsldap user and group.
     – Do not recreate the idsldap user and group, but let IBM Tivoli Directory Server installation recreate them (when you do step 3.)
3. Reinstall the base server package, and the server or proxy server packages depending on what type of server is required.

The base server package gets installed with IBM Tivoli Directory Proxy Server or with IBM Tivoli Directory Full Server package. If user is using InstallShield GUI, this is a hidden feature and is installed for the user whenever they choose IBM Tivoli Directory Proxy Server or IBM Tivoli Directory Full Server.

Note: You do not need to define the idsldap user and group before installation. If they do not exist, the base server installation creates the idsldap user and group.

Operating system utility uninstallation after InstallShield GUI installation

The *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* instructs you to use the InstallShield GUI to uninstall the IBM Tivoli Directory Server if the InstallShield GUI was used to install. If, however, you perform an operating system uninstallation after an InstallShield GUI installation, you must clean up any registry entries that might have been made by the installation process. For instructions for cleaning up the registry entries, see “Recovering from a failed installation” on page 27.

If default instance creation fails during the Typical installation

When using the Typical installation, if the default directory server instance fails to get created or if an error occurs, user must check the ldapinst.log file in the <*install_location*>/var directory to debug the problem.

On Windows operating system, installation might fail giving out message such as “DB2 Install was NOT successful”

If you are creating a new user ID and your system has "Password must meet complexity requirements" enabled, be sure that the password you supply meets the complexity requirements. If it does not, installation will fail. See the Windows documentation for information about password complexity requirements.

Here is an extract of Windows password complexity requirements from Windows Help:

Password must meet complexity requirements

Description: This security setting determines whether passwords must meet complexity requirements.

If this policy is enabled, passwords must meet the following minimum requirements:

- Not contain all or part of the user's account name

- Be at least six characters in length
- Contain characters from three of the following four categories:
  - English uppercase characters (A through Z)
  - English lowercase characters (a through z)
  - Base 10 digits (0 through 9)
  - Non-alphabetic characters (for example, !, $, #, %)
- Complexity requirements are enforced when passwords are changed or created

If the password complexity requirements are not met the installation will fail and logs messages to DB2setup.log located in the <*install_location*>\var directory. Here is an extract from the log:

```
Found echo string in C:\Program Files\IBM\LDAP\V6.3\var\db2inst.log file.
returnCode from DB2 Install is set to: 87
Return Code from DB2 install is: 87
Found failing return code in db2inst.log file.
DB2 Install was NOT successful.
```

To resolve this problem, set the user password to meet the Windows password complexity requirement and try again.

When installing using InstallShield GUI on an AIX system, native install packages might not get installed

On an AIX system, GSKit should be installed before the client or base server "max_crypto" packages are installed. If GSkit is not installed, then native install packages, such as max_crypto, for client and server might not get installed and features, such as SSL, cannot be used.

InstallShield GUI uninstallation

The following items relate to InstallShield GUI uninstallation.

Product directories still exist after uninstallation

If the *installationpath*/_uninst and *installationpath*/_jvm directories still exist and you think you have successfully uninstalled all features, run the InstallShield GUI uninstallation again and select the **Product Uninstallation** check box to remove the product completely. This should remove the _uninst and _jvm subdirectories.

Chapter 5. Troubleshooting migration

Migration refers to the process of installing IBM Tivoli Directory Server version 6.3 to replace an earlier version while preserving changes that were made to the data, schema definitions, and directory server configuration from the earlier version. The following sections contain troubleshooting information for migration.

Migration log files

Check the following log files for information about migration processes:

On AIX, Linux, and Solaris platforms:
Errors that occurred during migration are logged in the /var/idsldap/V6.3/idsadm.log file.

On Windows platforms:
Errors that occur during migration are logged in the *install_directory*\var\idsadm.log file.

Kerberos service name

In Tivoli Directory Server V6.1 and later versions of directory server look for **ldap** in the keytab file in which an **LDAP** service name was located and used by the previous versions of server. To avoid any problem arising because of this, user can do either of the following:

- Generate a keytab file by adding a lower case LDAP Kerberos service name and start using the new keytab file to communicate.
- Set the environment variable LDAP_KRB_SERVICE_NAME. If this environment variable is set, then the Kerberos service name is used with “LDAP” server service name in the keytab file and to communicate with its clients. The environment variable needs to be set on the client side as well to continue using the upper case LDAP service name to communicate with its server. If the environment variable LDAP_KRB_SERVICE_NAME is not set, then the Kerberos service name is used with "ldap".

Database instance or database in configuration file but no longer on system

If you are using the Instance Administration Tool to migrate and there is an ibm-slapdDbInstance or ibm-slapdDbName attribute in your backed-up configuration file, but that DB2 instance or database no longer exists on the system, you are not allowed to continue with migration. You receive an error message stating that the database instance or database is not present and migration cannot continue.

To recover from this problem, do one of the following:

- Comment out the database information from the configuration file and migrate using the Instance Administration Tool.
- Use the **idsimigr** command-line utility for migration. When you use this command-line tool, if the database instance from the ibm-slapdDbInstance attribute is no longer on the system, the information in the configuration file is ignored and information for a new database instance is inserted instead.

If it was the database that could not be found, the information is removed from the configuration file. You must then run **idscfgdb** to configure a database.

Format of backed-up schema files incorrect

If you receive an error at server startup that references definitions in the V3.modifiedschema file, verify that the format of the backed-up schema files is correct. For example, a newline in the middle of a definition in the V3.modifiedschema file from a previous release might result in incorrect definitions in the migrated V3.modifiedschema file.

ibm-slapdPlugin entry in configuration file changed

If a line in the ibmslapd.conf or slapd32.conf file for the ibm-slapdPlugin has been changed from its original form, it might be left in the migrated configuration file and cause an error at server startup. For example, the line in the original configuration file was:

```
ibm-slapdPlugin: database    /lib/libback-rdbm.so rdbm_backend_init
```

and the line was changed to:

```
ibm-slapdPlugin: database    /usr/ldap/lib/libback-rdbm.so rdbm_backend_init
```

The line in the second example is not removed by the migration tool and the server will not be able to load /usr/ldap/lib/libback-rdbm.so at startup because the path is not a valid path.

Considerations when performing migration

When a user intend to migrate from a previous version of Tivoli Directory Server to a current version, certain conditions that the user must consider before migrating are:

- Specify a valid backup directory, encryption key, and encryption salt for a instance that you intend to migrate.
- Ensure that the required files are available in the backup directory.
- The source and target versions of Tivoli Directory server are supported versions for migration.
- Ensure that the ldapdb.properties file in the <TDS_instancehome>/etc is present and valid.

If during migration error conditions are encountered, then the tool will attempt to restore schema and configuration files to their original state and will exist by displaying appropriate error messages. Some of the likely causes for exit of migration are:

- If parsing of schema or configuration files fail.
- On Windows, if the migration of services files fail.

To identify the reasons for migration failure, user should examine the following:

- The error messages that are displayed on console when migration fails, user can use this to determine the cause of failure.
- Run migration with trace mode ON and redirect the trace output to a file. This trace output can be used by the user to determine the reason for migration failure.

- If an ibm-slapdInvalidLine message is displayed during server startup, then user should check if the configuration file is corrupted.
- The idsdbmigr.log file should be checked for any failure in database migration.
- If the server fails to start after migration, user must set the trace mode ON and the following files should be examined to determine the cause of failure.
  - trace file (output from trace)
  - db2diag.log
  - db2cli.log

Chapter 6. Troubleshooting instance creation and configuration

If you install IBM Tivoli Directory Proxy Server or IBM Tivoli Directory Full Server, IBM Tivoli Directory Server requires instance creation and configuration (for the IBM Tivoli Directory Full Server only) after installation. No directory server instance is created by default. This chapter explains how to troubleshoot these processes by providing descriptions of instance creation and configuration options, instructions for avoiding common problems, and troubleshooting steps for instance creation and configuration-related errors.

Instance creation overview and common errors

The following sections discuss instance creation and possible errors you might encounter.

Instance creation overview

After you install a server, you must have an user account on the system and then create a directory server instance. You can create a user on a system either using the Instance Administration Tool (**idsxinst**), which has a GUI, or the **idsadduser** command-line utility. You can then create an instance using either the Instance Administration Tool (**idsxinst**) or the **idsicrt** command-line utility. When you create a directory server instance, a database instance is also created if the IBM Tivoli Directory Full Server package is installed on the computer. By default, the directory server instance and the database instance have the same name. The name must match the name of an existing user on the system that meets certain qualifications. See the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* for information about the necessary qualifications.

You can have multiple directory server instances on one computer. The files for each instance are stored in a path that includes the instance name.

After successful installation of the server, if you used the InstallShield GUI to install, the Instance Administration Tool runs. If you did not use the InstallShield GUI to install, you must run the Instance Administration Tool or use the **idsicrt** command-line utility.

You must perform the following configuration tasks before you can use the server:

- Create the directory server instance.
- Set the IBM Tivoli Directory Server administrator distinguished name (DN) and password. This operation can be compared to defining the root user ID and password on AIX, Linux, and Solaris systems.
- Configure the database, unless the server is a proxy server. (Be sure that you have created the user ID for the database owner first.) You do not need to configure a database for a proxy server.

You can also use the Instance Administration Tool for the following tasks:

- Edit the TCP/IP settings for an instance
- View all instances on the computer
- View details about a particular instance

- Delete an instance
- Migrate a server from a previous release to an IBM Tivoli Directory Server 6.1 instance or later versions of directory server instance

Common instance creation errors

The following section discusses possible errors you might encounter with instance administration.

Cannot create additional instance because of invalid IP address

On AIX, Linux, and Solaris systems, if you have two IP addresses configured, and you try to configure two directory server instances that use the two IP addresses, you might receive an error.

For example, assume that you have IP addresses 9.42.40.67 and 9.42.40.125 configured, and you use the following commands to create directory server instances that use these IP addresses:

```
idsicrt -I svtinst3 -i 9.42.40.67
idsicrt -I svtinst4 -i 9.42.40.125
```

You might receive an error message like the following one when you try to create the second instance:

```
[root@tvt5067 root]# idsicrt -I svtinst4 -i 9.42.40.125
GLPCTL062E The specified IP Address '9.42.40.125' is not a valid IP address for
this machine.
```

The problem might be one of the following:

- The Host IP addresses file does not have the correct entry for the second IP address. For example, on Linux systems, the /etc/hosts file must have the second IP entry in the correct format. For example:

  ```
  9.48.181.173    mymachine.mylocation.ibm.com      mymachine
  ```

- The system settings must be such that the system first checks the Host IP addresses file instead of performing a DNS lookup. The setting in the operating system Name service switch file must be changed to perform Host IP resolution lookup before going to the DNS. For example, on Linux systems, the /etc/host.conf file must have the line `multi on` to allow Host IP address file lookup first.

See the documentation for your operating system for information about setting the Name service switch.

On a system with sles10, the idssethost command fails to recognize the second IP address

On a system with sles10, to make Tivoli Directory Server support multiple IP addresses, add IP addresses in the configuration file under the entry "cn=Configuration" as:

```
ibm-slapdIPAddress: <IP_address1>
ibm-slapdIPAddress: <IP_address2>
```

Now, restart the directory server, the server will listen at the IP addresses specified in the configuration file.

Note: Users can provide any number of IP addresses.

Windows 2003 Enterprise Server: Two directory instances can use the same port number

On the Windows 2003 Enterprise Server operating system, two directory instances can run on the same port numbers. For example, a directory instance configured for "all" and another IP address configured for a specific IP address can use the same port.

This is not an error, but the behavior is unique to Windows 2003 Enterprise Server.

On Windows 2003, instance creation might fail during the instance owner creation stage if the user password does not meet the operating system password requirements

Reason:

During the installation of IBM Tivoli Directory Server, default options can be opted to create a default instance. Instance creation and configuration utilities of Tivoli Directory Server may take a password that does not meet the operating system password requirements.

On Windows 2003, the configuration fails with the following messages (this is as per the messages seen on console and is also available in the idsadm.log file of directory server):

```
Jul 16 10:21:42 2007 You have chosen to perform the following actions:

Jul 16 10:21:42 2007 GLPGRP019I System user will be created for directory
                                server instance.
Jul 16 10:21:42 2007 GLPGRP020I The system user 'idsinst' will be created.
Jul 16 10:21:42 2007 GLPGRP026I The user 'idsinst' will be a member of the
                                'administrators' group.
Jul 16 10:21:42 2007 GLPGRP005I The password for user 'idsinst' will be set.
Jul 16 10:21:42 2007 GLPGRP002I Creating system user 'idsinst'.
Jul 16 10:21:42 2007 GLPGRP006I Setting the password for user 'idsinst'
Jul 16 10:21:43 2007 GLPGRP043E Failed to create user 'idsinst'. Error code
                                return by 'NetUserAdd API' is 2245.
Jul 16 10:21:43 2007 GLPGRP010W The program did not complete successfully.
                                View earlier error messages for information
                                about the exact error.
```

Solution:

In this particular scenario, the explanation for the message, 2245 - NERR_PasswordTooShort, is that the password was shorter than the operating system password requirements. To rectify this, use a password that meets the operating system password requirements. To know more about password requirements for a Windows operating system, see Password must meet complexity requirements.

The instance creation can also fail for the following reasons:

- Password is too long
- Password could be the recent most in the change history
- Password does not have enough unique characters
- Password does not meet the password policy requirements

On a 32–bit Windows 2008 operating system, which is installed on a 64–bit hardware, the Administration server might fail to start after the creation of a Tivoli Directory Proxy Server instance

At times the Administration server fails to start even after displaying the following output from the command **idsicrt**.

```
GLPCTL074I Starting admin server for directory server instance: 'inst1'.
GLPCTL075I Started admin server for directory server instance: 'inst1'.
GLPICR029I Created directory server instance: : 'inst1'.
```

This behavior has been observed when creating a Tivoli Directory Proxy Server instance using the command **idsicrt -x** on a 32–bit Windows 2008 operating system, which is installed on a 64–bit hardware platform. In addition the following behavior were observed for the administration server with the above configuration:

- Might fail when the **idsicrt** command is run with the **-x** option in no-prompt mode
- Might fail when the **idsicrt** command is run with the **-x** option in prompt mode
- Might fail when the **idsicrt** command is run with the **-x** option in no-prompt mode and trace is on
- Starts when the **idsicrt** command is run with the **-x** option in prompt mode and trace is on
- Starts when the **idsicrt** command is run without the **-x** option

Configuration overview and common errors

The following sections discuss configuration and possible errors you might encounter.

Overview

If you do not set the Administrator DN and password or configure the database through the Instance Administration Tool, you can use the Configuration Tool (**idsxcfg**) for these and other tasks.

The Configuration Tool has a GUI, and it can be used for the following tasks:

- Setting or changing the IBM Tivoli Directory Server administrator distinguished name (DN) and password
- Configuring and unconfiguring the database
- Enabling and disabling the changelog
- Adding or removing suffixes
- Adding schema files to or removing schema files from the list of schema files to be loaded at startup
- Importing and exporting LDAP Data Interchange Format (LDIF) data
- Backing up, restoring, and optimizing the database

If you prefer to use the command line, all the tasks in the list can be done with the following command-line utilities.

- **idsdnpw** sets the administrator DN and password
- **idscfgdb** configures the database for a directory server instance
- **idsucfgdb** unconfigures the database
- **idscfgchglg** configures the change log for a directory server instance
- **idsucfgchglg** unconfigures the change log for a directory server instance
- **idscfgsuf** configures a suffix for a directory server instance
- **idsucfgsuf** unconfigures a suffix for a directory server instance
- **idscfgsch** configures a schema file for a directory server instance
- **idsucfgsch** unconfigures a schema file for a directory server instance
- **idsldif2db** or **bulkload** imports LDIF data

- **idsdb2ldif** exports LDIF data
- **idsdbback** backs up the database
- **idsdbrestore** restores the database
- **idsrunstats** optimizes the database

Common errors

The following section discusses possible errors you might encounter with configuration.

Interrupting Configuration Tool database tasks causes an incorrect status for the files

If you are using the Configuration Tool to configure, unconfigure, import, export, backup, restore, or optimize a database and the process is interrupted by, for example, a segmentation fault, the status of the files is returned incorrectly. When you try to restart the process, the message

```
Task is already running.
```

is displayed. This is because the status output for the process is monitored through files in the idsslapd-<*instance_name*>/tmp folder that were not deleted when the process was interrupted.

To restart the interrupted process, you must first manually delete all of the *.dat and *.stat files in the idsslapd-*instance_name*>/tmp directory (where *instance_name* is the instance name).

Failure when configuring an existing database instance and database

If you are using AIX, Linux, or Solaris and you are configuring an existing database and database instance using the **idscfgdb** command, a core dump might occur after the configuration is completed. This failure, however, can be ignored. The database is successfully configured.

Error when starting the Configuration Tool on AIX

The following error might occur when you start the Configuration Tool on AIX:

```
# idsxcfg exec(): 0509-036 Cannot load program idsxcfg
                          because of the following errors:
         0509-022 Cannot load module /usr/ldap/lib/libdbadmin.a.
         0509-150   Dependent module /usr/ldap/lib/libdb2.a(shr_64.o)
                    could not be found.
         0509-152   Member shr_64.o is not found in archive
```

If this error occurs, check the following:

- You have a supported version of DB2. (See the *IBM Tivoli Directory Server Installation and Configuration Guide* for information about supported versions of DB2.)
- You have 64-bit hardware.
- You are running a 64-bit kernel.
- You migrated your database to 64-bit.

Configuration programs terminate on AIX

If the configuration GUI tools terminate immediately when you start them, check the LIBPATH. If the *jre*/bin/classic directory of a JVM other than the one provided with IBM Tivoli Directory Server comes before the %LDAPHOME%/java/bin/classic directory, do one of the following:

- Remove the extraneous JVMs from the LIBPATH.
- Place the %LDAPHOME%/java/bin/classic directory in front of the other JVM directories in the LIBPATH.

DB2 does not configure properly

Note: Before configuring the database, be sure that the environment variable DB2COMM is **not** set.

If a failure occurs during database configuration, usually one of the following is the cause:

- The user ID was not set up correctly.
- The permissions for the user ID are not correct.
- Remnants of a previous database (database or table space directories) with the name you specified for the database are present on the system.
- There is not enough space in the location you specified.
- The location is not accessible.

Check to see if there are problems with any of these items, and then try to configure again after you fix the problem.

Note: If you use the Configuration Tool to configure and configuration fails, the Configuration Tool does some cleanup, and this can sometimes fix the problem. If you do not find any of the problems in the list, try configuring again.

Server does not start after making changes to configuration file attributes

The attributes defined in IBM Tivoli Directory Server configuration file are significant to only the first 18 characters. Names longer than 18 characters are truncated to meet the DB2 restriction.

If you want to index the attribute, the limit is further restricted to 16 characters. If you add attributes longer than 18 characters, the server might not start. For additional information, see the Web Administration Tool helps under **Reference**, Directory Schema.

Transaction log is full

The following messages might be displayed at IBM Tivoli Directory Server startup if the schema defines too many attributes:

```
SQL0965C The transaction log for the database is full
SQLSTATE=57011 slapd unable to start because all backends failed to configure
```

You might need to increase the DB2 transaction log sizes by typing the following:

```
db2 update db cfg for ldaptest using logprimary X
db2 update db cfg for ldaptest using logsecond X
```

where *X* is greater than the currently defined size. You can check the current log size by using the following command:

```
db2 get db cfg for dbname
```

Problems in Configuration Tool windows

The following sections describe problems that might occur on the Configuration Tool panels while you are using the Configuration Tool.

Translated titles might truncate in Configuration Tool: Titles in the pop-up windows in the Configuration Tool might truncate depending upon the language. If this problem occurs, you can resize the window accordingly, depending on your display.

Some keyboard commands fail on Browse windows: On Windows systems, for functions in the Configuration Tool (such as Import LDIF data) that contain a path field with a **Browse** button, you might not be able to use the Space, Enter or arrow keys on the keyboard to view the contents of the **Look in** menu on a **Browse** window. To work around this problem, press Alt+Down Arrow to display the **Look in** menu, and use the arrow keys to select a drive.

Task not highlighted when using keyboard: On AIX and Linux systems, in the Configuration Tool, when you use the arrow keys to move between tasks in the task list on the left, the tasks might not be highlighted, and the information in the window on the right might not change. To select a task on the left, move to the task you want using the arrow keys, and then press the Spacebar.

NullPointer exception when exiting the Configuration Tool: If you exit the Configuration Tool after entering an invalid database name, a NullPointer exception occurs in the command window where the **idsxcfg** command was executed. The exception does not affect the configuration process.

Bulkload messages continue to be displayed in the table after the data is imported: In the Configuration Tool, if you import LDIF data and select the **Bulkload** option, messages continue to be displayed in the table even after the data is imported. Some of these messages might be exceptions, but the import is successful.

Debugging configuration

During configuration, you might experience some problems with the configuration programs. There are some extra debugging steps that can help you and IBM Software Support determine the cause of these problems.

Database configuration: Because there are so many variables at play during configuration, errors can occur. Some of the factors that can affect this option are:

- Which platform, and which version of the operating system, you are using.
- Which version of DB2, and which fix packs have been installed for it.

  Note: DB2 comes in a wide variety of packages: Personal Edition, Enterprise Edition, Extended Enterprise Edition, and others. Many of these are supported across several versions of DB2, and each version can have several available fix packs.
- Amount of disk space available in affected drives and partitions.
- Third party software that alters commonly used environment variables.

If the database configuration fails, the bottom-line question is, "What failed, and how do I fix it?" The following sections describe sources of output that can be used to debug configuration problems.

Standard sources of output: There are several "standard" sources of information available:

- The output on the screen

  All of the configuration programs are either started from a console command line prompt or open a background console. As the database configuration

progresses, status messages (and limited error messages) are displayed in the associated console window. If a problem occurs, copy these messages to the system clipboard and then save them in a file for the IBM Software Support teams.

- DB2 log files

  If the error is a direct error from DB2, then DB2 often creates message or error files (in the /tmp directory on AIX, Linux, and Solaris platforms). If you have a database configuration problem on an AIX, Linux, or Solaris system, examine all of the files in the /tmp directory that were created around the time of the attempted configuration.

  On Windows systems, examine any DB2 error logs in your DB2 installation directory under the directory named for the instance you were trying to configure. For example, if your were trying to create an instance and database named ldapdb2, and if your DB2 was installed in D:\sqllib, examine the files in the D:\sqllib\ldadb2 directory if it exists. In particular, look for and examine the file named db2diag.log in that directory.

Creating advanced debug output: See "Server debug mode" on page 12 for information about using debugging tools that are provided.

Chapter 7. Troubleshooting DB2

This chapter contains information about problems related to DB2.

DB2 license file expired

If you see the following message during DB2 or server startup:

```
GLPCTL010E Failed to start database manager for database instance: <instance name>.
```

you might have a problem with your electronic DB2 license. To verify this, type the following at the command prompt:

```
db2start
```

If your license is correct, you see the message:

```
SQL1063N DB2START processing was successful.
```

Otherwise, you see a message indicating that your license has expired or will expire in some number of days.

If there is a problem with your electronic DB2 license, one of the following situations might be the cause:

- You have a demonstration license.
  1. To upgrade your DB2 product from a demonstration license to a product license, copy the license file from the DVD to the system where DB2 is installed; you do not need to reinstall DB2.

     If you installed the version of DB2 that is provided with IBM Tivoli Directory Server, the license file is in one of the following locations:

     - If you have a DVD: *<mount_point>*/db2/db2/license/db2ese_o.lic (or *<cdrom_drive:>*\db2\db2\license\db2ese_o.lic for Windows)
     - If you downloaded a zip file for installation: *directory_where_file_was_unzipped*\tdsV6.3\db2\db2\license\db2ese_o.lic
     - If you downloaded a tar file for installation: *directory_where_file_was_untarred*/tdsV6.3/db2/db2/license/db2ese_o.lic

     Note: Your Proof of Entitlement and License Information booklets identify the products for which you are licensed.
  2. After you have a valid license file on the system, run the following command to activate the license:

     ```
     db2licm -a license_filename
     ```
- You have purchased a different DB2 product.

  If you install a DB2 product as Try-and-Buy, and you buy a different DB2 product, you must uninstall the Try-and-Buy product and then install the new one that you have purchased. Type the following at a command prompt to upgrade your DB2 license:

  ```
  db2licm -a license_filename
  ```

Note: *license_filename* is the name of the license file; for example, db2udbee.lic.

Recovering from migration failure in DB2 9.x

The idsdbmigr tool uses DB2 backup and DB2 restore mechanism to recover from migration failure since direct recovery from migration failure is not available with DB2 9.x.

The DB2 database can be recovered from the DB2 database backup. The DB2 database can be backed up using the idsdbback utility shipped along with Tivoli Directory Server or by using the DB2 commands like DB2 BACKUP DATABASE <*database-alias*>. The DB2 database can be restored by using the idsdbrestore utility shipped along with Tivoli Directory Server or by using the DB2 commands like DB2 RESTORE DATABASE <*source-database-alias*>.

See the section "Overview of online backup and restore procedures for Tivoli Directory Server" in *IBM Tivoli Directory Server Version 6.3 Administration Guide* to know more about DB2 backup and restore.

Installing DB2 9.5 on Red Hat Enterprise Linux (RHEL) 5 64-bit or SuSE Linux Enterprise Server (SLES) 10 operating system for Intel Linux or zLinux

Problem

When installing DB2 9.5 on RHEL 5 64-bit or SLES 10 operating system for Intel® Linux or zLinux, an error message is displayed. For example,

```
/db2v9.5/ese/db2/linuxamd64/install/../bin/db2usrinf: error while
loading shared libraries: libstdc++.so.5: cannot open shared object file:
No such file or directory
```

Cause The shared library files, such as libstdc++.so.5, might be missing.

Solution

When installing DB2 9.5 on RHEL 64-bit or SLES 10 operating system for Intel Linux or zLinux, you must consider the following:

- DB2 9.5 requires RHEL 5 GA or above versions.
- Install all available compat-libstdc++ RPM-packages before installing DB2 9.5 and ensure that libstdc++.so.5 is available, this is required for DB2 servers and clients.

To know more about installation requirements on DB2 servers, see https://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=/com.ibm.db2.luw.qb.server.doc/doc/r0008865.html.

