

IBM Tivoli Directory Server



Problem Determination Guide

Version 6.0

IBM Tivoli Directory Server



Problem Determination Guide

Version 6.0

Note

Before using this information and the product it supports, read the general information under Appendix B, "Notices," on page 69.

First Edition (April 2005)

This edition applies to version 6, release 0, of IBM Tivoli Directory Server and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright International Business Machines Corporation 2005. All rights reserved.

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

Contents

| | |
|---|----------|
| Preface | v |
| Who should read this book | v |
| Publications | v |
| IBM Tivoli Directory Server library | v |
| Related publications | vi |
| Accessing publications online | vi |
| Ordering publications | vi |
| Accessibility | vii |
| Tivoli technical training | vii |
| Support information | vii |
| Conventions used in this book | vii |
| Typeface conventions | vii |
| Operating system differences | vii |

| | |
|