DB2 diagnostic information (db2diag.log)

The db2diag.log file contains DB2 diagnostic information. A user with appropriate privileges can set the fully qualified path for DB2 diagnostic information using the diagpath parameter. If this parameter is null, the diagnostic information will be written to the files in the following directories.

For DB2 v9.5

On Windows platforms

- In Windows Vista environment, the db2 diagnostic error logs are written to the *ProgramData\IBM\DB2\ directory*.
- In Windows 2003 and XP environment, the db2 diagnostic error logs are written to the *Documents and Settings\All*

Users\Application Data\IBM\DB2\Copy Name\<instance>, where *Copy Name* is the name of DB2 copy.

On AIX, Linux, and Solaris platforms
The db2 diagnostic error logs are written to *INSTHOME/sqllib/db2dump*, where *INSTHOME* is the home directory of the instance.

To know more about DB2 diagnostic information in DB2 v9.5, see Diagnostic data directory path configuration parameter.

An SQL0964C error, (transaction log full), might get displayed when loading large amounts of data from a file

When loading a file that contains a large number of entries, you might receive the following error messages:

```
SQL0964C SQLSTATE=57011
```

You can troubleshoot this error by increasing the transaction log size. For this, perform the following:

1. At command line issue the following command and the password for the user:
   ```
   su - <db2instownername>
   ```
   where, *<db2instownername>* is the name of the DB2 instance owner.
2. Determine the current log file size setting by issuing the command:
   ```
   db2 get db config for <db2instancename> | grep -i logfilsiz
   ```
3. Increase the size of the log file size setting by issuing the command:
   ```
   db2 UPDATE db cfg for <db2instancename> using LOGFILSIZ <new_value>
   ```
4. Stop the slapd process.
5. To apply the changes related to database, issue the following command:
   ```
   db2 force applications all
   ```
6. Restart the slapd process.

Note: You can also use the bulkload utility to load files with large amounts of entries.

A Tivoli Directory Server instance might start in config-only mode after applying DB2 fix pack

If you get an error or your directory server instance starts in the config-only mode after applying a DB2 fix pack, you must follow the post-install instructions provided in the current DB2 fix pack readme file to resolve the problem related to applying of DB2 fix packs. You must also check the ibmslapd.log and db2cli.log files for the error descriptions that may have been logged.

The following example error messages might appear in the ibmslapd.log file after applying the DB2 fix pack:

```
02/31/07 21:26:06 Error code -1 from odbc string:" SQLTables "  .
02/31/07 21:26:07 Error code -1 from odbc string:" SQLFetch "  .
02/31/07 21:26:07 Error code -1 from odbc string:" SQLFetch "  .
02/31/07 21:26:07 Error code -1 from odbc string:" SQLFetch "  .
```

The following example error messages might appear in the db2cli.log file after applying the DB2 fix pack:

```
02/31/07 21:26:06 native retcode = -443; state = "38553"; message =
  "[IBM][CLI Driver][DB2/6000] SQL0443N  Routine "SYSIBM.SQLTABLES"
  (specific name "TABLES") has returned an error SQLSTATE with diagnostic
  text "SYSIBM:CLI:-805". SQLSTATE=38553"
02/31/07 21:26:07 native retcode = -99999; state = "24000"; message =
  "[IBM][CLI Driver] CLI0115E  Invalid cursor state. SQLSTATE=24000"
02/31/07 21:26:07 native retcode = -99999; state = "24000"; message =
  "[IBM][CLI Driver] CLI0115E  Invalid cursor state. SQLSTATE=24000"
02/31/07 21:26:07 native retcode = -99999; state = "24000"; message =
  "[IBM][CLI Driver] CLI0115E  Invalid cursor state. SQLSTATE=24000"
```

To know more about the post-install steps to be followed after applying a DB2 fix pack, see Applying DB2 fix packs.

Chapter 8. Troubleshooting the Web Administration Tool and the application server

The IBM Tivoli Directory Server Version 6.3 Web Administration Tool is installed on an application server, such as Embedded WebSphere Application Server, which is included with the IBM Tivoli Directory Server and administered through a console. WebSphere Application Server (WAS) can also be used as the application server. This chapter explains how to troubleshoot the IBM Tivoli Directory Server Web Administration Tool and application server.

Troubleshooting the Web Administration Tool

The following sections contain troubleshooting information for the Web Administration Tool.

Corruption of data entered in the Web Administration Tool

If data that you enter in non-English languages in the Web Administration Tool is damaged, do the following:

On the embedded version of WebSphere Application Server - Express®

Edit the server.xml file in the following directory:

WAS_home/appsrv/config/cells/DefaultNode/nodes/DefaultNode/servers/server1

Add the text shown in bold to the stanza as shown:

```
<processDefinition xmi:type="processexec:JavaProcessDef"
   xmi:id="JavaProcessDef_1"
   executableName="${JAVA_HOME}/bin/java"
   executableTarget="com.ibm.ws.runtime.WsServer"
   executableTargetKind="JAVA_CLASS"
   workingDirectory="${USER_INSTALL_ROOT}">
<execution xmi:id="ProcessExecution_1" processPriority="20" runAsUser=""
   runAsGroup=""/>
<monitoringPolicy xmi:id="MonitoringPolicy_1" pingInterval="60"
   maximumStartupAttempts="3" pingTimeout="300" autoRestart="true"
   nodeRestartState="STOPPED" />
<ioRedirect xmi:id="OutputRedirect_1"
   stdoutFilename="${SERVER_LOG_ROOT}/native_stdout.log"
   stderrFilename="${SERVER_LOG_ROOT}/native_stderr.log"/>
<jvmEntries xmi:id="JavaVirtualMachine_1" classpath="" bootClasspath=""
   verboseModeClass="false" verboseModeGarbageCollection="false"
   verboseModeJNI="false" initialHeapSize="0"
   maximumHeapSize="256" runHProf="false" hprofArguments=""
   debugMode="false" debugArgs="-Djava.compiler=NONE -Xdebug -Xnoagent
   -Xrunjdwp:transport=dt_socket,server=y,suspend=n,address=7777"
   genericJvmArguments="">
<systemProperties xmi:id="Property_10"
   name="client.encoding.override" value="UTF-8" required="false"/>
</jvmEntries>
```

On WebSphere Application Server

On the WebSphere Administrative Console tree:

- Select **Servers**.
- Select **Application Server**.
- Select the server you want; for example, server1.
- Click **Process Definition**.

- Click **Java Virtual Machine**.
- Click **Custom Properties**.
- Click the appropriate button for making a new property.
- In the **Name** field, type `client.encoding.override`.
- In the **Value** column, type `UTF-8`.
- Click **Apply**.
- Stop and restart the WebSphere Application Server.

Migrating files when patching or migrating the Web Administration Tool

You must back up the following four files before uninstalling the IDSWebApp.war file (the Web Administration Tool) and restore them after you have reinstalled the war file:

- `console administartor login and password settings ${WASHome}/profiles/TDSWebAdminProfile/installedApps/DefaultNode/ IDSWebApp.war.ear/IDSWebApp.war/WEB-INF/classes/security/console_passwd`
- `# console servers & console properties / SSL key database settings ${WASHome}/profiles/TDSWebAdminProfile/installedApps/ DefaultNode/IDSWebApp.war.ear/IDSWebApp.war/WEB-INF/classes/IDSConfig/ IDSConfig/IDSServersConfig/IDSServersInfo.xml`
- `# console properties / component management settings ${WASHome}/profiles/TDSWebAdminProfile/installedApps/DefaultNode/ IDSWebApp.war.ear/IDSWebApp.war/WEB-INF/classes/IDSConfig/IDSAppReg/IDSAppReg.xml`
- `# console properties / session properties settings ${WASHome}/profiles/TDSWebAdminProfile/installedApps/ DefaultNode/IDSWebApp.war.ear/IDSWebApp.war/WEB-INF/classes/IDSConfig/ IDSSessionConfig/IDSSessionMgmt.xml`

Additional login panels fail

When using the Web Administration Tool, do not open additional login panels from the **File** options of the browser. Only one instance of the Web Administration Tool can function on a single browser instance. They cannot share the same cookies. Additional login panels must be opened from new instances of the browser.

For AIX, Linux, and Solaris systems:

Launch new windows from the command line using the `&` option. For example:

```
mozilla &
```

For Windows systems:

- Internet Explorer - Open additional Internet Explorer windows using the **Start** window or an Internet Explorer short cut from the desktop.
- Mozilla - The Mozilla Web browser does not support multiple Web Administration Tool sessions on Windows.

Note: Netscape browsers are no longer supported.

idsldapmodify command puts Web Administration Tool into inconsistent state

If you are logged into the Web Administration Tool and you change your password using the command line (**idsldapmodify** command), the Web Administration Tool changes the server status to stopped. This occurs because the Web Administration Tool opens new connections to the server every time it

launches a task. The Web Administration Tool tries to connect to the server with the old password because it is unaware that the password has been changed; consequently the connection fails. You must log out and log back in using the new password.

To avoid this situation, if you have sufficient access authority, use the **User properties -> Change password** option to change your user password when working in the Web Administration Tool.

Web Administration Tool tabs, table headers, and static list boxes are displayed in incorrect language

The following problem has been encountered only on the AIX operating systems; however, Solaris and Linux systems might encounter the same problem.

The environment variables **LC_ALL** and **LANG** must be set to a native locale supported by Java; for example en_US.iso88591. They must not be set to either POSIX or C.

```
export LC_ALL=<new language>
export LANG=<new language>
```

The translation of the tabs, table headers, and static list boxes are saved in the language that was first used by the application server the first time a user logs into the Web Administration Tool application. If you change the locale on your machine, you might see the following exception:

```
java.lang.InternalError: Can't connect to X11 window server using ':0.0'
as the value of the DISPLAY variable.
        at sun.awt.X11GraphicsEnvironment.initDisplay(Native Method)
        at sun.awt.X11GraphicsEnvironment.<clinit>
           (X11GraphicsEnvironment.java:58)
        at java.lang.Class.forName0(Native Method)
        at java.lang.Class.forName(Unknown Source)
        at java.awt.GraphicsEnvironment.getLocalGraphicsEnvironment
           (GraphicsEnvironment.java:53)
        at sun.awt.motif.MToolkit.<clinit>(MToolkit.java:63)
        at java.lang.Class.forName0(Native Method)
        at java.lang.Class.forName(Unknown Source)
        at java.awt.Toolkit$2.run(Toolkit.java:507)
        at java.security.AccessController.doPrivileged(Native Method)
        at java.awt.Toolkit.getDefaultToolkit(Toolkit.java:498)
        at java.awt.Toolkit.getEventQueue(Toolkit.java:1171)
        at java.awt.EventQueue.invokeLater(EventQueue.java:506)
        at javax.swing.SwingUtilities.invokeLater(SwingUtilities.java:1086)
        at javax.swing.Timer.post(Timer.java:337)
        at javax.swing.TimerQueue.postExpiredTimers(TimerQueue.java:190)
        at javax.swing.TimerQueue.run(TimerQueue.java:226)
        at java.lang.Thread.run(Unknown Source)
```

To correct this exception, you must export the **DISPLAY** variable so that it is a valid computer; for example, the computer on which the application server is running. Then perform **xhost +** on the application server computer.

On the computer to which you want to export the DISPLAY, issue the command:

```
export DISPLAY=<valid_computer_name>:0
```

On the *<valid_computer_name>* issue the command:

```
xhost +
```

Microsoft Internet Explorer browser problems

If you have problems running the Web Administration Tool with Microsoft® Internet Explorer, try making the following changes to the cache setup:

- Click **Tools —> Internet Options**, and select **General**. Then click **Settings**. Under **Check for newer versions of stored pages**, click **Every visit to the page**.
- If you have unpredictable results when using the browser, the cache might be storing pages with errors. On the General folder page, click **Delete files** and **Clear History** to clear the cache. Use these options as often as necessary.
- Shutting down and restarting the browser can also repair some intermittent problems.

HTML special characters are not displayed correctly

Special characters in read-only data coming from the server are not displayed correctly in the HTML page. This is because of the way that the HTML is rendered by the Web browsers. For example:

- A string containing multiple spaces such as "a b" is rendered as "a b".
- A string containing the special character '<' is truncated. For example, "abc<abc" is rendered as "abc".

This affects fields such as labels, drop-down boxes, tables, and captions.

Web Administration Tool requires IBM JDK on a Domino server

If you want to use the Web Administration Tool with a Domino® server you must use the IBM 1.4.2 JDK or later. Using the JDK from Sun results in communication exceptions.

The following are limitations on the Domino server:

- The Manage schema functions do not work.
- Domino does not support user-defined suffixes.

Note: The standard suffix on the Domino server is a blank. Consequently, to view entries, you must select the radio button with the plus sign (+) next to it and click **Expand**.

Web Administration Tool does not save templates created with an object class that has no attributes

You can create object classes for IBM Tivoli Directory Server that have no MAY or MUST attributes. Such object classes can be used to create entries using other auxiliary object classes. However, if you attempt to create a template through the Web Administration Tool using such an object class, you are unable to save the template.

Note: All of the object classes included with IBM Tivoli Directory Server contain MAY and MUST attributes. They can be used to create templates.

Using Ctrl+L to view links makes non-editable fields appear editable

If you open the Web Administration Tool using Home Page Reader **Ctrl+L** keystroke to view the links on a Web Administration Tool page, non-editable fields

might appear editable. A text box might appear next to the non-editable field. Although you can enter data in the non-editable fields, the data is not saved.

Internet browser Back and Forward buttons not supported for Web Administration Tool

The **Back** and **Forward** buttons on Internet browsers cannot be used to navigate the Web Administration Tool.

Logging on to the Web Administration Tool console on Internet Explorer

On Windows systems, Web Administration Tool errors occur if all the following conditions exist:

- The Web Administration Tool is installed locally.
- The Web Administration Tool runs on a locally installed version of Microsoft Internet Explorer.
- The Web Administration Tool uses the locally installed Embedded WebSphere Application Server.
- An IP address or hostname is part of the URL used to access the Web Administration Tool.

If these conditions exist on your computer, avoid errors by using localhost instead of an IP address or hostname when logging on to the Web Administration GUI console.

For example, open an Internet Explorer Web browser and type the following in the **Address** field:

```
http://localhost:12100/IDSWebApp/IDSjsp/Login.jsp
```

Difficulties encountered using the Web Administration GUI console on the Windows Server 2003 platform

Web Administration Tool errors occur if all the following conditions exist:

- The Web Administration Tool is installed locally.
- The Web Administration Tool runs on a locally installed version of Microsoft Internet Explorer.
- The Web Administration Tool uses the locally installed Embedded WebSphere Application Server.
- An IP address or hostname is part of the URL used to access the Web Administration Tool.

To avoid these errors:

1. If Embedded WebSphere Application Server is running locally, add **http://localhost** to the list of trusted sites.
2. If Embedded WebSphere Application Server is running on a remote machine, add the IP address or hostname of the computer on which the Web application server is running to the list of trusted sites. **http://<***IP address***>** or **http://<***hostname***>**

To add a Web address to the Trusted Site list:

1. Click **Tools -> Internet Options -> Security -> Trusted Site -> Sites**.
2. Type the Web address in the Web site field.

3. Click **Add**.
4. Click **OK**.

To log on to the Web Administration Tool on the local computer, open an Internet Explorer Web browser and type the following in the **Address** field:

```
http://localhost:12100/IDSWebApp/IDSjsp/Login.jsp
```

To log on to the Web Administration Tool on a remote computer, open an Internet Explorer Web browser and type the following in the **Address** field:

```
http://<IP address> or <hostname>:12100/IDSWebApp/IDSjsp/Login.jsp
```

A new user might fail to logon to Web Administration Tool for the first time, if the password policy is enabled and "User must change password after reset (pwdMustChange)" in set

If the password policy is enabled and "User must change password after reset (pwdMustChange)" in set on the Password policy settings 1 panel in the Manage password policies wizard, user might not be able to logon to Web Administration Tool.

To resolve the problem, user can use the ldapchangepwd command line utility to reset the password and then use the new password to logon.

When performing a backup using the Web Administration Tool to a backup location that is specified in an NLV string another folder gets created

If the browser locale on which the Web administration tool is running is different from the system locale on which a Tivoli Directory Server instance is running, then it is observed that an NLV string folder also gets created apart from the backup folder when the backup operation is initiated.

This is because the string entered (as backup location) is used as a file path, which must be representable in the system's local code page. When the Web administration tool attempts to translate the Unicode input to the local code page to create the file path, it encounters Unicode input characters that cannot be represented to the system's locale causing the above problem.

Troubleshooting the embedded version of WebSphere Application Server - Express

The following sections contain troubleshooting information for the embedded version of WebSphere Application Server - Express.

Error when starting the embedded version of WebSphere Application Server - Express on AIX

Starting the embedded version of IBM WebSphere Application Server - Express on AIX (**startServer.sh server1**), might not work because port 9090 is already being used. See the WAS_install_path/logs/server1 directory for the actual log files. Usually the SystemErr.log and SystemOut.log files are most helpful, although the other logs might have some useful information.

To change the port number for the embedded version of IBM WebSphere Application Server - Express from 9090 to 9091, which is the port used on AIX

computers, edit the *WAS_inst_path*/config/cells/DefaultNode/virtualhosts.xml file and change 9090 to 9091. Do the same thing in the *WAS_inst_path*/config/cells/DefaultNode/nodes/DefaultNode/servers/server1/server.xml file. *WAS_inst_path* is the path where the embedded version of IBM WebSphere Application Server - Express is installed.

Note: This path does have two subdirectories named DefaultNode.

Make one change in each file for a total of two updates.

Chapter 9. Troubleshooting replication

This chapter contains troubleshooting information about replication and errors commonly encountered during replication.

Replication overview

Directory servers use replication to improve performance, availability, and reliability. Replication keeps the data in multiple directory servers synchronized. Replication provides three main benefits:

- Redundancy of information - Replicas back up the content of their supplier servers.
- Faster searches - Search requests can be spread among several different servers, instead of a single server. This improves the response time for the request completion.
- Security and content filtering - Replicas can contain subsets of the data in a supplier server.

See the replication chapter in the *IBM Tivoli Directory Server Version 6.3 Administration Guide* for a more detailed overview of replication.

Diagnosing replication errors

The following sections provide information about identifying the source of replication errors.

Sample replication topology

The following is an example of a basic replication topology. If you are not sure if you have set up your topology correctly, you can compare it against this one. This topology assumes that there is a suffix in the server configuration for o=sample.

This example file sets up a master server called **masterhost** with a replica called **replicahost**:

```
version: 1

dn: cn=replication, cn=localhost
objectclass: container

dn: cn=simple, cn=replication, cn=localhost
replicaBindDN: cn=master
replicaCredentials: ldap
description: simple bind credentials
objectclass: ibm-replicationCredentialsSimple

dn: o=sample
objectclass: organization
objectclass: ibm-replicationContext

dn: ibm-replicaGroup=default,o=sample
objectclass: ibm-replicaGroup

dn: ibm-replicaServerId=masterhost-389,ibm-replicaGroup=default,o=sample
ibm-replicationserverismaster: true
cn: masterhost
description: master
```

```
objectclass: ibm-replicaSubentry

dn: cn=replicahost,ibm-replicaServerId=masterhost-389,\
    ibm-replicaGroup=default,o=sample
ibm-replicaconsumerid: replicahost-389
ibm-replicaurl: ldap://replicahost:389
ibm-replicaCredentialsDn: cn=simple, cn=replication, cn=localhost
description: masterhost to replicahost
objectclass: ibm-replicationAgreement
```

Add the example file to **masterhost** with following command:

```
ldif2db -r yes -i <in>
```

After the file is loaded, export the data from the database using the following command:

```
db2ldif -o <out>
```

The server configuration file for **masterhost** must contain:

```
dn: cn=Configuration

ibm-slapdServerId: masterhost-389
```

The configuration file for **replicahost** must contain:

```
dn: cn=Configuration

ibm-slapdServerId: replicahost-389
```

and the following entry

```
dn: cn=master server, cn=configuration
cn: master server
ibm-slapdMasterDn: cn=master
ibm-slapdMasterPW: ldap
ibm-slapdMasterReferral: ldap://masterhost:389
objectclass: ibm-slapdReplication
```

Both **masterhost** and **replicahost** require the replicated subtree suffix in their configuration files:

```
dn: cn=Directory,cn=RDBM Backends,cn=IBM Directory,cn=Schemas,cn=Configuration
...
ibm-slapdSuffix: o=sample
```

Monitoring replication status using idsldapsearch

Note: The **idsldapsearch** examples in this section are based on the sample replication topology provided earlier in this chapter. See "Sample replication topology" on page 55 for more information.

There are many operational attributes that provide replication status information when explicitly requested on a search. One of these attributes is associated with the entry that is the base of the replicated subtree, that is, the entry that the ibm-replicationContext objectclass was added to. If you do a base search of that entry and request that the ibm-replicationIsQuiesced attribute is returned, the return attribute indicates if the subtree has been quiesced; for example:

```
idsldapsearch -h <hostname> -p <port> -b "o=sample"  -s "base"
"objectclass=ibm-replicationContext" ibm-replicationIsQuiesced
```

The remainder of the status-related operational attributes are all associated with a replication agreement object. These attributes are only returned when explicitly

requested on the search; for example, the following idsldapsearch example requests replication agreement status information indicating the replication state for all the replication agreements:

```
idsldapsearch -h <hostname> -p <port> -b "o=sample" -s "sub"
"objectclass=ibm-replicationAgreement" ibm-replicationState
```

The available attributes are:

- **ibm-replicationLastActivationTime**: The time that the last replication session started between this supplier and consumer.
- **ibm-replicationLastFinishTime**: The time that the last replication session finished between this supplier and consumer.
- **ibm-replicationLastChangeId**: The change ID of the last update sent to this consumer.
- **ibm-replicationLastGlobalChangeId**: The change ID of the last update to a global entry sent to this consumer. Global entries are things like cn=schema or cn=pwdpolicy that apply to the entire contents of a DIT.

  This attribute is deprecated in version 6.0.
- **ibm-replicationState**: The current state of replication with this consumer. Possible values are:

  Ready
  In immediate replication mode, ready to send updates as they occur.

  Retry An error exists, and an update to correct the error is sent every 60 seconds.

  Waiting
  Waiting for next scheduled replication time.

  Binding
  In the process of binding to the consumer.

  Connecting
  In the process of connecting to the consumer.

  OnHold
  This replication agreement has been suspended or "held".

  Error log full
  More replication errors have occurred than can be logged. The amount of errors that can be logged is based on the configured value for ibm-slapdReplMaxErrors.
- **ibm-replicationLastResult** The results of the last attempted update to this consumer, in the form:

  <timestamp> <change id> <result code> <operation> <entry DN>

  This attribute is available only if the replication method is single threaded.
- **ibm-replicationLastResultAdditional**: Any additional error information returned from the consumer for the last update. This attribute is available only if the replication method is single threaded.
- **ibm-replicationPendingChangeCount**: The number of updates queued to be replicated to this consumer.
- **ibm-replicationPendingChanges**: Each value of this attribute gives information about one of the pending changes in the form:

  ```
  <change id> <operation> <entry DN>
  ```

Requesting this attribute might return many values. Check the change count before requesting this attribute.

- **ibm-replicationChangeLDIF**: Gives the full details of the last failing update in LDIF. This attribute is available only if the replication method is single threaded.
- **ibm-replicationFailedChangeCount**: Indicates the total number of failed changes logged for the specified replication agreement.
- **ibm-replicationFailedChanges**: Lists the IDs, DNs, update types, result codes, timestamps, numbers of attempts for failures logged for a specified replication agreement.
- **ibm-replicationperformance**: Give the operation counts per connection for multi-threaded replication.

Viewing replication errors using the Web Administration Tool

Using the Web Administration Tool, you can view replication updates that were not completed because of errors that occurred during replication. Viewing this information can help you identify the source of your replication problem.

To view replication errors:

1. Log into the Web Administration Tool.
2. Expand the **Replication management** category in the navigation area and click **Manage topology**.
3. Select the subtree that you want to view from the replicated subtrees list and click the **Show topology** button on the table.
4. Click the **View errors** button.

From the "View errors" panel you can:

- View the details of a specific error in the replication agreement.
- Attempt to perform the selected replication update again.
- Attempt to perform all failed replication updates again.
- Remove a replication error from the table.
- Remove all replication errors form the table.

To view the details of a specific error in the replication agreement:

1. Select the replication error you want to view from the **Replication error management** table and click the **View details** button on the tool bar. The **Replication error details** table contains the following information about the selected error.

Supplier
The host name or IP address of the supplier

Consumer
The host name or IP address of the consumer

Change ID
The unique ID of the failed update sent to the consumer

Update DN
The DN of the entry on which the update was attempted

Operation type
The type of update request; for example, add or delete

Details

The LDIF representation of the entry associated with the failed update, including all the operational attributes

Controls

The controls used during the update

Viewing replication errors using the idsldapsearch command

The replication errors can be displayed by two replication status attributes:

- ibm-replicationFailedChanges
- ibm-replicationFailedChangeCount

For example, use the **idsldapsearch** command to display replication errors:

```
idsldapsearch -D <adminDN> -w <adminPW> -h <servername>
-p <portnumber> -b " " -s sub objectclass=ibm*nt
    ibm-replicationfailedchanges ibm-replicationfailedchangecount
```

This command can return output similar to the following:

```
cn=<server>-1389,ibm-replicaServerId=<server>-389,
ibm-replicaGroup=default,o=sample
ibm-replicationfailedchanges=1 20050407202221Z 68 1
170814 add cn=entry-85,o=sample
ibm-replicationfailedchangecount=1
```

You can use the **idsldapexop** command to show data for the update, retry the update, or remove the update from the replication error log. Use the following **idsldapexop** command to show data for the failed update:

```
ldapexop -D <adminDN> -w <adminPW> -op controlreplerr -show 1 -ra
 cn=<server>-1389,ibm-replicaServerId=<server>-389,
   ibm-replicaGroup=default,o=sample
```

This command can return output similar to the following:

```
dn: entry-85,o=sample
cn: entry-85
objectclass: person
objectclass: eperson
objectclass: organizationalperson
objectclass: inetorgperson
objectclass: top
userpassword: {AES256}tDO9yQT540xpp7ZMIg95mA==
sn: user
ibm-entryuuid: bf201fcb-758e-41dc-bdea-1855fe0b860b
control: 1.3.6.1.4.1.42.2.27.8.5.1 false
control: 1.3.18.0.2.10.19 false::
  MIQAAADJMIQAAAAnCgEAMIQAAAAeBAxjcmVhdG9yc05hbWUxhAAAAAoECENOPUFETUlOMIQAAA
AxCgEAMIQAAAAoBA9jcmVhdGVUaW1lc3RhbXAxhAAAABEEDzIwMDUwMzMwMjMyNzQ3WjCEAAAAKA
oBADCEAAAAHwQNbW9kaWZpZXJzTmFtZTGEAAAACgQIQ049QURNSU4whAAAADEKAQAwhAAAACgED2
1vZGlmeVRpbWVzdGFtcDGEAAAAEQQPMjAwNTAzMzAyMzI3NDda
```

You can also use the **idsldapexop** command to retry the update. The following command:

```
ldapexop -D <adminDN> -w <adminPW> -op controlreplerr -retry 1 -ra
 cn=<server>-1389,ibm-replicaServerId=<server>-389,
   ibm-replicaGroup=default,o=sample
```

can return output similar to the following:

```
Operation completed successfully.
```

This result indicates only that it was possible to send the update again, not that the update was successful.

If you run the **idsldapsearch** command again:

```
idsldapsearch -D <adminDN> -w <adminPW> -h <servername>
-p <portnumber> -b " " -s sub objectclass=ibm*nt
ibm-replicationfailedchanges ibm-replicationfailedchangecount
```

the search can return output similar to the following:

```
cn=<server>-1389,ibm-replicaServerId=<server>-389,
ibm-replicaGroup=default,o=sample
ibm-replicationfailedchanges=2 20050407214939Z 68 2
 170814 add cn=entry-85,o=sample
ibm-replicationfailedchangecount=1
```

Notice that the update has failed again. The error ID is now 2, the number of attempts is 2, and the last time and result code have been updated.

Use the **idsldapexop** command to remove the failed update from the replication error log:

```
idsldapexop -D <adminDN> -w <adminPW> -op controlreplerr -delete 2 -ra
 cn=<server>-1389,ibm-replicaServerId=<server>-389,
   ibm-replicaGroup=default,o=sample
```

This command can return output similar to the following:

```
Operation completed successfully.
```

If you run the **idsldapsearch** command again:

```
idsldapsearch -D <adminDN> -w <adminPW> -h <servername>
-p <portnumber> -b " " -s sub objectclass=ibm*nt
    ibm-replicationfailedchanges ibm-replicationfailedchangecount
```

the search can return output similar to the following:

```
cn=<server>-1389,ibm-replicaServerId=<server>-389,ibm-replicaGroup=default,o=sample
ibm-replicationfailedchangecount=0
```

It is also possible to retry and delete all failures by using `all` in place of the error ID.

Note: Do not confuse the change ID, which is constant, with the error ID, which is changed on every failed attempt.

Lost and found log

The lost and found log (lostandfound.log) archives entries replaced due to replication conflict resolution. Logging these entries allows you to recover the data in the replaced entries if necessary. The information logged for each replaced entry includes:

- The DN of the entry that is archived as a result of conflict resolution
- The type of operation that results in the conflict; for example, add or delete.
- The time the entry was created
- The time the entry was last modified
- The TCP/IP address of the supplier whose update caused the conflict
- The LDIF representation of the entry associated with the failed update, including all the operational attributes, such as ibm-entryUUID.

The difference between write and replicated write messages

In a Tivoli Directory Server environment, you may see informational message, such as

```
GLPSRV202I During the last hour 40 updates were received from suppliers
and 10 updates were received from other clients.
```

This message in the ibmslapd.log file indicates that a directory server is participating as a peer server in a replication network of directory servers. A peer server is capable of receiving updates from other peer servers and from client applications. A standalone master server will show zero updates from other suppliers but can have some updates from clients depending on the update activity in a given hour. A directory server configured only as a replica will show some updates from suppliers and zero updates from clients. Updates sent to such a replica that were referred to a master server are not counted as updates from clients.

The message that shows updates from suppliers and clients can serve as a possible informational message to indicate that replication conflicts may occur. There can also be cases where the updates from clients and suppliers are for entries in different replication contexts and no conflicts may occur. Depending on the replication topology it is also possible that updates from clients are being routed to different master servers configured as peers. This has the potential for causing conflicts, particularly when the updates are for groups. Replication conflict resolution will ensure that the data across the multiple servers converges, but some updates will be overwritten. It is advisable to have updates for a particular replication context sent to a single server even when peer servers are available.

Understanding the behavior of the objectclass ibm-replicaSubentry (ReplicaSubEntry) in a replication topology

When a directory server that is in a replication environment starts, it compares its serverID against those in the replicaSubEntry entry. If the serverID matches, then as per the attribute value of ibm-replicationServerIsMaster, server either plays the role of a supplier or consumer. If the serverID does not match with any of the serverIDs given in the replicaSubEntry entry, then server assumes that it is consumer in a replication topology.

If replicaSubEntry is defined, then the respective serverID provided with the attribute ibm-replicaServerId becomes supplier or consumer depending on the attribute value of ibm-replicationServerIsMaster.

For example:

```
cn=server1,ibm-replicaGroup=default,o=ibm,c=us
objectclass : top
objectclass : ibm-replicaSubentry
ibm-replicaServerId : Server1
ibm-replicationServerIsMaster : TRUE
cn : server1
```

Here replicaSubEntry means, server with serverID, server1, is a supplier server in the replication topology.

```
cn=server2,ibm-replicaGroup=default,o=ibm,c=us
objectclass : top
objectclass : ibm-replicaSubentry
ibm-replicaServerId : Server2
ibm-replicationServerIsMaster : FALSE
cn : server2
```

Here replicaSubEntry means, server with serverID, server2, is a consumer server in the replication topology.

Note: If replicaSubEntry is not present for a server in a replication topology, then it is assumed that the server is a consumer in a replication topology.

Viewing replication status using command line utilities

To help determine issue pertaining to replication, it is important to view the status of a replication agreement. You can use command line utilities to view the appropriate replication associated operational attributes. The two special attributes , “+ibmrepl” and "++ibmrepli", are defined to request replication related operational attributes in a search. The "+" and "++" are subsets of the operational attributes. The single "+" is less expensive. The "++" includes all operational attributes shown in the "+" attribute list and the those listed in the "++" column as shown in the table below.

Table 3. Replication related operational attributes

Attribute	"+" Attribute list	"++" Attribute list
+ibmrepl	ibm-replicationChangeLDIF ibm-replicationLastActivationTime ibm-replicationLastChangeId ibm-replicationLastFinishTime ibm-replicationLastResult ibm-replicationLastResultAdditional ibm-replicationNextTime ibm-replicationPendingChangeCount ibm-replicationState ibm-replicationFailedChangeCount ibm-replicationperformance	++ibmrepl includes the attributes from +ibmrepl and add the following: ibm-replicationPendingChanges ibm-replicationFailedChanges

To search a specific replication agreement, issue the ldapsearch of the following format:

```
idsldapsearch -h <hostname> -p <port> -D cn=<adminDN> -w <adminPW> \
   -b <ReplicationAgreement> objectclass=* ++ibmrepl
```

For example,

```
idsldapsearch -h peer1 -p 1389 -D cn=root -w password -b cn=peer2:2389,\
  cn=peer1:1389,ibm-replicaGroup=default,O=sample objectclass=* ++ibmrepl
```

To search all agreements, issue the ldapsearch of the following format:

```
idsldapsearch -h <hostname> -p <port> -D cn=<adminDN> -w <adminPW> \
   -s sub -b " " objectclass=ibm-replicationagreement ++ibmrepl
```

For example,

```
idsldapsearch -h peer1 -p 1389 -D cn=root -w password -s sub -b " " \
  objectclass=ibm-replicationagreement ++ibmrepl
```

To know more about replication status, see Monitoring replication status in the *IBM Tivoli Directory Server Version 6.3 Administration Guide*.

Using the environment variable IBMSLAPD_REPL_UPDATE_EXTRA_SECS

The default timeout for any change to be completed through replication is 60 seconds. If replication updates involve large amount of changes, such as adding of a large group entry, or adding or modifying entries that contains large objects such as credentials, then the update operation may require more than 60 seconds for the operation to finish. If any single update (add, delete, modify, or modifydn) operation through replication takes more than 60 seconds, then the supplier server times out that update operation and retries again sending the same update through replication. Inorder to extend the timeout duration for update operations in replication, you can use the IBMSLAPD_REPL_UPDATE_EXTRA_SECS environment variable.

The IBMSLAPD_REPL_UPDATE_EXTRA_SECS environment variable need to be added to the supplier servers in a replication topology. A valid value must be provided for the environment variable to extend the timeout value, this value is added to the default timeout value of 60 seconds. The valid values for this variable is as follows:

- Minimum: 1
- Maximum: 2147483647

Note: For optimal result, a value of 180 is preferred for the variable. Setting the variable with a value greater than 600 is not preferred. Determine a value best suited for your environment by testing the same update from a direct client against the consumer server from the supplier server.

The IBMSLAPD_REPL_UPDATE_EXTRA_SECS environment variable can be set either be adding the variable to the configuration file or using the command prompt.

Adding the variable to configuration file

Using the LDAP client utility:

1. Issue the ldapmodify command of the following format against the supplier server:

```
idsldapmodify -p <port> -D <adminDN> -w <adminPW>
dn: cn=Front End, cn=Configuration
changetype: modify
add: ibm-slapdSetenv
ibm-slapdSetenv: IBMSLAPD_REPL_UPDATE_EXTRA_SECS=180
```

2. Restart the directory server instance.

Using the Web Administration Tool:

1. Ensure that ibmslapd and ibmdiradm processes are running for the directory server instance.
2. Log on to the directory server instance using the Web Administration Tool.
3. From the left navigation area, expand Directory management and then click Manage entries.
4. On the Manage entries panel, expand cn=configuration, and then select cn=Front End and click Edit attributes.
5. On the Edit and entry panel, click Next to open the Optional attributes panel.
6. Click the Multiple values button next to the ibm-slapdSetenv field.

7. In the resulting panel, enter IBMSLAPD_REPL_UPDATE_EXTRA_SECS=180 in the ibm-slapdSetenv field and then click Add.
8. To save, click OK.
9. To effect the changes made, restart the directory server instance.

From command prompt

Set the environment variable IBMSLAPD_REPL_UPDATE_EXTRA_SECS.

On AIX, Linux, and Solaris systems (ksh shell):

```
export IBMSLAPD_REPL_UPDATE_EXTRA_SECS=180
```

On Windows systems:

```
set IBMSLAPD_REPL_UPDATE_EXTRA_SECS=180
```

For the set value of the environment variable to be effective, restart the directory server instance from the same shell where the above environment variable was set.

Replication Troubleshooting

The following sections contain troubleshooting information about replication

Replicated suffix must have ibm-replicationcontext object class

Before loading your database, make sure the ibm-replicationcontext object class exists for the suffix. If you load your data before setting the object class, you might receive an error similar to the following

```
08/13/04 15:32:34 For the replica group entry
ibm-replicaGroup=default,o=sample, the parent entry
must be an ibm-replicationContext entry.
08/13/04 15:32:34 Parent entry does not exist for entry
cn=urchin,ibm-replicaGroup=default,o=sample.
08/13/04 15:32:34 Entry cn=replication,cn=localhost already exists.
08/13/04 15:32:35 Parent entry does not exist for entry
cn=superman.tivlab.austin.ibm.com,cn=urchin,
                      ibm-replicaGroup=default,o=sample.
```

To add the ibm-replicationcontext object class to the suffix, run the following command:

```
ldapmodify -D cn=root -w secret -f mod.ldif
```

where the mod.ldif file contains:

```
dn: o=sample
changetype: modify
add: objectclass
objectclass: ibm-replicationcontext
```

Verify that suffixes and replication agreements exist using idsldapsearch

If you are experiencing errors with replication, run the following commands to verify that your suffixes are configured to be replicated and that the replication agreements exist.

Run the following command to verify that the context exists with replication agreements:

```
idsldapsearch -D cn=root -w secret -b o=sample objectclass=ibm-repl*
```

where "o=sample" is the replication context.

If this command does not return any results, the suffix is not configured to be replicated. You must configure the suffix to be replicated. See the *IBM Tivoli Directory Server Version 6.3 Administration Guide* for instructions for configuring a suffix for replication.

Run the following command to verify that the replication agreements exist:

```
idsldapsearch -D cn=root -w secret -b <replctxt>
              objectclass=ibm-replicationAgreement
```

where *replctxt* is the location where the replication agreements for a replication context are stored; for example, o=sample. If the command does not return results, the replication agreement might not exist. In order to replicate correctly, the correct replication agreements must exist. See the *IBM Tivoli Directory Server Version 6.3 Administration Guide* for instructions for adding replication agreements.

Peer to peer replication returns error "No such object occurred for replica"

If you are running peer-to-peer replication, you might encounter an error similar to the following:

```
 09/07/04 12:57:10 Error No such object occurred for replica '<CN=SERVER2>,
<CN=SERVER3>,IBM-REPLICAGROUP=DEFAULT,O=IBM': modify failed for entry
'<CN=MISSING_ENTRY>' change ID 5109011.
```

where *CN=SERVER2* and *CN=SERVER3* are the peer servers and *CN=MISSING_ENTRY* is the entry on which the error occurred.

One common cause of this error is that peer-to-peer replication, by design, does not allow for conflict resolution.

To correct this error, do the following:

1. Locate the entry listed under the "No such object occurred for replica" error in the Server log (ibmslapd.log).
2. Use the **idsdb2ldif** command to export the entry or entries in the log from the peer server on which the error or errors occurred; for example:

   ```
   idsdb2ldif -o <out.ldif> -I <instance name> -s <subtree DN>
   ```

   where:

   - *out.ldif* is the name of the file to which you want to export the entry.
   - *instance name* is the name of the instance.
   - *subtree DN* is the DN of the entry you want to export.
3. Use the **idsldapadd** command to import the entry to the other peer server; for example:

   ```
   idsldapadd -D cn=root -w secret -i <out.ldif>
   ```

   where *out.ldif* is the name of the file containing the entry you want to import.

Replication returns error "Insufficient access"

When a replication topology extended operation is issued to a Release 6.0 and later version of directory servers and the server's consumer is a pre 6.0 server, then the operation fails. In the server trace, `Insufficient Access` can be identified as the cause of the failure.

When a replication topology extended operation is issued to a server, the server propagates all of its replication topology entries to its consumers. However, the consumers must be of Release 6.0 and later versions. For a consumer of pre 6.0 release to have exactly the same replication topology entries as its supplier, import and export tools, such as **idsdb2ldif** and **idsldif2db**, can be used.

Replication topology extended operation returns result code 80

You might see following message after running a replication topology extended operation:

```
Operation failed with result code 80.
Details: "x" servers replicated successfully out of "y" attempts.
```

where *x* is not equal to *y*.

If this occurs, check for the following:

- If the replication context entry exists on the consumer server, be sure that the replication context entry has an objectclass of ibm-replicationContext. Alternatively, you can delete the replication context entry so that the supplier can propagate all of its replication topology-related entries, including the replication context entry, to the consumer.
- After sending all the replication topology-related entries under a replication context to the consumer, the supplier of the extended operation sends the replication topology extended operation to the consumer in an effort to cascade the operation. If more than one tier of servers is involved in a replication topology, be sure that each supplier has the proper credential object to bind with its consumers.
- One of the consumer servers is down or not reachable at that instance.
- The replication context is a non-suffix entry and the consumer does not have the parent entry of the context.

  For example, suppose that cn=johndoe,cn=people,o=sample is the context for the topology you want to replicate. If o=sample is the suffix on the consumer and cn=people,o=sample does not exist, the operation will fail.
- The repltopology extended operation timed out on a heavily loaded consumer. (This results in message GLPRPL098E.)
- Suppose that a certain set of agreements already exists on the consumer. The repltopology extended operation attempts to delete these agreements and before that attempts to purge the queue associated with that agreement. If the purge fails, the extended operation fails. (This results in message GLPRPL093E.)

Replication command-line interface error (Windows systems only)

If you are using a Windows operating system and have a master server configured to do replication, you might see an error like the following in the ibmslapd error log during updates:

```
[IBM][CLI Driver] CLI0157E Error opening a file.  SQLSTATE=S1507
```

This problem can be resolved by adding the following entry to the \sqllib\db2cli.ini file:

```
[COMMON]
TempDir=x:\<your directory>
```

where x:\<your directory> specifies an existing directory on a drive that has space available. DB2 writes temporary files to this directory. The amount of space required depends on the size of the directory entries you are adding or updating, but generally, more space is required than the size of the largest entry you are updating.

Entries in LDIF file are not replicated

If you use the **idsldif2db** command with the **-r yes** option (to indicate that the entries in the file are to be replicated) and you find that entries are not being replicated, the following information might help you resolve the problem.

For the **-r yes** option to work for a server, the server must have a server ID defined in the configuration file. The server ID is created the first time the server starts if it is not already defined. In addition, the replication topology entries (especially the replication subentries) defined in the directory information tree in the LDIF file must match the server ID for the server to be able to replicate.

Ways in which problems can occur include the following:

- The server ID is not defined in the configuration file. This can happen when an instance is newly created and the **idsldif2db** command is used immediately after, before the server has started for the first time.
- The server ID is defined in the configuration file, but the replication subentries (attribute ibm-replicaServerId) defined in the directory information tree in the LDIF file do not match the server ID in the configuration file. If you change the ibm-replicaServerId attribute in the LDIF file to match the server ID in the configuration file and then run the **idsldif2db** command with the **-r yes** option, replication occurs correctly.

Problem when replicating or modifying the cn=ibmpolicies subtree

In Tivoli Directory Server V6.0 and later versions, a partial replication configuration entry is automatically added to the cn=ibmpolicies subtree. This is because the design of a directory server in V6.0 and later versions have the ibm-replicationcontext, ibm-replicagroup, and ibm-replicasubentry setup automatically for the cn=ibmpolicies subtree the first time the LDAP server is started. Also, the partial entries are created with the default ibm-slapdServerId value (randomly generated when the instance is first created). As users often modify the ibm-slapdServerId value in the ibmslapd.conf file after the initial configuration, this subtree often become read-only or might not get replicated properly. To resolve this, you should consider removing these partial replication entries from all machines in the topology:

To remove the entries, do the following:

1. Search the directory servers to get the entries. Issue the following command:

```
idsldapsearch -D cn=root -w secret -L -b cn=ibmpolicies objectclass=\
  ibm-replica*

dn: CN=IBMPOLICIES
cn: IBMpolicies
objectclass: container
```

```
objectclass: top
objectclass: ibm-replicationcontext

dn: ibm-replicagroup=default,cn=ibmpolicies
objectclass: top
objectclass: ibm-replicagroup
ibm-replicaGroup: default

dn: ibm-replicaserverid=ac1156c0-a214-1029-934c-cd9424fd6984,\
    ibm-replicagroup=
default,cn=ibmpolicies
objectclass: top
objectclass: ibm-replicasubentry
ibm-replicationserverismaster: TRUE
cn: V6.0 Migration
ibm-replicaServerId: ac1156c0-a214-1029-934c-cd9424fd6984
```

Note: The value, ac1156c0-a214-1029-934c-cd9424fd6984, for ibm-replicaserverid in the example output is a randomly generated serverID. For your system, the value will be different.

2. Delete the ibm-replicasubentry. Issue the following command:

```
idsldapdelete -D cn=root -w secret -k ibm-replicaserverid=\
ac1156c0-a214-1029-934c-cd9424fd6984,ibm-replicagroup=default,\
cn=ibmpolicies
```

3. Delete the ibm-replicagroup. Issue the following command:

```
idsldapdelete -D cn=root -w secret -k ibm-replicagroup=default,\
cn=ibmpolicies
```

4. Modify the cn=ibmpolicies entry to remove the entry "objectclass: ibm-replicationContext". Issue the following command:

```
ldapmodify -D cn=root -w secret -k
dn: cn=ibmpolicies
changetype: modify
delete: objectclass
objectclass: ibm-replicationContext
```

In the example above the -k option in the ldapmodify and ldapdelete command, which allows the admin user to modify objects in a read-only subtree. It might be necessary to also pass the -R option to not chase referrals.

After you have removed these entries, you can setup replication on the cn=ibmpolicies subtree just as you would on any other subtree.

Master server can become unstable or stop when serving to large number of replica servers

Master server can become unstable or stop when serving to large number of replica servers. This is because the master server might have run out of resources. To resolve this, you can set the Ulimits DN entry in the configuration file to the following:

```
dn: cn=Ulimits, cn=Configuration
cn: Ulimits
ibm-slapdUlimitDataSegment: -1
ibm-slapdUlimitDescription: Prescribed minimum ulimit option values
ibm-slapdUlimitFileSize: 2097151
ibm-slapdUlimitNofile: 500
ibm-slapdUlimitRSS: -1
ibm-slapdUlimitStackSize: -1
ibm-slapdUlimitVirtualMemory: -1
objectclass: top
objectclass: ibm-slapdConfigUlimit
objectclass: ibm-slapdConfigEntry
```

And then configure the system ulimit values to:

```
core file size      (blocks, -c)    unlimited
data seg size       (kbytes, -d)    unlimited
file size           (blocks, -f)    unlimited
max memory size     (kbytes, -m)    unlimited
open files          (-n)            30000
pipe size           (512 bytes, -p) 64
stack size          (kbytes, -s)    unlimited
cpu time            (seconds, -t)   unlimited
max user processes  (-u)            262144
virtual memory      (kbytes, -v)    unlimited
```

Restart the servers for the changes to take effect.

Stopping a multithreaded replication supplier

In a replication environment, abruptly stopping a supplier that uses multithreaded replication to accelerate replication between its consumers can cause problems. At any given time a supplier might have sent multiple updates to a consumer, which can be the number of consumer connections multiplied by the depth of replication receive queue. A supplier can also have multiple consumers. In such cases, the number of replication updates at any given time can be large. The replication status of the supplier is based on the most recent change replicated (successfully or otherwise) so far. If a supplier is restarted, it will use this replication status value to determine the changes that need to be sent to the consumer. If a supplier is stopped abruptly before it receives a response from its consumers for the updates that it had sent, then the updates will be sent again by the supplier. When these updates are sent again, it can cause errors to be reported. These errors will be logged by the supplier and can be cleared using the Web administration tool or the command line utility for managing the replication error log. If replication error logging is enabled, then these errors will be counted against the limit for logged errors so they should be cleared.

If replication error logging is not enabled, these kinds of errors that occur should not block replication. The replicated add operations will report that the entries already exist, modify operation will report that the attribute values exist or are not found, and delete operations will report that the entries no longer exist. The consumer server might log these errors but replication should continue.

In some cases, the administrator might not be aware of the problem and therefore may not be able to resolve the problem, if a replicated update from a supplier to its consumer results in an error and the supplier was not available to log the error. Depending on the last response received by the supplier, this update may not be replicated again.

To avoid having replication related errors, perform these steps before stopping a supplier server.

1. Find the server ID of a supplier. An example of the ldapsearch command with its output:

   ```
   #ldapsearch -D cn=admin -w password -p 2389 -s base objectclass=* ibm-serverID

   ibm-serverId=wingspread-2389
   ```

2. Find all the replication subentries with the server ID obtained in step 1.

   ```
   #ldapsearch -D cn=admin -w password -p 2389 -s sub \
                               ibm-replicaServerId=wingspread-2389 0.0

   ibm-replicaServerId=wingspread-2389,ibm-replicaGroup=default,O=SAMPLE
   ```

3. Find all the replication agreements for the server.

```
#ldapsearch -D cn=admin -w password -p 2389 -b ibm-replicaServerId=wingspread-2389,\
    ibm-replicaGroup=default,O=SAMPLE objectclass=ibm-replicationAgreement 0.0

cn=wingspread-1389,ibm-replicaServerId=wingspread-2389,ibm-replicaGroup=default,
O=SAMPLE
```

4. Verify the status for a given replication agreement.

```
#ldapsearch -D cn=admin -w password -p 2389 -b cn=wingspread-1389,\
   ibm-replicaServerId=wingspread-2389,ibm-replicaGroup=default,O=SAMPLE \
   objectclass=* +ibmrepl

cn=wingspread-1389,ibm-replicaServerId=wingspread-2389,
   ibm-replicaGroup=default,O=SAMPLE
ibm-replicationChangeLDIF=N/A
ibm-replicationLastActivationTime=20080707152436Z
ibm-replicationLastChangeId=4855
ibm-replicationLastFinishTime=20080707152436Z
ibm-replicationLastResult=N/A
ibm-replicationLastResultAdditional=N/A
ibm-replicationNextTime=N/A
ibm-replicationPendingChangeCount=0
ibm-replicationState=ready
ibm-replicationFailedChangeCount=0
ibm-replicationperformance=
  [c=0,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
ibm- replicationperformance=
  [c=1,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
ibm-replicationperformance=
  [c=2,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
ibm-replicationperformance=
  [c=3,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
```

   The “+ibmrepl” in search filter returns operational attributes related to replication. The attribute names appear to the left of the equal signs. In the example, there are 4 connections to the consumer. Some replication status information attributes are only used with single threaded replication, (displayed with the value 'N/A'), others are only for multiple threaded replication. Use "++ibmrepl" to show all the attributes including the pending changes and logged replication errors.

5. Suspend replication for the agreement.

```
# ldapexop -D cn=admin -w password -p 2389 -op controlrepl -action suspend \
   -ra cn=wingspread-1389,ibm-replicaServerId=wingspread-2389,\
   ibm-replicaGroup=default,O=SAMPLE

Operation completed successfully.
```

6. Verify the status of replication agreement.

```
#ldapsearch -D cn=admin -w password -p 2389 -b cn=wingspread-1389,\
   ibm-replicaServerId=wingspread-2389,ibm-replicaGroup=default,O=SAMPLE \
   objectclass=* ++ibmrepl

cn=wingspread-1389,ibm-replicaServerId=wingspread-2389,
   ibm-replicaGroup=default,O=SAMPLE
ibm-replicationChangeLDIF=N/A
ibm-replicationLastActivationTime=20080707152648Z
ibm-replicationLastChangeId=4855
ibm-replicationLastFinishTime=20080707152648Z
ibm-replicationLastResult=N/A
ibm-replicationLastResultAdditional=N/A
ibm-replicationNextTime=N/A
ibm-replicationPendingChangeCount=1
ibm-replicationState=on hold
ibm-replicationFailedChangeCount=0
ibm-replicationperformance=
  [c=0,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
```

```
ibm-replicationperformance=
  [c=1,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
ibm-replicationperformance=
  [c=2,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
ibm-replicationperformance=
  [c=3,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
ibm-replicationPendingChanges=4856 modify CN=WINGSPREAD-1389,
 IBM-REPLICASERVERID=wingspread-2389,IBM-REPLICAGROUP=DEFAULT,O=SAMPLE
```

The pending change reported in the output was caused by the operation to suspend replication.

7. Determine if there are any replicated updates that have been sent to the consumer. In the output from step 6 related to the replication status, the attribute "ibm-replicationperformance" can be used to determine the number of updates that have been sent to the consumer. This attribute only applies to replication agreements using multithreaded replications.

   The information about the data associated with the ibm-replicationperformance attribute in the example output of step 6 is as follows:

   - c: This indicates the connection number. In the output of step 6, there are 4 connections. The first connection will show the most traffic. The workload determines how often the other connections are used.
   - l: This indicates the size limit for each queue. In the example, the value is 10. For each connection, there are two queues of same length. A queue for updates to be sent on the connection called the send queue, and the other queue for updates that have been sent but no response has been received from the consumer called the receive queue.

     When updates are sent, they are moved from send queue to the receive queue. When the receive queue reaches its size limit, no more updates are sent until some responses from the consumer are received. When the send queue reaches its size limit, no more updates are assigned to the connection. When the size limit for all the send queues of connections are reached, the supplier waits for the consumer to process the backlog.
   - op: This specifies the replication ID of the last operation assigned to the send queue of the connection. Replication IDs are assigned to all updates that are to be replicated to a consumer. The process of assigning replication IDs should not stop even if replication is suspended.
   - q: This specifies the current size of the send queue. This value should not change if replication is suspended.
   - d: This specifies the count of dependent updates. For example, an add request for an entry followed by a modify request of the same entry is counted as a dependency. All dependent updates must be assigned to the same connection so that they can be applied in correct order.
   - ws: This indicates the number of times the send queue reached its size limit.
   - ds: This specifies the number of dependent updates sent.
   - wd: This specifies the number of times the send queue waited for a dependent update before sending additional updates.
   - wr: This indicates the number of times the receive queue reached its size limit.
   - r: This indicates the number of replicated updates that is waiting for a response from the consumer.
   - e: This specifies the number of replication errors reported by the consumer.

- ss: This indicates the session count of the sender thread. It is incremented when a connection to the consumer is established.
- rs: This indicates the session count of the receiver thread.

If the number of replicated updates waiting for a response is 0 (indicated by "r"), then for this consumer server it is safe to stop the supplier. The value of "r" will vary between the value of "l" (size limit of the queue, which is default to 10) and 0. If the value is not 0 for "r", you should wait for it to be 0. The value of "r" depends on the size and type of the replication update and the workload on the consumer. Once this value is 0, the supplier will send the current status of updates to the consumer. On restarting the supplier, it will replicate only the updates that have not been sent before.

8. Repeat the steps 4 through 7 for each replication agreement serviced by the supplier.
9. Stop the supplier server when the replicated update status is 0 for all the consumers.
10. After restarting the supplier, resume the replication that was suspended in step 5.

```
#ldapexop -p 2389  -D cn=admin -w password -op controlrepl -action resume \
-ra cn=wingspread-1389,ibm-replicaServerId=wingspread-2389,\
ibm-replicaGroup=default,O=SAMPLE

Operation completed successfully.
```

You can also use Web administration tool to see the status of a replication agreement and suspend or resume replication. Use a similar approach to determine when there are not any replicated updates whose status has not been reflected on the supplier.

Synchronizing Tivoli Directory Servers in a replicated environment

If the Tivoli Directory Servers in a replicated environment become out of synch, the replication queues might get blocked. To resolve this, you must resynchronize your replicated environment. Consider a scenario, where M1 is the master server with the most recent updated data, and R1 and R2 are the two replica servers of the master server, M1. To resynchronize the directory servers, perform the following steps:

1. Take R1 and R2 offline by stopping the R1 and R2 servers.
2. Quiesce M1 for all queues.
3. Clear the queues on M1 to R1 and M1 to R2. Repeat this process for all the queues. Using the GUI, the Web administration tool, click Manage queues under the Replication management category in the navigation area. On the Manage queues wizard, click Queue details. On the Queue details panel, click Pending changes and then click Skip All Blocking Entries.
4. Export M1's data to a file. Issue the following command:

   ```
   idsdb2ldif -o /tmp/M1.ldif
   ```
5. Unquiesce the M1 server.
6. Unconfigure and drop the database on R1 and R2. Make sure you answer yes to destroy the database. Issue the command of the following format:

   ```
   idsucfgdb -I <instance_name> -r
   ```
7. Configure the database on R1 and R2. Issue the command of the following format:

```
idscfgdb -I <instance_name> -a <dbadminDN> -w <dbadminPW> -t <databasename> \
  -l <dblocation> -n
```

8. Synchronize the modified schema. This can be done by copying the V3.modifiedschema from M1 over to R1 and R2. The modified schema, V3.modifiedschema, is located in the *<instance_home>/idsslapd-<instance_name>/etc* directory.
9. Synchronize the ibmslapddir.ksf file. To know more about Synchronizing two-way cryptography between server instances, see the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide*.

   Note: Only if the master and the replicas are on the same hardware and operating system, the ibmslapddir.ksf file can be copied over from master to replicas. The ibmslapddir.ksf file is located in the *<instance_honme>/idsslapd-<instance_name>/etc* directory.
10. Copy the M1.ldif file to replicas and load M1's data onto R1 and R2. Issue the following command:

    ```
    idsldif2db -i /tmp/Master.ldif -r no
    ```
11. Start the R1 and R2 servers.

Note: In the case of Windows platform, change the paths accordingly. Alternatively, you can use the ldapdiff or idsideploy utility to synchronize between a master and replica server depending on your Tivoli Directory Server environment. The ldapdiff utility identifies differences in a replica server and its master, and can be used to synchronize replicas. The idsideploy utility with the -r and -L options can be used to synchronize a peer-peer or peer-replica servers. User can create the target instance either as a peer or replica of the master server with the -r option. The -L option provides the restore location from which the source instance's backed up database can be restored on to the target instance (peer or replica). To know more about the ldapdiff or idsideploy utility, see the *IBM Tivoli Directory Server Version 6.3 Command Reference*.

Using multimaster configurations

In Tivoli Directory Server, when configuring with several peer masters the configuration should be such that updates for the same entry or set of entries can not occur to several peer masters at the same time. The replication system can be configured in such a way that all writes go to one master except in the case of failover, or the system can be configured so that all writes for a given subtree go to one master except in the case of failover. If writes of the same entry occur on several masters then before such write can be replicated, an update conflict might occur.

Tivoli Directory Server can be configured with conflict resolution. This ensures that for almost all update conflicts the latest change to a given entry will be preserved so that the content of all servers will converge to the same value for the entry. However, update conflicts should be avoided. This is because conflict resolution might cause inherent loss of data, the later change to the entry is preserved but the earlier change is discarded. Conflict resolution can also affect replication performance, if the number of conflicts observed is large.

Sometimes, it is not possible to avoid configurations where update conflicts can occur. For example, there may be Tivoli Directory Servers at two sites and because of a temporary loss of network connectivity between the sites, all writes occurring at a given site may occur on the server for that site. Update conflicts may occur as a result, and the Tivoli Directory Server conflict resolution procedures will then

converge the content of entries on the servers. However, in most configurations, nearly all update conflicts can be avoided.

If conflict resolution is to be used, it is important to load the directory server using a procedure that ensures that the timestamps for the created entries are the same on all of the servers in the topology at the outset. There are two ways of doing this, they are as follows:

- Loading a directory server using bulkload and then back up the database and restore that database backup on the other servers.
- Loading a directory server using bulkload and then extract an LDIF file including timestamps from this server using the db2ldif command. Thereafter, bulkload the resulting LDIF file onto the other servers in the replication topology.

On the Replication panel of the Web administration tool the options to specify replication filter and replication method are not available when creating a master server

When creating a master server in a replication topology using the web administration tool, the options to specify the replication filter and replication method for the master server is not available. This is a limitation in the Web administration tool. The replication filter contains the existing filters under the selected subtree and the replication method specifies the type of replication, single-threaded or multi threaded.

To specify the replication filter and replication method options in a peer-to-peer replication, click the button next to the peer server and then click Edit agreement. In the Edit agreement panel of the peer server, specify the replication filter and replication method that you want to set and then click Apply.

Replication between a Tivoli Directory Server version 6.3 supplier and a downlevel consumer server that does not support SHA-2

When replication is set up between a Tivoli Directory Server v6.3 supplier server that has SHA-2 family of encryption algorithm as the configured password encryption method and a downlevel consumer server that does not support SHA-2 family of encryption algorithm, then in this case the supplier will log a message stating that the replication operation cannot be started with the consumer and will set the replication state to "connecting". Similarly, if attribute level encryption is set with SHA-2 family of encryption scheme on Tivoli Directory Server v6.3, replication to downlevel consumer server will fail to start and the replication state will be set to "connecting".

Chapter 10. Troubleshooting performance

If you are experiencing problems with the performance of your directory server, refer to this section for possible fixes and workaround.

Identifying performance problem areas

This section contains some methods for identifying areas that might be affecting the performance of your directory server.

Server audit log

The Server audit log shows what searches are being performed and the parameters used in each search. The Server audit log also shows when a client binds and unbinds from the directory. Observing these measurements allows you to identify LDAP operations that take a long time to complete.

idsslapd trace

An idsslapd trace provides a list of the SQL commands issued to the DB2 database. These commands can help you identify operations that are taking a long time to complete. This information can in turn lead you to missing indexes, or unusual directory topology. To turn the idsslapd trace on, run the following commands:

1. `ldtrc on`
2. `idsslapd –h 4096`

After you have turned the trace on, run the commands that you think might be giving you trouble.

Running a trace on several operations can slow performance, so remember to turn the trace off when you are finished using it:

```
ldtrc off
```

Adding memory after installation on Solaris systems

Memory added to a computer after the installation of a Solaris operating system does not automatically improve performance. To take advantage of added memory, you must:

1. Update the shared memory (shmem) value in the /etc/system file:

   ```
   set shmsys:shminfo_shmmax = physical_memory
   ```

   Where *physical_memory* is the size on of the physical memory on the computer in bytes.

   You must restart the computer for the new settings to take effect.
2. From the command line, set the ulimit values for increasing process memory and file size to unlimited:

   ```
   ulimit -d unlimited
   ulimit -v unlimited
   ulimit -f unlimited
   ```

Setting the SLAPD_OCHANDLERS environment variable on Windows

On Windows, if you have clients that are generating many connections to the server and the connections are being refused, set the SLAPD_OCHANDLERS environment variable to 15 before starting the server.

Error messages similar to the following might be logged in the idsslapd.log file:

```
Feb 11 14:36:04 2004  Communications error: Exceeding 64
    connections/OCH - dropping socket.
```

If you see these errors, do the following:

1. Save a copy of your ibmslapd.conf file.
2. Insert the following in the section that starts with `'dn: cn=FrontEnd,cn=Configuration'`:

   ```
   ibm-slapdSetenv: SLAPD_OCHANDLERS=15
   ```
3. Stop and restart the server.

DB2 rollbacks and isolation levels

If you are experiencing rollback activities in DB2, check the isolation level. Rollbacks occur when one application process has a row locked while another application process tries to access that same row. Because the default isolation level, repeatable read, can result in more rows being locked than are actually required for the current read request, a more relaxed isolation level is normally required for LDAP applications.

For example, the read stability isolation level allows other applications to insert or update data in rows that have been read. If a second read is issued for that range of rows, the new data is reflected in the result set. Keep in mind, however, that the second read can return data that is different from the first read. If an application depends upon the same data being returned on multiple reads, the isolation level should be set to repeatable read.

To set the DB2 isolation level, type the following at a command prompt:

```
db2set <isolation_level>=YES
```

where *isolation_level* is the isolation level you want to apply, such as DB2_RR_TO_RS.

Note: All applications using the current database instance are affected by this setting.

Default value of LOGFILSIZ needs to be increased

If you are adding a very large group (more than 50,000 members) to your directory, and you have migrated your database from a previous release, modify the LOGFILSIZ parameter of your DB2 database to be at least 2000. On migrated databases, this value might currently be set to 750 or 1000.

You can verify this value by issuing the following commands. For this example the names of the user, instance, and database are **ldapdb2**.

For AIX, Linux, and Solaris platforms:

```
su - ldapdb2
db2start
db2 get database config for ldapdb2 | grep LOGFILSIZ
```

To increase this value, issue the following command:

```
db2 update database config for ldapdb2 using LOGFILSIZ 2000
db2 force applications all
db2stop
db2start
```

For Windows platforms:

```
db2cmd
set DB2INSTANCE=ldapdb2
db2 get database config for ldapdb2 <outputfile>
```

Find the value for LOGFILSIZ in the output file. To increase this value, issue the following command:

```
db2 update database config for ldapdb2 using LOGFILSIZ 2000
db2 force applications all
db2stop
db2start
```

Note: This value is already set correctly if you created or configured your database with the Configuration Tool.

Auditing for performance profiling

In order to identify performance bottlenecks during operation execution, you can check the server audit log for the summary figures indicating performance hotspots. The following hotspots are identified for auditing:

- When an operation has to wait in the worker thread queue for a long time before the worker thread actually starts executing the operation.
- The time spent for cache contention inside the backend needs to be tracked.
- The time spent in handling client I/O, that is, the time spent in receiving the request and returning the result. This value can also be used for detecting bottlenecks because of slow clients or network issues.

Using the audited performance hotspot data, directory administrators can use the system audit facilities to log the LDAP audit record with the system-defined record format.

While auditing the performance profiling, the following points should be considered:

- The configuration options can be enabled to auditing for a combination of different types of operations, for example, auditing for add and modify operations only, along with the auditing for performance.
- At the end of operation execution, the audit information is stored in the server audit logs only. In a scenario where the server is having performance bottlenecks and is in a hung state, the cn=workers, cn=monitor search can be issued. This search gives information about where each worker is stuck, which is obtained by accumulating information collected about the worker till that point in the audit records.

For each operation, performance data field in the audit records is controlled using the configuration option “ibm- auditPerformance”. Currently, the following performance data fields will be defined for each operation:

operationResponseTime
This field represents the time difference in milliseconds between the time

the operation was received and the time its response was sent. The operation received time and the response sent time of an operation are published in audit v3 header.

timeOnWorkQ
This field represents time in milliseconds spent in the worker queue before execution is initiated on the operation. The value of this field is the difference between the time execution was initiated and the time the operation was received.

rdbmLockWaitTime
This field represents time in milliseconds spent in acquiring locks over RDBM caches during operation execution. The value in this field helps administrators to determine the time spent for cache contention against real work.

The lock wait time over the following resources are also considered.

- Resource cache
- DN cache
- Entry cache
- Filter cache
- Attribute cache

  Note: Starting with the IBM Tivoli Directory Server 6.3 release, attribute cache is deprecated. Henceforth, users should avoid using attribute cache.
- Deadlock detector
- RDBM locks

This is implemented by introducing a field in the operation structure, which is updated when acquiring of lock is attempted during operation execution. In addition, wrapper functions are introduced for functions that attempt to acquire locks over RDBM caches. These wrapper functions take another operation pointer as parameter and update the operation's lock wait time field if ibm-auditPerformance is enabled.

clientIOTime
This field represents time in milliseconds that was spent in receiving the complete operation request and returning the complete operation response. This field is implemented in the operation structure and is updated on receiving the complete BER for operation request and on successfully returning the response BER message for the operation.

An example of the audit version 3 format for search operation with ibm-auditPerformance enabled will look like:

```
AuditV3--2006-09-09-10:49:01.863-06:00DST--V3 Search--
bindDN: cn=root--client: 127.0.0.1:40722--connectionID: 2--
received: 2006-09-09-10:49:01.803-06:00DST—Success
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: o=sample
scope: wholeSubtree
derefAliases: neverDerefAliases
typesOnly: false
filter: (&(cn=C*)(sn=A*))
operationResponseTime: 591
timeOnWorkQ: 1
rdbmLockWaitTime: 0
clientIOTime: 180
```

In order to control server performance hits while collecting information for performance data fields, the ibm-auditPerformance field is introduced in the audit configuration section. The value of the ibm-auditPerformance field is false, by default and therefore no performance data will be collected and published by default. When the value of the ibm-auditPerformance field is set to true, performance data will be collected and published in the audit logs for each operation that is enabled to be audited. If the ibm-auditPerformance field is enabled, that is, set to true, in audit record section the four performance data fields are audited: operationResponseTime, timeOnWorkQ, rdbmLockWaitTime, and clientIOTime. The values of these performance data fields are times in milliseconds.

Chapter 11. Troubleshooting scenarios

This chapter contains some troubleshooting scenarios you might encounter and provides some solutions.

Server is not responding

If the server appears to not respond, first verify whether the server is truly not responding, or simply performing very slowly.

To determine if the server is suffering from poor performance, follow the directions in the *IBM Tivoli Directory Server Version 6.3 Performance Tuning and Capacity Planning Guide* for monitoring performance. Compare the operations initiated and operations completed values, as well as the adds requested and adds completed values for a better understanding of what is happening on your system in regard to performance.

If you determine that the server is not responding, run the **IBM Support Assistant Lite** tool. This tool gathers information that you can provide to IBM Software Support to help identify the problem. See "Tools for troubleshooting IBM Tivoli Directory Server" on page 1 for information about the **IBM Support Assistant Lite** tool.

Memory leak suspected

If you suspect that you are experiencing a memory leak, run a script similar to the following one. This script gathers information about the memory sizes of the processes running on your system.

Note: This is an example for AIX. You might need to make modifications for your operating system.

When the script finishes, send the monitor.out text file generated by the script to IBM Software Support for analysis.

The script is as follows:

```
#!/bin/sh
instance=ldapdb2
port=389
binpath=/opt/IBM/ldap/V6.1/bin

while [ true ]; do
  echo | tee -a /tmp/monitor.out
  echo 'Begin Monitoring.....' | tee -a /tmp/monitor.out
  date | tee -a /tmp/monitor.out
  echo 'Process info via ps auwx command: ' | tee -a /tmp/monitor.out
  ps auwx | egrep '(slapd|$instance|PID)' | grep -v grep | tee -a /tmp/monitor.out

 echo 'Memory info via vmstat: ' | tee -a /tmp/monitor.out
  #<VMSTAT command-"#">
  vmstat -t 2 5 | tee -a /tmp/monitor.out

  echo 'Port activity via netstat: ' | tee -a /tmp/monitor.out
  netstat -an | grep $port | tee -a /tmp/monitor.out
  date | tee -a /tmp/monitor.out
```

```
echo 'cn=monitor output follows....' | tee -a /tmp/monitor.out
$binpath/ldapsearch -p $port -s base -b cn=monitor objectclass=* | tee
     -a /tmp/monitor.out 2>&1

date | tee -a /tmp/monitor.out

echo 'Sample LDAP query follow: ' | tee -a /tmp/monitor.out

##
date | tee -a /tmp/monitor.out
echo 'Same query but direct to db2: ' | tee -a /tmp/monitor.out
##
date | tee -a /tmp/monitor.out

sleep 600 #10minutes

done
```

SSL communications returning errors

If you are experiencing errors on SSL, run the following command to verify that SSL is set up correctly.

```
ldapsearch -Z -K <keyfile> -P <keyfilepw>
 -b suffix objectclass=*
```

Where

- *keyfile* is the name of the SSL database file
- *keyfilepw* is the SSL key database password
- *suffix* is the suffix being searched; for example, `-b o=sample`

Record and send any errors to IBM Software Support.

Recovering data from a directory server instance for which the encryption seed value has been lost

If an encryption seed value is lost for a directory server instance during an instance creation, then you can not recover the lost encryption seed value. However, you can recover the data from the directory server instance for which the encryption seed value is lost. The workaround for this is to create a new directory server instance with a new encryption seed value and then use the db2ldif and ldif2db utilities to export and import data. You can supply the new encryption seed and salt value of the new instance to these utilities, and thereby the data would be preserved (along with the passwords) on this new instance. The steps to recover data on a Linux machine are as follows:

1. Create a user for the instance. Issue the command of the following format:

   ```
   idsadduser -u newinst -w newinst -l /home/newinst -g idsldap
   ```

2. Create and configure a new directory server instance. Issue the commands of the following format:

   ```
   idsicrt -I newinst -e thisismynewencryptionseed -l /home/newinst -n
   idscfgdb -I newinst -a newinst -w newinst -t newinst -l /home/newinst -n
   idsdnpw -u cn=root -p root -I newinst
   idscfgsuf -s "o=sample" -I newinst
   ```

   Note: Save the encryption seed "thisismynewencryptionseed".

3. After setting up the new instance, newinst, find and save the salt value generated by the directory server instance. To find the salt value, issue the command of the following format:

```
idsldapsearch -p <port_number> -D cn=root -w root -b "cn=crypto,cn=localhost" \
-s base objectclass=* ibm-slapdCryptoSalt
```

For example, consider the salt value of the new instance, newinst, as "newsaltvalue".

4. To export data to an LDIF file from the directory server instance (for example, oldinst) for which the encryption seed is lost, use the db2ldif command of the following format:

   ```
   db2ldif -o mydata.ldif -I oldinst -k thisismynewencryptionseed -t newsaltvalue
   ```

   Note: After completion of this command successfully, the entire data from the directory server instance, oldinst, would be stored in the mydata.ldif file specified in the db2ldif command.

5. Finally, import the data from the LDIF file to the new directory server instance. Issue the ldif2db command of the following format:

   ```
   ldif2db -i mydata.ldif -I newinst
   ```

Attribute encryption should be avoided in an environment that includes versions of Tivoli Directory Server earlier than V6.1

Attribute encryption should not be used in a Tivoli Directory server environment that include server versions earlier than V6.1. This is because storing encrypted attributes on some servers and not storing encrypted attributes on other severs defeats the purpose of encrypting attributes. In such an environment, for interoperability between servers you should not encrypt attributes. If attribute encryption is used in such an environment the following situations might arise:

- Attempts to replicate schema definitions for encrypted attributes might fail because the target server will not recognize the new IBMAttributeTypes keywords.
- On a server earlier than V6.1 and servers that do not have matching encryption keys, for an attribute that is defined but not encrypted, data are decoded for replication and are stored in decoded format. If the servers have the matching keys, data are not decoded during replication rendering data useless on the servers that do not have matching keys.

In situations where RDBM startup encryption processing fails for a given attribute, the processing can be skipped for the attribute by deleting or commenting the ibm-slapdMigrationInfo line from the configuration file for that entry from the RDBM. For example:

```
dn:  cn=Directory, cn=RDBM Backends, cn=IBM Directory, cn=Schemas, cn=Configuration
...
#ibm-slapdMigrationInfo: encrypt secretAnswer
```

Limitation in using character sets larger than 7-bit ASCII in passwords

Portable characters (common character set) or 7-bit ASCII characters use the first 7 bits to form characters (128 characters, 0 through 127). These are used on most of the code pages. In Tivoli Directory Server, "userpassword" is a binary attribute and it is not converted from the client code page (for example, IBM-437, IBM-850, or Windows 1252) to UTF-8 and then back to the server code page like text attributes. Code pages differ from the portable character limitations. If you use non-portable ASCII characters (beyond the first 127 or 7-bit ASCII) in a user password, then the password will match only if it is always provided from the same code page in which it was originally created.

For example, if you use the Web Administration Tool to create the password, *as12÷÷qw*, for the entry, cn=Bob Garcia,ou=austin,o=sample, and then if you perform a search using the ldapsearch command from the command line as Bob Garcia, the following results are displayed:

```
ldapsearch -D "cn=Bob Garcia,ou=austin,o=sample"\
        -w as12÷÷qw -b "o=sample" "objectclass=*"
ldap_simple_bind: Invalid credentials
```

This occurs because the Web Administration Tool and the command line use different code pages and the password *as12÷÷qw* contains non-portable characters. Therefore, unless you can be sure that the client always authenticates using the same code page that was used when the password was created, you should limit passwords to portable characters (7-bit ASCII). Using non-portable characters is a permanent limitation.

User might experience premature expiry of user password

Comparisons pertaining to password policy, such as validation of maximum age of password (pwdMaxAge), are done in UTC time. However, in geographies that follow daylight saving, a user might experience premature expiry of password or expiry of password later than the due time if not monitoring the timestamps in UTC. If users convert the time in their timezone to UTC and then perform the password expiry calculations, the password would expire at the expected time that was set.

Troubleshooting the limitation in the idssethost command

When configuring a Tivoli Directory Server instance to listen on a specific interface using the idssethost command or GUI, Instance Administration Tool (idsxinst), idssethost configures the directory server instance to listen only on the specific IP address by adding an entry to the ibmslapd.conf file:

```
ibm-slapdIpAddress: <IP_address>
```

However, this leads to unexpected behavior from the perspective of the user because the ibmslapd no longer listens on the loopback address (127.0.0.1). All LDAP client utilities that run locally attempt to connect to ibmslapd over the loopback interface. As a result, when the commands that reside on the local machine are run they fail to contact the directory server, unless the -h option is used to point specifically at the interface that ibmslapd is listening on.

Additionally, the idssethost command does not allow configuring the directory server to listen on the loopback interface. Any attempt to do this will return the following error:

```
idssethost -I ldapdb2 -i 127.0.0.1 -n
GLPCTL062E The specified IP Address '127.0.0.1' is not a valid IP address
for this machine.
```

For the current versions of Tivoli Directory Server, if ibmslapd is required to listen on a specific interface and the loopback interface, the directive to listen on loopback must be added manually. Perform the following steps:

1. Add the IP addresses that you want the server to listen to. Issue the ldapmodify command of the following format:

   ```
   idsldapmodify -p <port> -D cn=<adminDN> -w <adminPW> -f <filename>
   ```

   where, <filename> contains:

```
dn: cn=Configuration
changetype: modify
add: ibm-ibm-slapdIpAddress
ibm-slapdIpAddress: 10.10.10.10
-
add: ibm-ibm-slapdIpAddress
ibm-slapdIpAddress: 127.0.0.1
```

2. Query the DN entry "cn=Configuration" in the ibmslapd.conf file to see the existing IP addresses to which ibmslapd listens to. Issue the ldapsearch command of the following format:

   ```
   idsldapsearch -p <port> -D cn=<adminDN> -w <adminPW> -s sub \
   -b "cn=Configuration" -L objectclass=*
   ```

   An example excerpt of the output of the command is as follows:

   ```
   n: cn=Configuration
   cn: Configuration
   ibm-slapdAdminDN: cn=root
   ibm-slapdAdminGroupEnabled: true
   ibm-slapdAdminPW: {AES256}ohtCABBYFbFo7jREOPz/zQ==
   ibm-slapdCryptoSync: vDydxlFDW0xKtWBL
   ibm-slapdDerefAliases: always
   ibm-slapdIpAddress: 10.10.10.10
   ibm-slapdIpAddress: 127.0.0.1
   ibm-slapdPort: 389
   #ibm-slapdPwEncryption must be one of:
   # none/aes128/aes192/aes256/crypt/sha/ssha/md5/
   # sha224/sha256/sha384/sha512/ssha224/ssha256/ssha384/ssha512
   ibm-slapdPwEncryption: aes256
   ibm-slapdServerBackend: RDBM
   ...
   ```

3. Restart the directory server instance.

If there are no ibm-slapdIpAddress directives the default behavior for ibmslapd is to listen on all available interfaces. Once a specific (or multiple) ibm-slapdIpAddress entries are added to the ibmslapd.conf file, ibmslapd will no longer listen to any interfaces not explicitly listed in the configuration file. To reset a directory server so that it listens on all available interfaces, you can remove all the ibm-slapdIpAddress entries from the ibmslapd.conf and restart the server.

Troubleshooting the environment in which an SNMP agent is configured

Sometimes, you might need to fine tune the environment in which an SNMP agent is configured. IBM Tivoli Directory Integrator is set up to get the desired result when using an SNMP agent for monitoring directory server instances for performance and availability. Some of the likely scenarios and their workaround are listed:

- When the idssnmp tool is running for a long period of time, it is observed that the LDAP_HOME/idstools/snmp/logs/ibmdi.log file grows large in size.

  If a user wants idssnmp to generate or keep less amount of log, the user can modify the TDI_HOME/etc/log4j.properties file and configure an appropriate log appender. For more information on list of appenders, see *IBM Tivoli Directory Integrator version 7.1 Installation and Administrator Guide*.
- To run the idssnmp tool over SSL, user must edit the LDAP_HOME/idstools/snmp/solution.properties file and specify the certificate information.
- If a user wants to run idssnmp over SSL and the solution.properties file is not present, the user can create the solution files required by idssnmp by running the following command:

```
TDI_InstallDirectory/ibmdisrv -s LDAP_HOME/idstools/snmp –v
```

This command will create the solution.properties file in the LDAP_HOME/idstools/snmp directory.

- The idssnmp tool parses log files sequentially. For example, the idssnmp tool will parse slapd.log for the newly generated logs and then proceed parsing next log file, audit.log. This causes the traps to be sent in an order, which is not the same as the messages were generated in the log files. The trap messages contain information about the time with the log line when it was generated and the log file identifier. The user is required to identify the occurrence order of traps based on the timestamp.

Working with a Tivoli Directory Server instance after reinstalling the IBM Tivoli Directory Server

Consider a scenario, where a user creates a directory server instance and then uninstall the IBM Tivoli Directory Server product. And, later at some point of time reinstalls the product.

On Windows systems

The directory server and administration server services are created for a directory server instance when a directory server instance is created, and are removed when the instance is dropped or when the Tivoli Directory Server product is uninstalled.

If a user reinstalls the Tivoli Directory Server product, for the existing directory server instance the services do not get automatically added. To restore the services for the existing directory server instance, enter the following command at the command prompt:

```
ibmslapd –I <instance_name> -i
ibmdiradm –I <instance_name> -i
```

Note: The above commands create the services entry in a Windows operating system and does not create the directory server instance.

To start the services automatically for a directory server instance when the computer is started, open the Services window from Administrative Tools and set the Startup Type to Automatic for the services associated with the directory server instance.

On non-Windows systems

In the case of non-Windows operating systems, such as AIX and Linux, the entry that causes the ibmdiradm daemon to start automatically for a directory server instance gets removed. This entry is not restored for an existing directory server instance when you reinstall the IBM Tivoli Directory Server product.

To restore an entry, compare the inittab file that you saved with the existing /etc/inittab file. Copy any lines in the saved file that are not in the existing file, and add them to the inittab file.

An entry in the /etc/inittab file is of the following format:

```
<unique_id>:<run_states>:<action>:<path_to_script_to_be_run>
```

An example entry:

```
ids0:2345:once:/opt/ibm/ldap/V6.3/sbin/ibmdiradm -I myinst1 > /dev/null 2>&1
```

Working with the tombstone entries on Tivoli Directory Server

In Tivoli Directory Server V6.2 and above versions, a subtree is created that can hold the to-be deleted entries with operational attributes before they are permanently deleted from the RDBM database. The to-be deleted entries are moved to the tombstone subtree, cn=Deleted Objects, and the attribute table is updated for the entry to mark the entry as deleted by adding attribute such as isDeleted. This feature is supported only on directory server's primary RDBM backend. Tombstones are not supported in configuration, schema, or change log backend.

There may be situations for a possible data inconsistency that might get introduced by entry deletions when this feature is enabled, which might require directory administrator's intervention. For performance reasons, no check is provided that can possibly prevent tombstones entries with the attribute isDeleted set to TRUE from being accidentally created or modified under any subtrees.

One of the ways to identify these entries in a RDBM backend database is by comparing the search results returned by a normal search with a search base to that returned by a null base search.

For example, consider a RDBM backend database with two subtrees: o=sample and cn=Deleted Objects. Where, o=sample, contains three entries: cn=A, cn=B, and cn=C (with isDeleted=TRUE). The subtree cn=Deleted Objects containing entries, cn=X, cn=Y, and cn=Z (without isDeleted=TRUE).

When searches using a search base and null base are requested without including the return deleted object control, the following results are displayed.

- In the case of a search with a search base o=sample and search filter, objectclass=*, all entries under the search base including entries with isDeleted=TRUE are displayed.
- In the case of a null base search with search filter, objectclass=*, all entries except for those entries with isDeleted=TRUE are displayed.

*Table 4. The results of different search base with search filter, objectclass=**

Subtree search base	Search filter	With control	Search results	Remarks
o=sample	objectclass=*	No	cn=A cn=B cn=C	cn=C is a normal entry with isDeleted=TRUE
null	objectclass=*	No	cn=A cn=B cn=Z	List LDAP_ENTRY table with isDeleted!=TRUE. cn=C is not qualified.

It is possible that the isDeleted attribute is accidentally deleted or is set to FALSE for entries under the tombstone subtree. When searches using a search base and null base are requested with the return deleted object control, the following results are displayed.

- In the case of a search with a search base, cn=Deleted Objects, and search filter, objectclass=*, all entries under search base are returned. However, when a search with a search base, cn=Deleted Objects, and search filter, isDeleted=TRUE, is requested, entries with isDeleted=FALSE are not returned.

- In the case of a null base search with search filter, objectclass=*, all entries in the database are displayed. However, when a null base search with search filter, isDeleted=TRUE, is requested, only the entries with attribute isDeleted=TRUE in the database are displayed.

Table 5. The results of different search base with different search filters

Subtree search base	Search filter	With control	Search results	Remarks
cn=Deleted Objects	objectclass=*	Yes	cn=X cn=Y cn=Z	cn=Z is a tombstone without isDeleted=TRUE
cn=Deleted Objects	isDeleted=TRUE	Yes	cn=X cn=Y	
null	objectclass=*	Yes	cn=A cn=B cn=C cn=X cn=Y cn=Z	List the LDAP_ENTRY table including those with isDeleted=TURE
null	isDeleted=TRUE	Yes	cn=C cn=X cn=Y	

Note: Deletion of schema attributes may fail due to the fact that some of tombstone entries may still reference them. A delete, rename, or restore of a tombstone entry will not be replicated. This may result in data inconsistencies on replicas, particularly in rename and restore case.

Working with directory server instance backup

Consider a scenario, where a user has taken multiple backups (both offline and online) at multiple locations at different point of time. For example, if myinst1 is the directory server instance and <instance-location>/idsslapd-myinst1/backup1, <instance-location>/idsslapd-myinst1/backup2, and <instance-location>/idsslapd-myinst1/backup3 are the locations where backups are stored at T1, T2, and T3 time (where, T1<T2<T3).

When the instance, myinst1, is dropped or the database instance for the instance is unconfigured, the database configuration details (dbbackuponline and clbackuponline) are set to FALSE in the <instance-location>/idsslapd-myinst1/backup3/dbback.dat file.

Scenario1

If the instance, myinst1, is recreated and configured with the backup location set to <instance-location>/idsslapd-myinst1/backup2, where the offline backup of the previous instance was stored before it was dropped. In this case, the results of the monitor search, "cn=backup,cn=monitor", will be consistent for searches performed before starting the directory server instance (server state as stopped) and after starting the directory server instance (server state as running).

Scenario 2

If the instance, myinst1, is recreated and configured with the backup location set to <instance-location>/idsslapd-myinst1/backup1, where the online backup of the previous instance was stored before it was dropped. In this case, the results of the monitor search, "cn=backup,cn=monitor", will be not be consistent for searches performed before starting the

directory server instance (server state as stopped) and after starting the directory server instance (server state as running).

The reason is when the server is in stopped state monitor search refers <instance-location>/idsslapd-myinst1/backup1/dbback.dat file for the online status of database and change log (which is true since the previous online backup was stored at <instance-location>/idsslapd-myinst1/backup1), and when the server is in running state monitor search refers directory server for the online status of database and change log (which is false as database is not configured for online backup). If a user starts the directory server and then performs online backup based on the monitor search results that the user received when the directory server was in stopped state, user would get unexpected behavior because the database is not configured for online backup.

If a user's intention is to perform restore from an existing backup files for a recreated directory server instance, it is alright to configure to a previous backup location. However, if user's intention is to back up the recreated directory server instance, then to avoid the situation as mentioned in scenario 2, it is advisable to remove previous backup files from the location or specify a location that does not have any backup image.

Configuring preaudit records for serviceability

Sometimes, when Tivoli Directory Server locks up or stops, the audit log might not have a record of the operation that has caused the problem. This is because the audit logs are updated after the directory server backend completes the operation. So, any problems that occur before the audit records get updated are not logged and the result of the operation is unknown.

You can configure auditing of preaudit records that is, recording of operations that have not completed. When preaudit records are enabled, the audit plug-in is called to update audit records before the operation completes. When preaudit is enabled, the thread ID is also logged in the audit header. To enable pre-auditing, you must set the value of the IBMSLAPD_PREOP_AUDIT environment variable to "YES". This can be done by accessing the environment variable or by using the ldapmodify command with the following format:

```
ldapmodify -D <adminDN> -w <adminPW>
dn: cn=Front End, cn=configuration
changetype: modify
add: ibm-slapdSetEnv
ibm-slapdSetEnv: IBMSLAPD_PREOP_AUDIT=YES
```

An example of a pair of diagnostic audit records when preaudit is enabled, where the sequence identifier is 3: <"PREOP: 3" and "POSTOP: 3">, is as follows:

```
AuditV3--2007-08-29-11:44:32.912-06:00DST--V3 PREOP: 3 threadId:  1161116592
Search--bindDN: cn=root--client: 127.0.0.1:1044--connectionID:
3--received: 2007-08-29-11:44:32.912-06:00DST--Success
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: o=sample
scope: baseObject
derefAliases: neverDerefAliases
typesOnly: false
filter: (objectclass=*)

AuditV3--2007-08-29-11:44:33.092-06:00DST--V3 POSTOP: 3 threadId:  1161116592
Search--bindDN: cn=root--client: 127.0.0.1:1044--connectionID:
3--received: 2007-08-29-11:44:32.912-06:00DST--Success
```

```
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: o=sample
scope: baseObject
derefAliases: neverDerefAliases
typesOnly: false
filter: (objectclass=*)
```

No entries are displayed to root and anonymous users when logged on to IBM Tivoli Directory Proxy Server using the Web administration tool

Even with the ibm-slapdAllowAnon attribute under the DN entry "cn=Connection Management, cn=Front End, cn=Configuration" set to true, the root and anonymous users are not able to view entries using the Web administration tool when logged on to Tivoli Directory Proxy Server. This is because the Web administration tool uses page control to browse through entries. If paging is enabled only for administrators by setting the ibm-slapdPagedResAllowNonAdmin attribute to false under the DN entry "cn=ProxyDB, cn=Proxy Backends, cn=IBM Directory, cn=Schemas, cn=Configuration", then non administrators will not be allowed to browse through entries using the Web administration tool. This restriction is also applicable in the case of RDBM servers.

For the root and anonymous users to view the entries using the Web administration tool, the following must be considered:

- Set the ibm-slapdAllowAnon attribute under the DN entry "cn=Connection Management, cn=Front End, cn=Configuration" to true.
- Set the ibm-slapdPagedResAllowNonAdmin attribute under the DN entry "cn=ProxyDB, cn=Proxy Backends, cn=IBM Directory, cn=Schemas, cn=Configuration" to true.

After setting the attributes, restart the directory server instance.

Note: The root user is an anonymous user in the case of proxy server for DIT-related operations.

When performing the restore operation on a directory server instance, the directory server instance gets restored to latest consistent state and not to the point when the backup was performed

Observed

1. Create and configure a directory server instance for online backup.
2. Stop the directory server instance and perform the initial offline backup either using the Web administration tool or the **idsdbback -u -k** command.
3. Add the suffix, `o=sample`, and start the directory server instance.
4. Add the entry `o=sample`.
5. Verify that the database parameter `LOGARCHMETH1` is correctly set.
6. Perform a restore operation either using the Web administration tool or the **idsdbrestore -k** command.
7. Verify that the suffix, `o=sample`, is not present (since we have backed up before adding the suffix).
8. Add the suffix, `o=sample`, and start the directory server.

9. Perform **ldapsearch** for the entry `o=sample` and you will observer that the entry is present.

Expected result

The entry, `o=sample`, should not be present because we have added only the suffix, `o=sample`, after restore on a clean database (no data).

Reason

During roll-forward, DB2 scans the current logs in the `newlogpath`location. Because of the options specified in the roll-forward, DB2 scans the logs until the end, and restores a database to the latest consistent state.

For example, at the time of backup, if you have 100 entries in the directory server and after the backup operation if you delete 5 entries and then perform the restore operation, you might still find the 95 entries in the directory rather than the 100 entries that you have backed up. This is because, the latest consistent state of the database was after the deletion on 5 entries.

However, it is possible to modify the options in the roll-forward recovery operation such that the database is restored to the point where it had 100 entries. For this, we need to specify the timestamp of the last committed change, which is the timestamp at which the 100th entry was added and to obtain this value of timestamp is difficult.

Online backup and restore limitation

When performing online backup and restore, it is observed that if the folder name (backup location) to which online backup was initially configured is changed for the subsequent backups, no error will be thrown during the backup operation (idsdbback) but during the restore operation (idsdbrestore) the following error messages might be displayed:

```
...
GLPCTL101I Restoring backup database rdsdb to configured database rdsdb.
GLPCTL103E Failed to restore backup database rdsdb to configured database rdsdb.
GLPDBR004E Failed to restore directory server instance 'tdsadmin'.
GLPDBR028W The program did not complete successfully. View earlier error messages
 for information about the exact error.
...
```

Reason

As per the idsdbback and idsdbrestore (also available as dbback and dbrestore) design for online backup, the first-time backup must be a complete offline backup while the ibmslapd process is in stopped state. After the first offline backup, the online backup feature can be used while the ibmslapd process is running.

During the first offline backup, the idsdbback command takes the following options:

```
idsdbback -I <instance_name> -k /path/backupfolder1 -u [-a /path/logarchivefolder]
```

If the optional path for logarchive folder is not provided, the command will use a folder inside the *backupfolder1* folder (as per the example) to configure the *logarchivefolder* and sets this value in the corresponding db2 database's configuration parameter "LOGARCHMETH1".

If the backup folder is changed for a subsequent online backup, idsdbrestore will fail if the previous backup folder does not exist, since the LOGARCHMETH1 still points to the previously configured value.

Confirming the problem

To confirm, verify the LOGARCHMETH1 variable for the corresponding database's configuration.

```
su - <instance_name>
db2 list db directory
db2 get db configuration for <databasename> | grep -i LOGARCHMETH1
```

Note: Replace the <instance_name> and <databasename> with the appropriate names.

Resolving the problem

If there is a need to change the backup location after the first off line backup (or even after subsequent online backups) follow the procedure below to update the backup folder and logarchive folder values:

1. Stop the ibmslapd process.

   ```
   ibmslapd -I <instance_name> -k
   ```

2. Use the idsbackup command to update both the backup folder and logarchive folder.

   ```
   idsdbback -I <instance_name> -k /path/backupfolder2 \
       -a /path/backupfolder2/INACTIVE_LOGS -u -n
   ```

3. Start the ibmslapd process.

   ```
   ibmslapd -I <instance_name> -n -t
   ```

When user attempts to stop the log management service after starting the service using the Web administration tool, it fails to stop

Observed

The log management service gets started when a user uses the Start/Stop log management panel of the Web administration tool to perform this operation. However, when the user attempts to stop the log management service using the Start/Stop log management panel, the panel displays the service has been stopped but it has been observed that the log management service is running in the background.

Expected behavior

The expected result is that the log management service should have stopped.

Reason

For the log management to work, Tivoli Directory Integrator is required. It has been observed that when the IBM Tivoli Directory Integrator v7.1 GUI installer prompts for a location for the TDI Solutions Directory, and if the user opts for the default option, which is the preselected option, "Use a subdirectory named TDI under my home directory" the problem with log management service is observed.

Workaround

If the user opts for the Use Install Directory option from the Tivoli Directory Integrator v7.1 GUI installer for the Solutions Directory, the mentioned problem with log management service does not occur. For example, if the Tivoli Directory Integrator is installed in the location */opt/IBM/ldap/V7.1/TDI*, and the user opts to provide the TDI Solutions Directory within the */opt/IBM/ldap/V7.1/TDI* directory, then the problem with log management service is not observed.

IBM Tivoli Directory Server instance does not start and returns error GLPCRY007E

Scenario

1. Create a Tivoli Directory Server instance, *inst1*, configure the instance, and start the instance. The encryption seed used to create the instance, *inst1*, is *thisismyseed*.
2. Drop the instance, *inst1*, without dropping the database associated with it.
3. Recreate the instance with the encryption seed, *thisismyseed*, and configure the instance with the existing database.
4. Start the instance.

Observed

The instance does not start and returns error:

```
GLPCRY007E The directory key stash file is inconsistent with the
 associated encrypted data.
```

Reason

When a directory server instance is created and is started, some information from keystash file (.ksf), is stored in the database. Therefore, an existing database cannot be used with a keystash file that gets created when a instance is recreated.

Workaround

In such case, if a user intends to use an existing database with a new instance, then at the time of instance creation the user must use **-e** and **-g** options to specify the encryption seed and encryption salt values for the new instance. This encryption seed and salt value must be same as that of the dropped instance.

If the user has not provided the salt value with the **-g** option for the instance that the user is intending to drop, then the salt value must be determined before dropping an instance. Issue the **idsldapsearch** command of the following format to retrieve the salt value.

```
idsldapsearch -h <IP address> -p <port> -s base -b "cn=crypto,cn=localhost" \
   objectclass=* ibm-slapdCryptoSalt
```

Chapter 12. Interoperability

This chapter contains information on interoperability between Tivoli Directory Server and other directory servers.

Interoperability with Novell eDirectory Server

When performing simple bind using Tivoli Directory Server client utilities against Novell eDirectory Server, error message such as "ldap_bind: Confidentiality required" might get displayed

If you get error message such as "ldap_bind: Confidentiality required" when performing simple bind using Tivoli Directory Server client utilities against Novell eDirectory Server, you must run the following command:

```
#ldapconfig set "Require TLS for Simple Binds with Password=no"
```

Interoperability with Microsoft Active Directory

Making Tivoli Directory Server configured over SSL using serverClientAuth authentication to work with Microsoft Active Directory client LDP.exe

To make Tivoli Directory Server configured over SSL using serverClientAuth authentication to work with Microsoft Active Directory client LDP.exe perform the following steps.

1. Select Internet Information Services (IIS) Manager from Administrative Tools in Control Panel.
2. On the left navigation panel, select the Web Site node.
3. Under the Web Site node, right-click Default Web Site, and then select Properties.
4. On the Default Web Site Properties dialog box, select the Directory Security tab.
5. To request for a new certificate, click the Server Certificate button under the Secure communications area. This opens Web Server Certificate Wizard.
   a. On the Server Certificate page in the IIS Certificate Wizard dialog box, select the Create a new certificate option and click Next.
   b. On the Delayed or Immediate Request page, enter the required options and click Next.
   c. On the Name and Security Settings page, in the Name field enter the host name of the machine and click Next.
   d. On the Organization Information page, specify appropriate names and click Next.
   e. On the Your Site's Common Name page, in the Common name field, enter the host name of the machine and click Next.
   f. On the Geographical Information page, specify appropriate values and click Next.

g. On the Certificate Request File Name page, in the File name field specify the path name and file name for the certificate request and click Next.
h. The summary of the values provide is displayed. Click next.
i. Click Finish.

6. Send the certificate request using the steps mentioned above to any Certificate Authority (CA) to issue a certificate.
7. After receiving the server certificate add the certificate using IIS Certificate Wizard.
   a. On the Pending Certificate Request page, select the Process the pending request and install the certificate option and click Next.
   b. On the Process a Pending Request page, in the Path and file name field specify the path name and file name of the certificate. You can also use Browse to select the certificate. Click Next.
8. Export the personal certificate to pfx or p12 format using IIS Certificate Wizard.
   a. On the Modify the Current Certificate Assignment page, select the Export the current certificate to a .pfx file option and click Next.
   b. On the Export Certificate page, in the Path and file name field enter the path name and file name where pfx certificate to be stored. Click Next.
   c. On the Certificate Password page, in the Password and Confirm password fields enter the password and click Next.
   d. On the Export Certificate Summary page, the summary of the provided values are displayed. Click Next.
   e. Click Finish.
9. To import the certificate, double-click the stored pfx certificate. This opens Certificate Import Wizard.
   a. On the File to Import page, in the File name field enter the path and file name of the pfx certificate and click Next.
   b. On the Password page, enter the password and click Next.
   c. On the Certificate Store page, select the Place all certificate in the following store option and the click Browse and select Personal from the Select the certificate store you want to use list in the Select Certificate Store dialog box. Click Next.
   d. Click Finish.
10. To export the personal certificate in BER format, perform the following steps.
    a. Open Internet Explorer, select Internet Options from the Tools menu, select the Content tag in the Internet Options dialog box, and select Certificates under the Certificates area.
    b. On the Personal tab in the Certificates dialog box, select the certificate and click Export. This opens Certificate Export Wizard.
    c. On the Export File Format page, select the Base-64 encoded X.509 (.CER) option and click Next.
    d. On the File to Export page, in the File name field enter the file name you want to export and click Next.
    e. Click Finish .
11. On a machine on which a Tivoli Directory Server instance is running, open the Tivoli Directory Server's key database file using GSKit's key management application, ikeyman
12. Add the exported certificate as a signer in the server key database.

Chapter 13. Known limitations and general troubleshooting

This chapter contains miscellaneous problem determination information.

Known limitations

The following sections describe known limitations in IBM Tivoli Directory Server 6.1 and later versions of directory servers.

Command line utilities allow an option to be entered more than once

You can run a command that specifies an option more than once. If an option is specified more than once, the option entered last is used. For example, if you enter the following command, the `-I inst1` option is ignored and the `-I inst2` option is used.

```
idsdnpw -p root -n -I inst1 -I inst2
```

Some types of invalid data entered on command line utilities do not produce an error

If you enter a command that contains invalid data after all required options have been specified, you will not receive an error message. For example, the following command contains the required options for the **idsdnpw** command, but the '--' characters following the required option are invalid.

```
idsdnpw -p root -n -I inst1 --
```

Even though the '-' characters are invalid, no error is returned.

No locking mechanism for conflicting commands on the same directory instance

No locking mechanism exists at this time to prevent conflicting commands from running at the same time for the same directory instance. For example, you can run a command to configure a database and drop the database at the same time.

Unable to drop database

On Windows systems, if all of the following are true, you might not be able to drop the database immediately after you stop a directory server instance.

- The directory server instance is started from the console and not as a service.
- You stop the directory server instance by using the **ibmslapd -k** command.
- You try to drop the database immediately after stopping the directory server instance with the **ibmslapd -k** command.

The Instance Creation Tool and the **idsidrop** and **idsucfgdb** commands are able to unconfigure the database but fail to drop it if all the listed conditions are satisfied. If you encounter this problem, you can manually delete the database directory after running the **idsidrop** or **idsucfgdb** commands. Alternatively, wait at least two minutes after stopping the server, and then drop the database.

Partial replication

Partial replication is a replication feature that replicates only the specified entries and a subset of attributes for the specified entries within a subtree. The entries and attributes that are to be replicated are specified by the LDAP administrator. Using partial replication, an administrator can enhance the replication bandwidth depending on the deployment requirements. For instance, an administrator may choose the entries of the object class person with cn, sn, and userPassword attributes to be replicated and description attribute not to be replicated.

There are situations when administrator's intervention is required for the smooth running of partial replication. These scenarios are listed.

Creating missing parent entries on the consumer

In filtered replication, an entry addition might fail displaying "No such object' error because the parent entry does not exist on the consumer. This happens because the parent entry did not match the filter and was not replicated. In such cases, if the ibm-replicationCreateMissingEntries attribute is set to TRUE, the supplier should detect this error case and then generate and submit an add request for the missing entry before retrying the add operation instead of processing this case as an error. The missing entry should have the same DN as that of the immediate parent of the entry whose add failed. The missing entry belong to the objectclass extensibleObject and will contain operational attributes for create and modify timestamps as present on master server, that is, the timestamps will not be modified when the entry is created on consumer. The missing entry should have ACL's as on the supplier server and should also have the description attribute value set to "Missing entry created by <master server>".

Scenario

Sometimes the method to generate and submit a request to add a missing entry will be recursive and the end condition would be either a successful add of all missing ancestors in the chain or a failure might occur while adding any of the missing ancestors (for any reason other than NO_SUCH_OBJECT). In case of a failure, the change cannot be replicated and administrator intervention will be required.

Workaround

The administrator should manually take care of handling errors when the ibm-replicationCreateMissingEntries attribute is set to FALSE. Administrators can also use error logs to identify the replication failure error messages that are logged into error logs.

Modification in replication filter

Scenario

In partial replication, changes to replication filter can be dynamic. When a replication filter is changed, the data on the consumer would be in sync with the supplier cannot be assured.

Workaround

In cases where replication filter is changed, the administrator should take of such changes and reinitialize the consumer as per the new replication filter.

Note: The replication filter entry cannot be deleted if it is in use.

Replication is not initiated if the password encryption settings of a supplier are not supported by the consumer

In a replication environment, if a supplier is using a password encryption setting that is not supported by the consumer, then replication will not be initiated. Also, the supplier logs a message and sets the replication state to “error xxxx” where xxxx is the id of the message that describes the problem.

Migrating from IBM Tivoli Directory Server V6.0 or later version of directory server to V6.3

When migrating from IBM Tivoli Directory Server V6.0 or later version of directory server to V6.3, the existing instances can be migrated using the **idsimigr** command line utility. This tool retrieves the schema and configuration files of the instances from the standard location specific to the instances. While migrating from Tivoli Directory Server V6.0 or later version of directory server to V6.3, certain checks need to be performed, otherwise, the tool might exit displaying error messages.

- If an instance already exists, the backup directory should not be specified. If the backup directory is specified, the tool will exit displaying appropriate error messages.
- If the Tivoli Directory Server instance earlier that V6.3 that needs to be migrated has been dropped before running **idsimigr**, the backup directory should be specified. In such a scenario, the encryption key is not required but if the encryption key is specified, the tool will exit displaying appropriate error messages.
- During migration, the Windows service entry for each directory server and the Directory Administration server are migrated to Tivoli Directory Server 6.3. In such scenarios, to avoid any unforeseen errors, it is required that the user take backup of the schema, configuration, and key stash files before migration, even if, the user has not dropped the instances.

In Tivoli Directory Server V6.1 and later versions, alias dereferencing might not work when persistent search is run on a server with no alias entries

If persistent searches are run before any alias entries are added to the server, then persistent searches will not dereference aliases. That means, only if alias entries exist on the server before running persistent searches, the dereferenced aliases will be displayed.

When both proxy and back-end servers are configured to use PKCS#11 mode and need to communicate with a remote nCipher crypto hardware for SSL operation, the operation times out

In order to increase the operation timeout duration, you need to increase the number of retries that a proxy server should attempt to establish a connection. This is because

```
the total time for which a proxy server waits to establish a connection =
maximum time for which proxy waits to establish connection *
number of retries by a proxy server to establish a connection
```

To increase the number of retries, export the environment variable, SERVER_ATTEMPT_TIME, with the required retry count. Set the retry count to greater than 12, if the crypto hardware used for SSL operation is at a remote location.

Tivoli Directory Server V6.2 instance stops when nCipher crypto hardware client is restarted

Scenario

The below mentioned steps describe the situation in which a Tivoli Directory Server instance might stop.

1. Start a Tivoli Directory Server V6.2 instance configured over SSL with server client auth to use PKCS#11 in keystorage and accelerator mode.
2. Perform search operation using an LDAP client in SSL mode.
3. Restart the crypto hardware used.
4. Perform search operation using an LDAP client in SSL mode.

Reason

Users must not restart the crypto hardware if a Tivoli Directory Server V6.2 instance is configured to use PKCS#11 in keystorage or accelerator mode. If crypto hardware is reset that is used by a Tivoli Directory Server instance for cryptographic operations, then the instance will stop logging appropriate messages in trace file.

Querying an entry of large size using the idsldapdiff tool might throw an exception

The Java implementation of the idsldapdiff tool has limitation because of which it is unable to handle entries on Tivoli Directory Server that have more than 50 MB size. As a result of this, the tool might throw an Out of Memory exception when dealing with entries with more than 50 MB size.

The idsadsrun utility might fail when synchronizing a large number of entries with size-limit or time-limit like exception

To avoid exception like size-limit or time-limit, you need to consider the following:

1. Before performing synchronization, configure Microsoft Active Directory setting parameters to the following:

```
MaxPoolThreads              4
MaxDatagramRecv             4096
MaxReceiveBuffer            10485760
InitRecvTimeout             120
MaxConnections              5000
MaxConnIdleTime             900
MaxPageSize                 1000000
MaxQueryDuration            1000
MaxTempTableSize            10000
MaxResultSetSize            262144
MaxNotificationPerConn      5
MaxValRange                 1500
```

2. In the ibmdisrv file, tune the JVM parameters, for example, for a machine with 1 GB RAM, the parameter values can be Xms254m-Xmx1024m. You can tune the parameters based on your machine configurations. For best results, use a machine with high-end configurations to run the Active Directory synchronization tool.
3. Also, synchronizing approximately 100000 entries using the “Run full synchronization of the entries from Active Directory Server to IBM Tivoli

Directory Server followed by real time synchronization" mode while running the idsassrun tool gives best results. With the "Run real-time synchronization" mode, up to 400000 entries can be synchronized.

The idsadsrun utility fails if a Tivoli Directory Server instance is run on a different port using the –p option

Presently, the Active Directory synchronization tool detects the Tivoli Directory Server admin DN, password, LDAP URL, and port number from the instance name. Therefore, when a Tivoli Directory Server instance is run on a different port using the –p option, the tool is unable to detect the port number specified using the –p option.

Operations error is displayed when null base search is performed against a proxy server

IBM Tivoli Directory Proxy Server does not support null base search and gives an operations error if null base search is fired against it.

When installing using InstallShield GUI, a change in disk space on the system does not get refreshed on the tool

When using InstallShield GUI, the tool does not refresh the information when there is a change in disk space on the system. If user modifies disk space allocation on a system, the changed information does not get reflected on the tool. In order to use the changed disk space allocation, user has to cancel the current installation and start a fresh using InstallShield GUI.

When the pwdLockout attribute is set to true, user account might get locked even if the number of invalid bind attempts is less than the pwdMaxFailure value

A user account might get locked when all the invalid bind attempts are made within a given time interval that is set in the pwdFailureCountInterval attribute. For example, consider the following attributes are set to:

```
ibm-pwdPolicyStartTime=20070217044605Z
pwdInHistory=0
pwdCheckSyntax=1
pwdGraceLoginLimit=0
pwdLockoutDuration=0
pwdMaxFailure=3
pwdFailureCountInterval=0
passwordMaxRepeatedChars=0
pwdMaxAge=99
pwdMinAge=0
pwdExpireWarning=0
pwdMinLength=5
passwordMinAlphaChars=0
passwordMinOtherChars=0
passwordMinDiffChars=0
ibm-pwdPolicy=true
pwdLockout=false
pwdAllowUserChange=true
pwdMustChange=false
pwdSafeModify=false
ibm-pwdGroupAndIndividualEnabled=true
```

With this setting, if a user makes three invalid bind attempts, the user can still continue with bind attempts because the pwdLockout attribute is set to false.

However, pwdFailureTime is registered even when pwdLockout is false therefore if user has done three invalid bind attempts with pwdLockout=false, pwdFailureTime will have timestamps of the consecutive authentication failures.

Set the pwdLockout attribute to true:

```
# idsldapmodify -D <cn=RDN_value> -w <password>
                -p <port_number> -h <host_name>
dn:cn=pwdpolicy,cn=ibmpolicies
pwdLockout:true
```

Now, when the pwdLockout attribute is set to true another invalid or valid bind attempt will cause lockout of user account. This is because the invalid bind attempts made when "pwdLockout=false" is also taken into account depending on the number of values in the pwdFailureTime attribute that are younger than pwdFailureCountInterval.

Description attribute for groups is not syncing from Active Directory to Tivoli Directory Server

Active Directory synchronization solution only synchronizes the user entry attributes provided with TDSOptionalAttributes in the adsync_public.prop file.

When configuring Tivoli Directory Server over SSL to use PKCS#11 SYMMETRIC acceleration support, there are chances for memory leak

On configuring Tivoli Directory Server over SSL to use PKCS#11 SYMMETRIC acceleration support for performing cryptographic operations using nFast crypto library, memory leak is observed during operations.

Note: nFast cryptographic library is a third party library. It is used for PKCS#11 support provided by Tivoli Directory Server.

Importing LDIF files containing SHA-2 encrypted password or encrypted attributes to versions earlier than 6.3

The db2ldif utility exports user password data and other encrypted data as it is stored in the directory. This indicates that if the data is encrypted, then it is exported in the same format to the LDIF file.

The LDIF import utilities, ldif2db and bulkload, check the format of the data in the LDIF file to verify whether it is in a recognizable encryption format. If any data that is not recognized from the supported encryption method is assumed to be in clear text format, and is encrypted or not encrypted based on the configured value of the ibm- slapdPwdEncryption configuration attribute. Therefore, when the earlier versions of Tivoli Directory Server import Tivoli Directory Server V6.3 LDIF file containing data encrypted using the SHA-2 family of encryption scheme, the earlier version of servers assume the data is in clear text since SHA-2 family of encryption scheme is unknown encryption format, and will encrypt the data depending on the configured value of the ibm-slapdPwdEncryption configuration attribute.

Multivalued attributes in a virtual list view search

Explanation

In virtual list view searches, the search filter resolutions are done in database. This is because in virtual list view searches, the entire result set are not read from the database at one time. Whereas, in a normal search operation a list of EIDs is maintained in the memory to ensure that duplicate entries are not returned to clients, even if DB2 returns duplicate EIDs.

Since in virtual list view searches, the entire result set (list of EIDs) is not read into the memory, the constraint of identifying and preventing duplicates exist. Therefore, if a virtual list view search is performed with a primary sort key attribute having multiple values, then the entries returned might not be in sorted order. Additionally, duplicate entries might also be returned.

Example

Consider a directory server with the following data set:

Table 6. Entries and multivalued attribute of the entries

EID	Values of the cn attribute
1	A, Y
2	C, J
3	E

In a normal search with "cn" as the sort key, the entries will be returned in the following order: 1, 2, 3. However, the search filter resolution for the DB2 query will return EIDs in the following order: 1, 2, 3, 2, 1, based on the values of "cn". In this case, the duplication is prevented by maintaining the list of EIDs at the server end.

In the case of a virtual list view search, the entire result set is not maintained in memory and therefore preventing duplication is not possible.

Consider a virtual list view search sent with the following values: before count = 1, after count = 1, offset = 3, and content count = 0. If the Virtual list view control is applied over the result set of DB2, the entries returned will be 2, 3, 2, where the entry with EID=2 is returned twice. The result shows that there is a possibility of returning duplicate entries in a virtual list view search if the sort key is a multivalued attribute.

In a distributed directory environment, only base scope search is supported with ibm-allMembers

If distributed group and dynamic distributed group are enabled in the configuration file, then only base scope search is supported with ibm-allMembers. If onelevel or subtree scope search is attempted with ibm-allMembers, then an appropriate error message is logged in the ibmslapd.log file and LDAP_UNWILLING_TO_PERFORM is returned.

However, if distributed group and dynamic distributed group are disabled in the configuration file, then a search for ibm-allMembers is forwarded to a single backend server. In this case, the search returns group members for all search scopes.

A Tivoli Directory Server instance might fail to start if the system date is modified

If significant change is made to the system date, that is, set to a previous date (for example, one month), from the date when the Tivoli Directory Server instance was configured on the system, then the directory server instance might fail to start and would give the following error messages:

```
GLPRDB001E Error code -1 from function:" SQLTables "  .
GLPRDB001E Error code -1 from function:" SQLFetch "  .
GLPRDB001E Error code -1 from function:" SQLFetch "  .
GLPRDB001E Error code -1 from function:" SQLFetch "  .
GLPSRV064E Failed to initialize be_config.
```

In this scenario, when the directory server instance is run with server trace set to ON and debug level set, the following error messages can be seen in the server trace:

```
188:22:35:24 T1 retrieving SQLGetDiagRec info
188:22:35:24 T1 Error - map_rc_fnc: henv=0,hdbc=0,hstmt=10001,native
retcode = -443; state = "38553"; message = "[IBM][CLI Driver][DB2/SUN64]
SQL0443N Routine "SYSIBM.SQLTABLES" (specific name "TABLES") has returned
an error SQLSTATE with diagnostic text "SYSIBM:CLI:-727". SQLSTATE=38553"
```

The above error messages can also be seen in the db2diag.log file.

In the configuration file, the format of the DN gets changed when a composite DN is added as suffix

If a composite DN is added as suffix, the format of DN that gets added to the configuration file is different from the DN value that was provided. For example, a composite DN, *o=sample+c=in* gets updated in configuration files as *c=i\20 + o=sample*.

The idsdbmaint tool throws Unable to estimate the database size error message

When the idsdbmaint tool is run with root or administrator privileges the tool inherits those privileges and therefore, the tool will be able to access a directory even if it does not have write permissions or sufficient privileges for the directory instance owner. The idsdbmaint tool attempts to estimate the directory size with instance owner's privileges. In this case, if the instance owner does not have sufficient privileges to run the operation, the tool will throw the following error.

```
GLPDBA054E Unable to estimate the database size.
```

An error message Error opening filename.cat gets displayed when running Tivoli Directory Server

If Tivoli Directory Server is set to a locale that does not have corresponding message files for that locale, then an error message "Error opening filename.cat" will be displayed along with an appropriate message in English locale.

The reason for this error can be the following:

- An incorrect language pack is installed on the system
- Tivoli Directory Server does not support that particular locale

The values "TRUE" and "FALSE" are not translated

In the translated versions of IBM Tivoli Directory Server, the GUI tools, such as the Web administration tool, and the directory server messages will not translate the values "TRUE" and "FALSE" to the corresponding locales of the translated version.

In the Web administration tool some schema related keywords are not translated

In the Web administration tool for the translated versions of IBM Tivoli Directory Server, values of some schema related keywords such as syntax and matching rules are not translated.

The date field is not getting displayed properly for the Russian locale in the Web administration tool

For the translated version of the Web administration tool in the Russian locale, sometimes the date format either gets displayed in wrong format or the last character of the month name gets truncated. This is a limitation with the tool.

The date and time values are displayed in the English locale on certain panels in the translated versions of the Web administration tool

On certain panels, such as Manage backup/restore, in the translated versions of the Web administration tool, the date and time values displayed on the panels are in the English locale instead of the locale of the translated version.

The Error logo is not displayed with error messages in the Web administration tool

If you access panels on the Web administration tool when the directory server is in the stopped state, an Error panel is displayed with error messages. However, on this Error panel, the Error logo is not displayed. This is a limitation with the Web administration tool.

Mnemonics missing from the panels of Instance Administration Tool and Configuration Tool

On the panels of Instance Administration Tool and Configuration Tool, the mnemonics for the buttons, such as Help, Finish, and Cancel, might not be available. This limitation in the tools are specific to the French and Korean translated versions.

In the case of the French and Korean translated versions of the tools, you can use the following keys for the buttons:

- Help - The mnemonic key is "H".
- Finish - The mnemonic key is "F".
- Cancel - The mnemonic key is "C".

Alternatively, you can use Hot keys (a combination of keys) to access the following buttons:

- Alt + H - Help
- Alt + F - Finish
- Alt + C - Cancel

Able to encrypt an attribute that is present in the RDN of an entry

When adding an entry to a directory server instance with an RDN that has encrypted attribute in it, the following error message is displayed:

```
GLPWDM003E An error occurred while adding entry uid=5user,uid=5user,o=ibm,c=us :
uid=5user,uid=5user,o=ibm,c=us: [LDAP: error code 34 - GLPSRV156I
Encrypted attributes are not allowed in entry distinguished names.].
```

An attribute that is already present in the RDN of an entry can be encrypted without getting any error message. However, on restarting the directory server instance, the entry displays the RDN in clear text rather than displaying the encrypted attribute of the RDN in the encrypted format.

This inconsistency is a limitation in the existing design.

LDAP search filters exceeding 4K are not supported

If an LDAP search filter exceeds the 4K limit, then the server might throw an `ldap_search:bad search filter` error and an error message might also be logged in the db2cli.log file indicating a syntax error in the query sent to DB2. For example, the error message logged in the db2cli.log can be of the following format:

```
12/04/07 10:38:24 native retcode = -104; state = "42601"; message =
"[IBM][CLI Driver][DB2/6000] SQL0104N  An unexpected token
"END-OF-STATEMENT" was found following ".ORGANIZATIONDN WHER". Expected
tokens may include:  ")".  SQLSTATE=42601
```

Also, the ibmslapd.log file might contain the following error:

```
12/04/07 10:37:44 AM GLPRDB001E Error code -1 from function:" SQLExecute "  .
```

To avoid these errors, you must use search filters that do not exceed the 4K limit.

If the DB2 versions on source and target server are different, the idsideploy tool displays error when creating a directory server instance from an existing directory server instance

During the creation of a directory server instance from an existing directory server instance, the idsideploy tool takes online backup of the source database including the logs. At the target server, the database is restored with rollforward of logs to bring the database to a consistent state. However, there is a limitation when the target database is DB2 v9.5 and the source database is a previous version, DB2 v9.1. This is because the rolling forward of logs from a previous level to DB2 v9.5 is not supported.

Therefore, when using the idsideploy tool, you must use the same DB2 versions on the source and target server.

To know more about supported platforms for DB2 backup and restore operations, see Backup and restore operations between different operating systems and hardware platforms.

The idsideploy tool might fail to restore a database if the backup location has backup images of the database

When the idsideploy tool is run to create a copy of a directory server instance with data of an existing directory server instance, you must ensure that the directory path specified with the -L option does not already have backup image of the

database, which the tool is attempting to restore. If a backup image is already present, then the restore operation of idsideploy will fail.

The idsdbback command might fail to create an online backup image of a directory server instance created by the idsideploy tool

When the idsideploy tool is used to create a copy of a directory server instance (along with database), the tool backs up the source database and restores it on the target server. During this, all the internal database settings are also copied on to the target server as it is. The error messages that might get displayed are:

```
GLPDBB051E Failed to create path '/export/home/mybkup/back/INACTIVE_LOGS'
for logging inactive log files.
GLPDBB010E Failed to back up directory server instance 'inst2'.
```

One of the reasons because of which the error might have occurred is that the target database is using the same settings as that of the source database. If a user want to perform online backup for the target server instance, then the user must set the archive path for the target server instance before performing the online backup operation, otherwise online backup might fail. To know more about the idsdbback and idsideploy commands, see *IBM Tivoli Directory Server Version 6.3 Command Reference*.

Unable to connect from an OpenLDAP client over DIGEST-MD5 to Tivoli Directory Server

Tivoli Directory Server V6.2 fails to authenticate an OpenLDAP client that attempts to connect using the DIGEST-MD5 SASL mechanism if the version of OpenLDAP client used is 2.4.11. However with Tivoli Directory Server V6.2, you can use OpenLDAP clients version 2.3.33.

Possibility of inconsistent data on Tivoli Directory Server when transaction updates are replicated in an environment with failover setup

When transactional updates are replicated by a supplier, the updates are not replicated in a transactional manner by the supplier to its consumers. In a replicated environment the supplier only replicates the transactional updates to its consumers when the transaction is complete (committed or rolled back state). If a supplier goes offline when replicating the transactional updates, it is possible that only a part of the update is replicated to its consumers. In this case, when the supplier is brought online the remaining updates that are in its replication queue is replicated automatically.

However, in the case of a replication environment with failover, if the primary master fails when replicating the updates, the proxy server will failover to the peer server. In this case, the data might not be entirely consistent because it is possible for the peer server to not get all the updates made to master.

IBM Tivoli Directory Server might fail to create the default directory server instance

On AIX, Solaris, and Linux systems, if Tivoli Directory Server fails to create the default directory server instance it might be due to one of the following reasons:

- Not enough disk space in the /home or /export/home directory

- The root user might not have write permission on the /home or /export/home directory

Unable to log on to a system, when migrated users use LDAP - operating system authentication mechanism

Tivoli Directory Server does not support the password encryption mechanism that UNIX supports because of which the migrated users are not able to log on to the system using LDAP - operating system authentication mechanism.

The UNIX system uses a mix of MD5 and CRYPT password encryption scheme, which Tivoli Directory Server does not support. Though, Tivoli Directory Server do support CRYPT and MD5 encryption schemes.

The workaround for this problem are the following:

- LDAP administrator can reset the user password for the migrated users on the LDAP machine.
- Create new users on the LDAP machine for LDAP - operating system authentication.

If the backend server configured as primary write server is from earlier versions of Tivoli Directory Server, then the backend server rejects the propagated schema updates with an error

When schema updates are requested by a global administrator group member, the schema updates are first applied to the Tivoli Directory Proxy server v6.3, and then the updates are propagated to its backend servers. It is observed that if the backend server created using the earlier versions of Tivoli Directory Server is configured as the primary write server, then the backend server rejects the schema updates with an appropriate error.

In an environment where backend servers are configured using multiple versions of Tivoli Directory Servers, to ensure that the schema updates propagated from the Tivoli Directory Proxy server v6.3 are applied to the backend servers, user must configure the backend server that need to be set as primary write server using the Tivoli Directory Server v6.3.

The Accessibility tool, JAWS, is not able to read the message displayed on two dialog boxes of the Configuration Tool

When using the JAWS tool on Configuration Tool, the JAWS tool is not able to read out the message displayed on two dialog boxes because of the limitation in design that is used in implementing the messages. The following messages that are displayed on the dialog boxes are not read by the JAWS tool:

- Configuration of the instance has changed, causing this task to become invalid. Would you like to dispose this task?
- Are you sure you want to close this window?

General troubleshooting

The following sections describe general problems and solutions in IBM Tivoli Directory Server 6.1 and later versions of directory servers.

Instance owner unable to access core file for core that occurred during server initialization

If the root user starts the server and a core file is produced early during initialization of the server, the core file might not be accessible to the instance owner user. Instead, the root user has access to the core file.

If this error occurs, the root user can manually set the core file's ownership to the instance owner user if desired.

This problem occurs only on AIX, Linux, and Solaris operating systems.

Key label in .kdb file and ibmslapd.conf file do not match.

If the key label in the SSL key database certificate does not match the key label in the IBM Tivoli Directory Server configuration file (ibmslapd.conf), you will receive the following error:

```
The default SSL key database certificate is incorrect in file
c:/keytabs/pd_ldapkey.kdb.
```

Check the key label in the configuration file and the SSL key database certificate. If they do not match, create a self-signed SSL key database certificate that matches the key label in the configuration file. For more detailed information about how to create a self-signed key database certificate, see the *IBM Tivoli Directory Server Version 6.3 Administration Guide*.

GSKit certificate error

If you are trying to import a signer or personal certificate from an external certificate authority (CA) such as Entrust and the GSKIT fails with the error,

```
An error occurred while receiving the certificate from the given file.
```

the problem might be that the certificate returned from Entrust is a chain certificate, not a root certificate. You must have a root certificate to start a certificate chain. A chain certificate cannot start a certificate chain.

If you do not already have a root certificate, the following is one method of obtaining one.

An example of a root certificate is GTE Cybertrust, which is included in Internet Explorer (IE) 5.5; however, it is not included by default in the GSKit kdb database. To obtain this certificate you must:

1. Export one of the GTE Cybertrust certificates (there are 3) from Internet Explorer as Base64 encoded.
2. Add the certificate as a trusted root certificate.

   Note: In order to use the GSKit option to set a certificate as a trusted root, the certificate must be self-signed.
3. Add the chain CA certificate from Entrust.
4. Receive the SSL certificate from Entrust.

Server instance fails to start because of incorrect file permissions

On AIX, Linux, and Solaris systems, file permissions are frequently altered inadvertently by the actions of copying or editing a key database file. Because

these actions are generally done as the user ID **root**, file permissions are set for the user **root**. For the directory server instance to make use of this file, you must change the file permissions so that it is readable by the user ID **idsldap**. Otherwise the directory server instance fails to start.

```
chown idsldap:idsldap <mykeyring>.*
```

Server instance fails to start because localhost hostname is set incorrectly

The localhost hostname must correspond to the local loopback address of 127.0.0.1. If localhost is renamed or the TCP/IP address has changed, the directory server instance does not start.

Server instance cannot be started except by instance owner

On AIX, Linux, and Solaris systems, if a user other than the directory server instance owner cannot start the directory server instance, be sure that the following are true:

- The user who is attempting to start the directory server instance is a member of the primary group of the directory server instance owner.
- The directory server instance owner's primary group has Write access to the location where the database was created.

See "Setting up users and groups: directory server instance owner, database instance owner, and database owner" in the *IBM Tivoli Directory Server Version 6.3 Installation and Configuration Guide* for information about requirements for the directory server instance owner, database instance owner, and database owner.

Error opening slapd.cat file on Windows systems

On Windows systems, you might receive an error message that includes the following:

```
Error opening slapd.cat
Plugin of type DATABASE is successfully loaded from
   C:/Program Files/IBM/LDAP/V6.2/bin/libback-config.dll.
Error opening rdbm.cat
```

If this occurs, check the NLSPATH environment variable. The installation program sets the NLSPATH environment variable as a system environment variable. However, if the system also has the NLSPATH variable set as a user environment variable, the user NLSPATH environment variable overrides the system setting.

To correct this, append the NLSPATH information from the system environment variable to the information in the user environment variable.

DSML file client produces error

The DSML file client produces the following error when it is set up to communicate using SSL and a user tries to connect to an LDAP server that does not use SSL.

```
SSL IS ON
javax.naming.CommunicationException: 9.182.21.228:389.  Root exception is javax.
net.ssl.SSLProtocolException: end of file
    at com.ibm.jsse.bd.a(Unknown Source)
    at com.ibm.jsse.b.a(Unknown Source)
    at com.ibm.jsse.b.write(Unknown Source)
    at com.sun.jndi.ldap.Connection.<init>(Connection.java:226)
    at com.sun.jndi.ldap.LdapClient.<init>(LdapClient.java:127)
```

```
at com.sun.jndi.ldap.LdapCtx.connect(LdapCtx.java:2398)
at com.sun.jndi.ldap.LdapCtx.<init>(LdapCtx.java:258)
at com.sun.jndi.ldap.LdapCtxFactory.getInitialContext(LdapCtxFactory.java:91)
at javax.naming.spi.NamingManager.getInitialContext(NamingManager.java:674)
at javax.naming.InitialContext.getDefaultInitCtx(InitialContext.java:255)
at javax.naming.InitialContext.init(InitialContext.java:231)
at javax.naming.InitialContext.<init>(InitialContext.java:207)
at javax.naming.directory.InitialDirContext.<init>(InitialDirContext.java:92)
at com.ibm.ldap.dsml.DsmlRequest.processRequests(DsmlRequest.java:767)
at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:253)
at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:402)
at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:373)
at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:296)
at com.ibm.ldap.dsmlClient.DsmlFileClient.main(DsmlFileClient.java:203)
```

The error is not fatal and the output XML file is generated.

Non default log files need valid path

If you want to store your log files in a nondefault path, you must ensure that the path exists and is valid. You must create the directory before you can configure the log files.

Null searches retrieve entries of deleted suffixes

A null search (`ldapsearch -s sub –b "" objectclass=*`) returns all the entries found in the database. If you have deleted a suffix without first removing its entries from the database, those entries are returned by the null search even though the suffix no longer exists.

The idsldapsearch command with -h option gives error with the DIGEST-MD5 mechanism

The DIGEST-MD5 SASL bind mechanism requires that the client be able to resolve the fully-qualified host name of the server. If the client cannot resolve the server's fully-qualified hostname the bind fails with an LDAP_PROTOCOL_ERROR. To correctly resolve the host name, you might need to make system changes or make DNS configuration changes, such as enabling reverse DNS mapping.

For example, AIX, Linux, Solaris, and HP-UX (Itanium) systems have lines in the /etc/hosts file with the syntax:

```
<IP address><fully qualified distinguished name><alias>
```

This syntax is used to define the local hostname to the IP address mappings.

If the syntax is something like:

```
127.0.0.1 localhost
```

when localhost is resolved, it is seen as the fully qualified distinguished name of the system. This causes DIGEST-MD5 to fail.

For the DIGEST-MD5 mechanism to work correctly, the syntax must be something like:

```
127.0.0.1 ldap.myserver.mycompany.com  localhost
```

The syntax of the line is now such that `ldap.myserver.mycompany.com` is a valid fully qualified distinguished name for the localhost system.

After enabling language tags, do not disable language tags

After enabling the language tag feature, if you associate language tags with the attributes of an entry, the server returns the entry with the language tags. This occurs even if you later disable the language tag feature. Because the behavior of the server might not be what the application is expecting, to avoid potential problems, do not disable the language tag feature after it has been enabled.

Create the key database certificate before setting up SSL

Before setting up SSL communications on your server, you must use the GSKit utility, **ikeyman**, to create the necessary certificates. See "Using ikeyman" and "Secure Sockets Layer" in the *IBM Tivoli Directory Server Version 6.3 Administration Guide*.

idsbulkload appears to hang during parsing phase

The **idsbulkload** utility has special code to handle nested groups, and the extra processing takes time.

For example, if an LDIF file contains 50,000 nested groups with 100 membergroups in each of the nested groups, **idsbulkload** might need about 1 to 2 seconds to process each one of the nested groups during the parsing phase.

In this case, **idsbulkload** appears to hang before showing any progress.

An environment variable, BULKLOAD_REPORT_CHUNK, can be used to increase the frequecy of progress reporting.

Set the variable to a positive integer value; for example, 100. Use the following commands:

- On AIX, Linux, and Solaris systems: `export BULKLOAD_REPORT_CHUNK=100`
- On Windows systems: `set BULKLOAD_REPORT_CHUNK=100`

idsbulkload will then report parsing progress at 100 entry interval. For example:

```
...
GLPBLK061I Parsing entries ...
GPBLK004I 100 entries parsed successfully out of 100 attempts.
LPBLK004I 200 entries parsed successfully out of 200 attempts.
..
```

Tivoli Directory Server may crash if the size of any log file exceeds the system file size limit

When the size of any log file grows beyond the system file size limit, Tivoli Directory Server may crash. This typically occurs when tracing is enabled on the server.

Not able to connect to Tivoli Directory Server over SSL while copying an instance using the idsxinst tool

The reason for this problem could be incorrect configuration. To resolve this problem, perform the following steps:

1. Verify that GSKit is installed on the server.
2. Verify that the gskikm.jar file is present in the <tds_ldap_home>/java/jre/lib/ext directory.

3. In the java.security file under the <tds_ldap_home>/java/jre/lib/security directory, check if the CMS provider entry exists. If the entry does not exist, add this entry in the java.security file by entering the following:

   ```
   security.provider.X=com.ibm.spi.IBMCMSProvider
   ```

   where, X is the next number in the order.
4. Ensure that /lib exists in the system path.
5. While connecting to source server over SSL, providing the 'Key name' is not mandatory and can be left blank.

Tivoli Directory Server fails to start or displays error when performing ldap operations after bulkload is done

After performing bulkload, if Tivoli Directory Server fails to start or displays error when performing LDAP operations, it could be because of one of the following reasons:

- Check the log file, db2diag.log, if there is an error that states "ACCESS TABLE WHEN IN RESTRICTED STATE". This means that loading data or bulkload was not complete or was unsuccessful.
- The table is in the “Load Pending” or “Locked” state. A previous LOAD attempt on the table might have resulted in failure. Accessing the table is not allowed until the LOAD operation is restarted or terminated.

Consider the following options to rectify the problem:

- Stop or restart the failed LOAD operation on the table by issuing LOAD with the TERMINATE or RESTART option.
- Check if the bulkload_status file is present. This file is created in the home directory of the instance. If this file is present, it means that bulkload was unsuccessful. Check the file for errors and rectify it, and try running the bulk load utility again.

Migration fails if Tivoli Directory Server V6.0 is configured with DB2 v8 and the environment variables are set for a different version of DB2

Migration might fail if either one of these conditions exists:

- If Tivoli Directory Server V6.0 is configured with DB2 v8 and the environment variables are set for a different version of DB2.
- If Tivoli Directory Server V6.0 is configured with DB2 v9 and the environment variables are set for a different version of DB2.

To resolve this, you must ensure that:

- When you migrate Tivoli Directory Server V6.0 configured with DB2 v8 the environment variables set on the system are of DB2 v8.
- When you migrate Tivoli Directory Server V6.0 configured with DB2 v9 the environment variables set on the system are of DB2 v9.

The following environment variables must be updated depending on the DB2 version in use:

- PATH
- CLASSPATH
- INCLUDE

- LIB
- DB2INSTANCE

The idsadsrun tool might fail for some instances when run simultaneously for multiple instances on the same machine

When running the idsadsrun tool simultaneously for multiple instances on the same machine, if the user gets the following exception: "org.apache.derby.client.am.DisconnectException: java.net.ConnectException : Error opening socket to server <host_name> on port <port_number> with message : Connection refused", then user must apply the latest available fixpack.

To get this fix, go to Tivoli Directory Integrator support site: http://www-306.ibm.com/software/sysmgmt/products/support/IBMDirectoryIntegrator.html?S_CMP=rnav

On Windows operating system, Tivoli Directory Server startup messages might get displayed in two different locales when a language other than English is specified for Tivoli Directory Server

If Tivoli Directory Server startup messages are being displayed in two different locales, the most likely reason for the problem is that the currently logged in user and the Tivoli Directory Server instance owner have different set of locale configured on the system.

You can consider one of the following ways to rectify this problem:

- Set the LANG environment variable explicitly to the language you want use. For example, set LANG=de_DE (or any other supported language). You must then start the server from the same window.
- Modify the regional and language settings on the Regional and Language Options dialog box to ensure that both the currently logged in user and the instance owner have the same set of regional and language settings to view the server messages in the same language.

Unable to open a new connection for an LDAP client to connect to Tivoli Directory Server running on a Linux or Solaris operating system

On Linux and Solaris operating system, there is a limit on the maximum number of file descriptors that can be opened by a process. For Linux and Solaris operating systems, the default value of the maximum number for open file descriptor is 1024 and 256, respectively.

A Tivoli Directory Server V6.1 instance uses 15 file descriptors for the purpose of logging messages. So on Linux, a Tivoli Directory Server V6.1 instance stops accepting new connections after 1009, that is 1024 – 15, concurrent client connects. Whereas on Solaris, a Tivoli Directory Server V6.1 instance stops accepting new connections after 241, that is 256 – 15, concurrent client connects. If an error is encountered while accepting new connections appropriate message is logged. This error does not affect any existing connections only new LDAP clients will fail to connect to the Directory server.

In order to increase the maximum open file descriptors, user should issue the following command and restart the server from the same command prompt.

```
#ulimit –Hn <number of connections>
```

Note: The performance with very high number of concurrent client connections depends on the hardware and the operations being performed. With thousands of concurrent client connections sending operations simultaneously the performance of the directory server may decrease.

When deploying a replica or a peer in a replication environment using the idsideploy tool, if the tool detects more than one entry with same replica serverID and ibm-replicationServerIsMaster=true, the tool throws an error

When deploying a replica or a peer in a replication environment using the idsideploy tool, if the tool detects more than one replication subentry containing the same serverID value for the attribute ibm-replicaServerId with the attribute ibm-replicationServerIsMaster set to true, the tool throws error.

For any given replication context, multiple replication subentries are not required, only one replication subentry is required. For example, if the entries are made as shown in the below example, idsideploy will fail.

```
dn: ibm-replicaServerId=Peer1,ibm-replicaGroup=default, ou=ounit1, o=sample
objectclass: top
objectclass: ibm-replicaSubentry
ibm-replicaServerId: Peer1
ibm-replicationServerIsMaster: true
cn: Peer1
description: Peer1

dn: cn=Peer1_entry,ibm-replicaGroup=default, ou=ounit1, o=sample
objectclass: top
objectclass: ibm-replicaSubentry
ibm-replicaServerId: Peer1
ibm-replicationServerIsMaster: true
cn: Peer1_entry
description: Peer1
```

In the above example, to rectify the problem users should create only one entry.

The idsadscfg, idssnmp, and idslogmgmt tools might throw error if the environment variable values contain spaces

If you have installed IBM Tivoli Directory Integrator in a different location other that the default location, set the following environment variable:

- For the Log management (idslogmgmt) tool, Active Directory synchronization (idsadscfg), and SNMP (idssnmp) tools function correctly, you must explicitly set the IDS_LDAP_TDI_HOME environment variable to point to the directory where you installed Tivoli Directory Integrator.

The value set using the environment variable IDS_LDAP_TDI_HOME must not have space or double quotes, then the tools will not work properly. On Windows, the tools work properly when tilde, "~" (that is, short path or path with no spaces) is used.

The idsadsrun tool stops and exists when attempting to synchronize Active Directory and Tivoli Directory Server after restarting Tivoli Directory Server

When running idsadsrun tool in full synchronization mode after configuring Active Directory synchronization solution, it is observed that even if Tivoli Directory Server instance that is associated with Active Directory synchronization stops the idsadsrun tool remains in active state. However, on updating Active Directory for entries it has been observed instead of synchronizing with Tivoli Directory Server instance the tool stops and exists.

The idsadsrun tool after performing a pass of full synchronization of Tivoli Directory Server instance with Active Directory, it then performs real-time synchronization. In the above situations, if a pass of full synchronization is done and real-time synchronization is running then the user can restart the Active Directory synchronization solution to run in real-time synchronization mode.

The idsadsrun utility might fail to synchronize as is from Active Directory Server to Tivoli Directory server

When running the idsadsrun utility to synchronize a large number of users and groups entries (for example, above 100000 users and 10000 groups) from Active Directory server to Tivoli Directory Server, you might observe that all entries in Active Directory Server have not completely synchronized in Tivoli Directory Server. It is observed that some user entries and their corresponding group entries might fail to sync.

This might occur depending on the system configuration and the JVM parameters tuned on your system . For example, for a machine with 1 GB RAM, the parameter values can be Xms254m-Xmx1024m. You can tune the parameters based on your system configurations.

For best results, users can also configure Microsoft Active Directory setting parameters to the following:

```
MaxPoolThreads 4
MaxDatagramRecv 4096
MaxReceiveBuffer 10485760
InitRecvTimeout 120
MaxConnections 5000
MaxConnIdleTime 900
MaxPageSize 1000000
MaxTempTableSize 10000
MaxResultSetSize 262144
MaxNotificationPerConn 5
MaxValRange 1500
```

The idscfgdb command might fail to configure a database for a directory server instance on Red Hat Enterprise Linux (RHEL) 4 64-bit operating system

When configuring a database for a directory server instance using the idscfgdb command on RHEL 4 64-bit operating system, the tool might exit with the following error messages:

```
GLPCTL026I Creating database: 'mydata'.
GLPCTL028E Failed to create database: 'mydata'. The failure might
have occurred because the system was not set up correctly before using the tool.
GLPCTL011I Stopping database manager for the database instance: 'mydata'.
GLPCTL012I Stopped database manager for the database instance: 'mydata'.
```

```
GLPCDB004E Failed to add database 'mydata' to directory server instance:
'mydata'.
GLPCDB026W The program did not complete successfully. View earlier error messages
for information about the exact error.
```

In this situation, the db2diag.log file might contain the following information:

```
2008-03-20-11.33.32.471455+330 I5410E982          LEVEL: Error
PID     : 3214                 TID  : 182960860768 PROC : db2fm
INSTANCE: mydata                NODE : 000
FUNCTION: DB2 Common, Generic Control Facility, gcf_stop, probe:30
MESSAGE : ECF=0x9000036D=-1879047315=ECF_FM_DB2FMD_PROCESS_NOT_EXIST
          There is no fault monitor daemon running
CALLED  : OS, -, open
RETCODE : ECF=0x9000001A=-1879048166=ECF_FILE_DOESNT_EXIST
          File doesn't exist
CALLSTCK:
  [0] 0x0000002A956FF982 /opt/ibm/db2/V9.5/lib64/libdb2osse.so.1 + 0x1A7982
  [1] 0x0000002A956FF82F ossLogRC + 0x6B
  [2] 0x0000002A9BCBDB01 gcf_stop + 0x42B
  [3] 0x0000002A95DB9559 _ZN9GcfCaller4stopEP12GCF_PartInfomP11GCF_RetInfo + 0x105
  [4] 0x0000000000405708 main + 0x17F0
  [5] 0x0000003049D1C3FB __libc_start_main + 0xDB
  [6] 0x0000000000403E7A __gxx_personality_v0 + 0x9A
  [7] 0x0000000000000000 ?unknown + 0x0
  [8] 0x0000000000000000 ?unknown + 0x0
  [9] 0x0000000000000000 ?unknown + 0x0


2008-03-20-11.33.32.472141+330 I6393E948          LEVEL: Error
PID     : 3214                 TID  : 182960860768 PROC : db2fm
INSTANCE: mydata                NODE : 000
FUNCTION: DB2 Common, Fault Monitor Facility, db2fm, probe:170
MESSAGE : ECF=0x90000349=-1879047351=ECF_FM_FAIL_TO_STOP_GCF_FM
          Failed to stop the GCF fm module
CALLED  : DB2 Common, Generic Control Facility, GcfCaller::stop
DATA #1 : signed integer, 8 bytes
0
DATA #2 : unsigned integer, 8 bytes
1CALLSTCK:
  [0] 0x0000002A956FF982 /opt/ibm/db2/V9.5/lib64/libdb2osse.so.1 + 0x1A7982
  [1] 0x0000002A956FF883 ossLogRC + 0xBF
  [2] 0x0000000000405772 main + 0x185A
  [3] 0x0000003049D1C3FB __libc_start_main + 0xDB
  [4] 0x0000000000403E7A __gxx_personality_v0 + 0x9A
  [5] 0x0000000000000000 ?unknown + 0x0
  [6] 0x0000000000000000 ?unknown + 0x0
  [7] 0x0000000000000000 ?unknown + 0x0
  [8] 0x0000000000000000 ?unknown + 0x0
  [9] 0x0000000000000000 ?unknown + 0x0
```

To resolve the above problem, you can do the following:

- Update the kernel parameter, kernel.shmmax, and run the idscfgdb tool again. For example, kernel.shmmax = 3221225472

The idscfgdb command might fail while creating a database with error code GLPCTL028E

On AIX, Linux, and Solaris systems, the idscfgdb command might fail while creating a database.

An example of the db2cli.log file with the information logged:

```
retcode = 1478; state = "01626"; message = "SQL1478W
The defined buffer pools could not be started.
Instead, one small buffer pool for each page size supported by DB2 has been started.
SQLSTATE=01626
```

An example of the db2diag.log file with the information logged:

```
MESSAGE : ZRC=0x850F0005=-2062614523=SQLO_NOSEG
          "No Storage Available for allocation"
          DIA8305C Memory allocation failure occurred.
DATA #1 :
Unable to attach 3 segments totalling 2478440448 bytes starting at address
0x0000000000000000.  One possible cause may be an improper setting for the
shmmax Linux kernel tuneable.
```

To know more about tuning kernel parameters on Linux systems, see Modifying kernel parameters (Linux). To know more about tuning kernel parameters on Solaris, see Modifying kernel parameters (Solaris operating system).

The idscfgdb command might not extend DMS cooked tablespace size exactly in the multiples of the value provided with the -z option

When using the idscfgdb command to provide a value with the -z option by which DMS cooked tablespace should be extended, this might not be reflected exactly when running the command. The actual value used to extend the tablespace might be slightly lesser or greater than the value specified. This is because the database manager strives to maintain consistent growth across tablespace containers.

Compatibility issue with Common Auditing and Reporting Service (CARS) 6.0.1 server

The CARS logging feature provided with IBM Tivoli Directory Server 6.2 uses the CARS 6.1 client. Therefore, the CARS 6.1 server is required for using CARS 6.1 clients. Any version of CARS server other than 6.1 will not be compatible with the CARS logging tool.

User might face problem when monitoring Tivoli Directory Server instances on a Solaris machine using an SNMP agent

On a Solaris machine, user might face problem in monitoring Tivoli Directory Server instances using an SNMP agent trying to logon using SSH from Tivoli Directory Integrator.

In such situations, user is required to start an rsh session on the Solaris machine and then try logging using rsh on to the Solaris machine. After logging on to the Solaris machine, user can monitor Tivoli Directory Server instances using an SNMP agent.

The idsdbrestore utility displays error messages if the ldapdb.properties file is modified

The idsdbrestore utility refers the /opt/ibm/ldap/V6.2/etc/ldapdb.properties file and not the instance specific ldapdb.properties file located in the <install-home>/idsslapd-<instance-name>/etc directory during a restore operation.

If a user has updated or modified the currentDB2InstallPath parameter in the ldapdb.properties file to a different DB2 installation path or to a different DB2

major version after the directory server instance creation, error messages are displayed when performing a restore operation.

To resolve this problem, user can temporarily copy the <install-home>/idsslapd-<instance-name>/etc/ldapdb.properties file to the /opt/ibm/ldap/V6.3/etc directory before performing a restore request using the idsdbrestore utility. After idsdbrestore completes the request, restore the original ldapdb.properties file.

The idsxinst tool fails to run and generates a coredump file or the tool does not display the directory server instances present on the system if it gets launched

If the system does not have the requisite filesets or has corrupted filesets for the configured locale, then either of the following might happen:

- The idsxinst tool might fail to run and might generate a coredump file.
- If the idsxinst tools launches, it might not display the directory server instances present on the system.

To resolve this problem, ensure that all the filesets required for the configured locale are installed.

Backup and restore using the Configuration tool do not function when provided with path in Unicode string

When entering paths on the GUI tools, ensure that the path specified can be represented on the system. This is because the string provided for file path must be representable in the system's local code page as GUI translates the Unicode input to local code page. For example, if the Unicode input for the file path contains Chinese characters that is provided on a system with French locale, the translated file path will not be a valid path.

Tivoli Directory Sever starts in config-only mode when migrating from an earlier version using the Instance administration tool

When providing values for migration using the GUI, Instance administration tool (idsxinst), the input values should be such that it can be representable in the system's local code page. Otherwise, you might experience problems when starting the Tivoli Directory Server.

In a Tivoli Directory Server environment, a warning message with message code GLPSRV147W might get displayed

In a Tivoli Directory Server environment, a warning message with code GLPSRV147W might get displayed. This might be because of the default value of write timeout that is set to 10 seconds.

If you see this error frequently for your Tivoli Directory Server environment, you must consider increasing the write timeout value by modifying the ibm-slapdWriteTimeout attribute under the entry DN "cn=Connection Management, cn=Front End, cn=Configuration".

You can either use the Web administration tool or the ldapmodify command to change the value of ibm-slapdWriteTimeout. To change the value, issue the ldapmodify command of the following format:

```
#idsldapmodify -D <adminDN> -w <password> -i <filename>
```

where <filename> contains:

```
dn: cn=Connection Management,cn=Front End, cn=Configuration
changetype: modify
replace: ibm-slapdWriteTimeout
ibm-slapdWriteTimeout: 120
```

Platform specific problems

This information applies to the following operating systems:

For AIX only

The following information applies only to the AIX operating system.

Problem with MALLOCTYPE=buckets

Before setting MALLOCTYPE to **buckets** on the AIX 5.2 operating system, ensure that you have installed the patch for APAR IY50668. Otherwise the LDAP server might fail with a core file.

Verifying that AIX hardware is 64-bit

The server on AIX requires 64-bit hardware. To verify that your AIX hardware is 64-bit, run the following command:

```
bootinfo -y
```

If the command returns 32, your hardware is 32-bit.

In addition, if you type the command `lsattr -El proc0`, the output of the command returns the type of processor for your server.

Verifying that the AIX kernel is 64-bit

To verify that you have the 64 bit kernel (/usr/lib/boot/unix_64) installed and running, run the following command:

```
bootinfo -K
```

In addition, if you type the command `lsattr -El proc0`, the output of the command returns the type of processor for your server.

Note: If the hardware is 32-bit, then you can only have a 32-bit kernel. You cannot have a 64-bit kernel. If the hardware is 64-bit, then you can have either a 32 or 64-bit kernel.

To switch between a 32-bit and 64-bit mode at the operating system level on AIX 5.3:

When you install the operating system, go to Additional features and specify 64-bit mode. (The default is 32-bit mode.) To switch from 32-bit mode to 64-bit mode, use the following commands:

```
# ln -sf /usr/lib/boot/unix_64 /unix
# ln -sf /usr/lib/boot/unix_64 /usr/lib/boot/unix
# bosboot -ad /dev/ipldevice
# shutdown -Fr
# bootinfo -K
```

The kernel is now in 64-bit mode.

To switch from 64-bit mode to 32-bit mode, use the following commands:

```
# ln -sf /usr/lib/boot/unix_mp /unix
# ln -sf /usr/lib/boot/unix_mp /usr/lib/boot/unix
# bosboot -ad /dev/ipldevice
# shutdown -Fr
# bootinfo -K
```

The kernel is now in 32-bit mode.

Error on AIX when running db2start

The following error might occur when you try to run **db2start**:

```
0509-130 Symbol resolution failed for /usr/lib/threads/libc.a(aio.o)
because:
         0509-136   Symbol kaio_rdwr (number 0) is not exported from
                    dependent module /unix.
         0509-136   Symbol listio (number 1) is not exported from
                    dependent module /unix.
         0509-136   Symbol acancel (number 2) is not exported from
                    dependent module /unix.
         0509-136   Symbol iosuspend (number 3) is not exported from
                    dependent module /unix.
         0509-136   Symbol aio_nwait (number 4) is not exported from
                    dependent module /unix.
         0509-192 Examine .loader section symbols with the
                  'dump -Tv' command.
```

If this occurs on AIX, you have asynchronous I/O turned off.

To turn on asynchronous I/O:

1. Run **smitty chgaio** and set **STATE to be configured at system restart** from **defined** to **available**.
2. Press Enter.
3. Do **one** of the following:
   - Restart your system.
   - Run **smitty aio** and move the cursor to **Configure defined Asynchronous I/O**. Then press Enter.

The **db2start** command now works.

On AIX 6.1 with workload partitions (WPARs) configured, starting a Kerberos service might fail

When WPARs are configured on an AIX 6.1 system, starting a Kerberos service might fail, and would display the following message:

```
Starting krb5kdc...
Unable to bind server socket on port 88.
Unable to initialize network.
        Status 0x44 - The socket name is not available on this system..
krb5kdc could not be started.
```

This problem is because of a limitation with IBM Network Authentication Service (NAS) 1.4. To resolve this, use versions above IBM NAS 1.4.

IBM Tivoli Directory Server utilities like ldif2db and bulkload might stop or exit on an AIX 6.1 system

When using server utilities such as ldif2db and bulkload with directory server on an AIX 6.1 system, the utility might stop or exit with an error. An example trace of the ldif2db utility is as follows:

```
253:20:40:16 T1 K2244835 do_iconv_open: local_codepage=NULL
253:20:40:16 T1 K2244835 xlate_utf8_to_localcp: inlen=8
253:20:40:16 T1 K2244835 xlate_utf8_to_localcp: rc=0
GLPRDB002W ldif2db: 50 entries have been successfully added out of 50 attempted.
253:20:40:16 T1 K2244835 vPrintMessage: catid=2, level=2, num=2.
253:20:40:16 T1 K2244835 do_iconv_open: local_codepage=NULL
253:20:40:16 T1 K2244835 xlate_utf8_to_localcp: inlen=8
253:20:40:16 T1 K2244835 xlate_utf8_to_localcp: rc=0
253:20:40:16 T1 K2244835 close_one_backend: calling be->be_close
253:20:40:16 T1 K2244835 close_one_backend: calling be->be_close
253:20:40:16 T1 K2244835 calling config_close...
253:20:40:16 T1 K2244835 close_one_backend: calling be->be_close
253:20:40:16 T1 K2244835 calling rdbm_close...
253:20:40:16 T1 K2244835 leaving rdbm_close...
./ldif2db[1040]: 880864 Segmentation fault(coredump)
```

This is a problem with AIX 6.1 base level. To resolve this, use AIX 6.1 with FP1 or later fix levels.

The idsidrop command generates java core when a Tivoli Directory Server instance is dropped from IBM Power7 system with AIX 7.1

IBM Tivoli Directory Server v6.3 uses IBM DB2 v9.7 fixpack 2, which includes JRE that is used by IBM Tivoli Monitoring Agent for DB2. This JRE version is not supported on IBM Power7 system.

Workaround

If a user has already installed DB2 v9.7 manually on the system for use with Tivoli Directory Server, then user should uninstall the IBM Tivoli Monitoring Agent for DB2.

Procedure

1. Using the DB2 instance owner credentials stop all the Monitoring Agent for DB2 processes. Run the following command:

   ```
   <DB2_Directory>/itma/bin/itmcmd agent -o instance stop ud
   ```

   where, *<DB2_Directory>* is the directory where DB2 copy of Tivoli Monitoring Agent is installed.

   Note: If multiple DB2 instances are being monitored then there can be multiple kuddb2 processes that need to be stopped.
2. On AIX and Linux operating systems, uninstall Tivoli Monitoring Agent for DB2. Run the following command:

   ```
   <DB2_Directory>/itma/bin/uninstall.sh REMOVE EVERYTHING
   ```

For more information about uninstalling IBM Tivoli Monitoring for databases, see Uninstalling IBM Tivoli Monitoring for Databases: DB2 Agent with the DB2 installer.

For Windows 2000, Windows Server 2003 Enterprise, Windows XP, Windows Server 2003 R2 Datacenter Edition, Windows Server 2008, and Windows 7 only

The following sections apply only to the Windows 2000, Windows Server 2003 Enterprise, Windows Server 2003 R2 Datacenter Edition, Windows Server 2008, and Windows XP and Windows 7 client platforms.

Setting LANG and LC_ALL system environment variables for nonEnglish InstallShield GUI installation

For the InstallShield GUI installation to bring up the same language that the operating system is using, two variables must be set in the system environment

- LANG = <*locale*>
- LC_ALL = <*locale*>

where <*locale*> is the locale that the operating system is using.

Go to http://www.microsoft.com/globaldev/ for a list of Microsoft locale values.

Certain UTF-8 supplementary characters do not display correctly

IBM Tivoli Directory Server supports UTF-8 (Unicode Transformation Format, 8-bit form) to use Unicode characters, which contains MS932 (Shift JIS) characters plus supplementary characters not defined in MS932. Supplementary characters might be displayed as square box in Internet Explorer running on Windows 2000. See Figure 1.

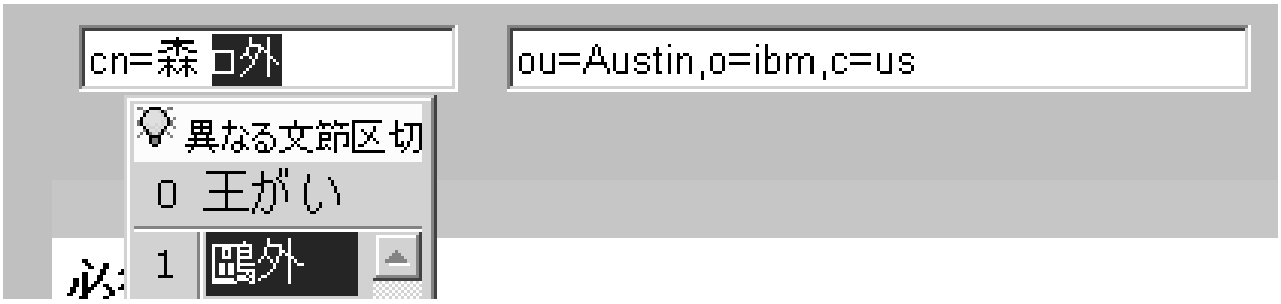

Figure 1. Unicode Code Point U+9DD7 displayed as a square

If this occurs, install one of the East Asian language kits. Depending on your environment, install the Japanese, Korean, Simplified Chinese or Traditional Chinese language kit which is included in your Windows CDs. For example, Unicode code point U+9DD7 is one of the supplementary characters in the Japanese environment. With the correct language kit installed, the supplementary character is displayed correctly. See Figure 2.

Figure 2. U+9DD7 displayed correctly

Note: This problem has not been observed in Windows XP.

Communications error: Exceeding 64 connections/OCH

On Windows, if you have clients that are generating many connections to the server and the connections are being refused, the server might log error messages similar to the following in the ibmslapd.log file:

```
Feb 11 14:36:04 2004 Communications error:
                Exceeding 64 connections/OCH - dropping socket.
```

If you see these errors, do the following:

1. Stop the server.
2. Save a copy of your ibmslapd.conf file.

3. Insert the following in the section that starts with 'dn: cn=FrontEnd,cn=Configuration':

   ```
   ibm-slapdSetenv: SLAPD_OCHANDLERS=5
   ```

4. Restart your server.

If you continue to receive error messages, increase the value of the SLAPD_OCHANDLERS environment variable by 5 until you stop receiving error messages.

Starting IBM Tivoli Directory Server at operating system startup

In IBM Tivoli Directory Server, the server (the **ibmslapd** process) is started manually through the Services window or by the **ibmslapd** command. If you try to start the server automatically by updating the **Startup Type** in the Services window to **Automatic**, errors occur when you restart the computer. This is because DB2 must be running before the **ibmslapd** process can start.

If you want the server to start automatically, you can create a batch file to start the **ibmslapd** process. The batch file should be invoked after all the services are started, so that DB2 will be completely up and running before the **ibmslapd** process starts.

The following is an example of commands in a .bat file that you can add to the Startup folder to start the server:

```
@echo off
%LDAPHome%\bin\ibmdirctl [-h <hostname>] [-D <adminDN>] [-w <password>]
         [-p <portnumber>] start -- [ibmslapd options]
```

Note: Be sure that the **Startup Type** for the **IBM Tivoli Directory Administration server** entry in the Services window is set to **Automatic**. If it is not, the administration server control program (**ibmdirctl**) will not work.

Backup and restore

Scenario

On Windows platform, when performing a backup, restore, or load to or from a directory mapped as remote drive using DB2 utilities fails giving error message “SQL2036N The path for the file or device "<file_or_devicename>:\" is not valid”.

Reason

When a user tries to perform a backup, restore, or load to or from a directory mapped as remote drive, for instance, H:\MyFolder\test; user gets the error.

There are two different reasons for getting this error:

1. The user is specifying an invalid shared drive in the command.

   ```
   db2 backup db mydata to H:\
   ```

   Error message displayed:

   SQL2036N The path for the file or device "H:\" is not valid.

2. The user is specifying a valid UNC name for the mapped drive but he is not using the right USERID.

   ```
   db2 backup db mydata to \\MyFolder\test
   ```

   Error message displayed:

SQL2036N The path for the file or device "\\MyFolder\test "is not valid.

For the db2 backup failing the db2diag.log looks like:

```
database_utilities sqlubcka Probe:0 Database:mydata

Starting a full database backup.

2006-06-10-10.42.09.175000 Instance:DB2 Node:000
PID:2404(db2syscs.exe) TID:2500 Appid:none
database_utilities sqluMCTestDevType4Backup Probe:60

Media controller -- invaliddevice path: H:\MyFolder\test


2006-06-10-10.42.09.253000 Instance:DB2 Node:000
PID:2404(db2syscs.exe) TID:2576 Appid:*LOCAL.DB2.030610154019
database_utilities sqlubcka Probe:0 Database:mydata

Backup terminated.
```

Workaround

This error is because of Windows restriction. You need to consider the following:

- The user must to start DB2 server by using an existing USERID instead of the default "Local System Account".
- To specify read and write permissions on the network drives, select Services in Administrative Tools under Control Panel. Next, in the Services window, select Properties from the Action menu. Select Log On tab and update "Log on as" to indicate a specific user who has the read and write permissions on the network drives.
- In Windows 2000 and Windows XP, you need to perform backup, restore, or load action by specifying the full qualified UNC name instead of the network share.
- Use the command:

  ```
  db2 backup db mydata to \\ MyFolder\test
  ```

  instead of

  ```
  net use H: \\ MyFolder\test
  db2 backup db mydata to H:
  ```

On Microsoft Windows Server 2003 R2 Datacenter Edition SP1 or above versions, the idsicrt tool fails to catalog instance node and exits

On a system with Microsoft Windows Server 2003 R2 Datacenter Edition SP1 or above versions on it, before creating a directory server instance set the DB2INSTPROF environment variable. For example:

```
SET DB2INSTPROF=<PATHNAME>
```

where, *PATHNAME*, specifies the absolute path of the directory where the profile of the instance node will be created.

Note: The length of the absolute path should not exceed 70 characters.

You can also download DB2 9.5 Fix Pack 2 that has the fix to the above problem.

On Microsoft Windows operating systems, if DB2 v9.5 is installed using the InstallShield GUI then the idsdbmaint tool might display error indicating DB2 diagnostic path could not be determined

On Microsoft Windows operating systems, if DB2 v9.5 is installed using the InstallShield GUI then the idsdbmaint tool is unable to fetch DB2 diagnostic path. This problem does not occur if DB2 is installed separately (without using the InstallShield GUI). If the idsdbmaint tool is not able to fetch DB2 diagnostic path then the index reorganization operation will not have the desired impact. This is because the idsdbmaint tool requires DB2 diagnostic path for the index reorganization and row compression features to function. The idsdbmaint tool might displays the following error message when this problem is encountered.

```
GLPDBA048E The diagpath for database instance 'dsrdbm01' could not be determined.
```

The following is observed when issuing the idsdbmaint command:

```
c:\>idsdbmaint -I dsrdbm01 -r
GLPWRP123I The program 'C:\Program Files\IBM\LDAP\V.2\sbin\32\dbmaint.exe'
 is used with the following arguments '-I dsrdbm01 -r'.
GLPSRV200I Initializing primary database and its connections.
GLPDBA037I Row compression task will be performed.
GLPDBA048E The diagpath for database instance 'dsrdbm01' could not be determined.
```

Workaround

To resolve the issue of DB2 diagnostic path when installing DB2 v9.5 using the InstallShield GUI on Windows operating system, issue the following commands:

```
db2start
db2 connect to <instance name>
db2 update dbm cfg using DIAGPATH C:\<instancename>
db2 terminate
db2stop
```

On Microsoft Windows Server 2008, the Web Administration Tool and White Pages might not get launched properly with Internet Explorer

On Windows Server 2008, it has been observed that the Web Administration Tool panels and the White Pages configuration might not work properly with Internet Explorer unless the URLs of the Web Administration Tool and White Pages are part of Trusted sites of the browser.

To resolve this issue, users must make the URLs of the Web Administration Tool and White Pages part of Trusted sites. To do this, perform the following steps:

1. Launch Internet Explorer.
2. On the Tools menu, click Internet Options.
3. On the Internet Options dialog box, select the Security tab.
4. On the Security tab, select Trusted sites and then click the Sites button.
5. In the Add this Web site to the zone box, enter the URL for the Web Administration Tool and then click the Add button.
6. In the Add this Web site to the zone box, enter the URL for White Pages and then click the Add button.
7. Click OK.

Launch the Web Administration Tool and White Pages using the Internet Explorer.

Tivoli Directory Server clients run on Windows PowerShell might not function as expected

Windows PowerShell is not a supported shell for Tivoli Directory Server. Therefore, when Tivoli Directory Server clients are run on PowerShell, they might not function as expected. Tivoli Directory Server clients function properly when run on command prompt.

For example, when the following command is run on command prompt, it returns root DSE search results. However, when run on PowerShell, it displays the command usage.

```
idsldapsearch.cmd -p 389 -D cn=root -w root -s base -b "" objectclass=*
```

On Windows-based operating systems, the log management tool fails to start from Tivoli Directory Server V6.3 Web Administration Tool

On Windows-based systems, the recommended way to start Tivoli Directory Server and Administration server is through services. However, when the processes are started through services regardless of what the login credentials of the user is, the processes will always run using the SYSTEM user's privileges.

For example, if a user creates Tivoli Directory Server instance, myinst1, and log in as the myinst1 user and then starts the server and administration server using the services then the processes will run with the SYSTEM user's privileges. In this scenario, if the log management tool, idslogmgmt, will not work if it is started or stopped from the Web Administration Tool. This is because log management requires that the Web Administration Tool be started or stopped using the instance owner's credentials. Since the start/stop log management request from the Web Administration Tool goes to the administration server, which is running as the SYSTEM user and therefore starting or stopping the log management tool through the administration server will fail.

If the administration server is started from the command line using the credentials of myinst1, then the process would run with the myinst1's privileges. Therefore, starting or stopping the log management tool using the Web Administration Tool will work.

Uninstalling IBM Tivoli Directory Server client packages from Microsoft Windows 7

If a user attempts to uninstall Tivoli Directory Server client packages from a Windows 7 system as a non-privileged user, then the user might get the error message "Access is denied" . In such case, the user must uninstall the packages as an administrator by selecting the "Run as administrator" option from the User Account Control shield on command button or link. The Run as administrator option will prompt for administrator user name and password.

To know more about User Account Control on Windows, see http://msdn.microsoft.com/en-us/library/aa511445.aspx.

On 64-bit Windows Server 2008, Tivoli Directory Server V6.3 instance fails when stressed with large number of schema updates

When schema updates are made continuously for prolonged period of time on Tivoli Directory Server V6.3 instance that is running on 64-bit Windows Server 2008 operating system, then the directory server instance fails. This has been observed only on 64-bit Windows Server 2008 systems.

For Solaris only

On Solaris 10 with zones configured, removal of server components might fail if configured on a different zone

When installing IBM Tivoli Directory Server 6.2 on Solaris big-zone, all dependent client and server install packages must also be installed and run from big-zone.

The propagated installation of clients and srvbase from Solaris global-zone, and the installation of server components on big-zone does function. However, if the installed Tivoli Directory Server components (clients and srvbase) are uninstalled from global-zone, it will also result in the removal of these components from big-zone and small-zone, after which, the removal of server component from big-zone might fail.

Using the DB2 utility, db2osconf, on Solaris systems

The DB2 utility, db2osconf, makes recommendations for kernel parameter values based on the size of a system. This command is currently available only for DB2 on Solaris SPARC systems with 64-bit instances. If zones are configured on a Solaris SPARC system, then the db2osconf utility is available only in global zone. In the case of Solaris non-global zones and Solaris x86-64 systems, the db2osconf utility is not available instead, you can use the projmod command to set the values for kernel parameters, such as limits for shared memory, semaphore ids, and total shared memory.

On Solaris SPARC global zone, use db2osconf to obtain recommended values for kernel parameters. An example of the db2osconf command and its output is as follows:

```
#db2osconf
set msgsys:msginfo_msgmni = 6144
set semsys:seminfo_semmni = 7168
set shmsys:shminfo_shmmax = 9578697523
set shmsys:shminfo_shmmni = 7168

Total kernel space for IPC:
0.98MB (shm) + 1.71MB (sem) + 2.08MB (msg) == 4.77MB (total)
```

An example of the projmod command on Solaris SPARC (using values generated by db2osconf) is as follows:

```
projmod -s -K "project.max-shm-memory=(privileged,9578697523,deny)" user.db2inst1
projmod -s -K "project.max-shm-ids=(privileged,7168,deny)" user.db2inst1
projmod -s -K "project.max-msg-ids=(privileged,6144,deny)" user.db2inst1
projmod -s -K "project.max-sem-ids=(privileged,7168,deny)" user.db2inst1
```

An example of the projmod command on Solaris x86-64 (use values suitable for your environment) is as follows:

```
projmod -a -K "project.max-shm-ids=(priv,4k,deny)" user.db2inst1
projmod -a -K "project.max-sem-ids=(priv,4k,deny)" user.db2inst1
projmod -a -K "project.max-shm-memory=(priv,4G,deny)" user.db2inst1
projmod -a -K "project.max-msg-ids=(priv,4k,deny)" user.db2inst1
```

The values of these limits should be set in accordance with the available system resources in your environment. For more information see Memory management and related concepts in the DB2 v9.x Information Center.

On Solaris system, the idsadsrun utility does not exit when error is encountered

On a system with Solaris 10 as the operating system, the idsadsrun utility fails to exit when errors are encountered. In such cases, the user need to use Ctrl+C to stop the process.

On Solaris-Opteron system, migration of a Tivoli Directory Server v6.1 instance with DB2 v9.1 to Tivoli Directory Server v6.3 with DB2 v9.7 fails

On an AMD Opteron platform with Solaris 10 operating system, migration of a Tivoli Directory Server v6.1 instance with DB2 v9.1 to Tivoli Directory Server v6.3 with DB2 v9.7 fails. When migration of the directory server is initiated using the idsimigr command, it gives the following error message.

```
GLPMIG041E The database name listed in the backed up configuration
           file cannot be found on the system.
```

This problem is observed when DB2 v9.1 associated is at fix pack level 7 or lower.

However, if you use DB2 v9.1 fix pack 8 or higher levels, migration of Tivoli Directory Server v6.1 instance to Tivoli Directory Server v6.3 is successfully performed.

Appendix A. Common Base Event (CBE) features

In an effort to create self-managing environment, IBM has taken initiative in introducing "Autonomic Computing". Autonomic computing is an open standard based architecture that allows systems to configure, heal, optimize, and protect itself. In order to determine the conditions of the different components of the system, it is necessary to standardize the format of the event data so that the system can resolve its current conditions.

To standardize the format of data for the problem determination architecture IBM introduced a common format for log and trace information called the Common Base Event (CBE) format. This format creates consistency across similar fields and improves the availability to correlate across multiple logs. CBE is based on a 3-tupple structured format, which includes:

- Component impacted by a situation, or the source
- Component observing a situation
- Situation data, the properties describing the situation including correlation information

The 3-tupple format makes it possible to write and deploy resource-independent management functions that can isolate a failing component.

In an effort to align IBM Tivoli Directory Server to autonomic computing space, it is a must to have the logs such as error log, server audit log, and so on, produced by the Tivoli Directory Server product to provide these logs in CBE format.

The IBM Common Auditing and Reporting Service (CARS) component leverages CBE, which is a common format for events proposed by IBM, and IBM Common Event Infrastructure (CEI) technologies to provide an audit infrastructure. The purpose of CBE is to facilitate effective intercommunication among disparate components within an enterprise. In order to effectively process audit data, the CARS component requires the audit data to be in the CBE format. CEI is an IBM strategic event infrastructure for submission, persistent storage, query, and subscription of the CBE events. The CARS component uses the CEI interfaces for submission of events. These events can be denoted as auditable by using configuration options at the CEI Server that stores them in a CEI XML Event store that meets the auditing requirements.

The CARS component allows staging of data from the CEI XML Event store into report tables. IBM products and customers can provide audit reports based on auditable events staged into report tables. The CARS component also supports managing the lifecycle of auditable events, which includes archive, restore, and audit reports on restored archives.

In Tivoli Directory Server, auditing capability is implemented using the Tivoli Directory Server audit plug-in. A user can implement the audit enhancements to write audited data to CBE format. An example is listed:

- The audit data could be read and transformed to CBE format by an external application such as, IBM Tivoli Directory Integrator, and then sent over to CARS using the CEI API or the CARS embeddable Java client.

To implement the example, the settings for this feature in the ibmslapd configuration file are retrieved. If the settings are specified the audit data files are

read periodically and converted into CBE format by the log management tool. Depending on the settings, the CBE formatted data could be written to a file, to a CEI server, or both. The data sent to a CEI server is stored in the CEI database and CARS will move the audit data into a CARS database. The data will then move into a staging area for CARS reports or into a database archive for long term storage.

CBE related scenarios

There are some special case scenarios that should be considered for the CBE feature.

Attribute related special case scenarios

Unspecified attribute settings

If a value is set for ibm-slapdLogEventFileSizeThreshold and the value of ibm-slapdLogEventFileMaxArchives is not specified either in the default entry or in the specific log entry then in such case archiving will occur but the number of archive files will be unlimited.

Attribute settings given in wrong format

- If the value provided for ibm-slapdLogEventFileSizeThreshold is in the wrong format then an error message is logged and no archiving will occur.
- If the value provided for ibm-slapdLogEventFileMaxArchives is in the wrong format then an error message is logged but archiving will occur and the number of archive files will be unlimited.
- If the value provide for ibm-slapdLogEventFileArchivePath is invalid then the archived file path is in the same directory as that of the original file's path.

CBE file and Log management actions related scenarios

Out of disk space

If the disk gets full the log management activity will fail this is indicated by displaying an error message on the standard output and if possible it is also logged in the log.

Archive path errors

- If the archived file cannot be written to the path specified than an error message is logged on the idslogmgmt log.
- If a file with the same name already exists in the mentioned archive path then an error message is logged in the idslogmgmt log and the archiving will fail. When the next log management occurs, for the operation to succeed the timestamp should be different.

Log archiving and CBE activity interference

When the Tivoli Directory Server log management tool is configured to send CBE data to a CEI server and log archiving is also enabled then there is a possibility that the archiving threshold is reached but the log data has not been sent to the CEI server. The reason for this could be that the CEI server is down or the transmission rate to CEI server is lesser than the original log write rate. In such

cases the log archiving is suspended in order to not lose data that data that should be sent to CEI server. The expected behaviors when this situation occurs with the log management settings are listed.

- When CBE formatted logs are enabled and the value of ibm-slapdLogEventFileMaxArchives is set to zero, then the CBE file that should have been deleted is kept and the file continues to grow.
- When CBE formatted logs are enabled, the value of ibm-slapdLogEventFileMaxArchives is set to a number greater that zero, and the set maximum number of CBE log files have been reached, then the CBE file that should have been deleted is kept and the number of CBE archives continues to grow.
- When CBE formatted logs are disabled and the value of ibm-slapdLogMaxArchives is set to zero, then the log file that should have been deleted is kept and the file continues to grow.
- When CBE formatted logs are disabled, the value of ibm-slapdLogMaxArchives is set to a number greater than zero, and the set maximum number of log files have been reached, then the oldest archived file that should have been deleted is kept and the number of archives continues to grow.

Log activity overlapping cycles

When a current cycle of log activities are running for a log and if the next cycle of log activities are triggered, the tool should not allow multiple cycles to overlap. The next cycle should only start after the completion of the first cycle. This prevents different log activity cycles from interfering and causing data loss.

Appendix B. Support information

This section describes the following options for obtaining support for IBM products:

- "Searching knowledge bases"
- "Obtaining fixes"
- "Contacting IBM Software Support" on page 136

Searching knowledge bases

If you have a problem with your IBM software, you want it resolved quickly. Begin by searching the available knowledge bases to determine whether the resolution to your problem is already documented.

Search the information center on your local system or network

IBM provides extensive documentation that can be installed on your local computer or on an intranet server. You can use the search function of this information center to query conceptual information, instructions for completing tasks, reference information, and support documents.

Search the Internet

If you cannot find an answer to your question in the information center, search the Internet for the latest, most complete information that might help you resolve your problem. To search multiple Internet resources for your product, expand the product folder in the navigation frame to the left and select **Web search**. From this topic, you can search a variety of resources including:

- IBM technotes
- IBM downloads
- IBM Redbooks®
- IBM developerWorks
- Forums and newsgroups
- Google

Obtaining fixes

A product fix might be available to resolve your problem. You can determine what fixes are available for your IBM software product by checking the product support Web site:

1. Go to the IBM Software Support Web site (http://www.ibm.com/software/support).
2. Under **Products A - Z**, select your product name. This opens a product-specific support site.
3. Under **Self help**, follow the link to **All Updates**, where you will find a list of fixes, fix packs, and other service updates for your product. For tips on refining your search, click **Search tips**.
4. Click the name of a fix to read the description and optionally download the fix.

To receive weekly e-mail notifications about fixes and other news about IBM products, follow these steps:

1. From the support page for any IBM product, click **My support** in the upper-right corner of the page.
2. If you have already registered, skip to the next step. If you have not registered, click register in the upper-right corner of the support page to establish your user ID and password.
3. Sign in to **My support**.
4. On the My support page, click **Edit profiles** in the left navigation pane, and scroll to **Select Mail Preferences**. Select a product family and check the appropriate boxes for the type of information you want.
5. Click **Submit**.
6. For e-mail notification for other products, repeat Steps 4 and 5.

For more information about types of fixes, see the *Software Support Handbook* (http://techsupport.services.ibm.com/guides/handbook.html).

Contacting IBM Software Support

IBM Software Support provides assistance with product defects.

Before contacting IBM Software Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have:

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus®, and Rational® products, as well as DB2 and WebSphere products that run on Windows, AIX, Linux, Solaris, and HP-UX operating systems), enroll in Passport Advantage® in one of the following ways:
  - **Online**: Go to the Passport Advantage Web page (http://www.lotus.com/services/passport.nsf/WebDocs/ Passport_Advantage_Home) and click **How to Enroll**
  - **By phone**: For the phone number to call in your country, go to the IBM Software Support Web site (http://techsupport.services.ibm.com/guides/contacts.html) and click the name of your geographic region.
- For IBM eServer™ software products (including, but not limited to, DB2 and WebSphere products that run in System z, System p, and System i environments), you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for eServer software products, go to the IBM Technical Support Advantage Web page (http://www.ibm.com/servers/eserver/techsupport.html).

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States or, from other countries, go to the contacts page of the IBM Software Support Handbook on the Web (http://techsupport.services.ibm.com/guides/contacts.html) and click the name of your geographic region for phone numbers of people who provide support for your location.

Follow the steps in this topic to contact IBM Software Support:

1. Determine the business impact of your problem.
2. Describe your problem and gather background information.

3. Submit your problem to IBM Software Support.

Determine the business impact of your problem

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you need to understand and assess the business impact of the problem you are reporting. Use the following criteria:

Table 7. Severity level and their description

Severity	Business impact
Severity 1	**Critical** business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	**Significant** business impact: The program is usable but is severely limited.
Severity 3	**Some** business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	**Minimal** business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.

Describe your problem and gather background information

When explaining a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently. To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? (For example, hardware, operating system, networking software, and so on.)
- Are you currently using a workaround for this problem? If so, please be prepared to explain it when you report the problem.

Submit your problem to IBM Software Support

You can submit your problem in one of two ways:

- **Online**: Go to the "Submit and track problems" page on the IBM Software Support site (http://www.ibm.com/software/support/probsub.html). Enter your information into the appropriate problem submission tool.
- **By phone**: For the phone number to call in your country, go to the contacts page of the IBM Software Support Handbook on the Web (techsupport.services.ibm.com/guides/contacts.html) and click the name of your geographic region.

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround for you to implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the IBM product support Web pages daily, so that other users who experience the same problem can benefit from the same resolutions.

For more information about problem resolution, see Searching knowledge bases and Obtaining fixes.

Appendix C. Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript®, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine™ and Cell/B.E. are trademarks of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

Intel, Intel logo, Intel Inside®, Intel Inside logo, Intel Centrino®, Intel Centrino logo, Celeron®, Intel Xeon®, Intel SpeedStep®, Itanium, and Pentium® are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

IT Infrastructure Library® is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

ITIL® is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Java 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, Windows NT®, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark in the United States and/or other countries licensed exclusively through X/Open Company Limited.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Index

T

U

V

W

Printed in USA

GC27-2752-00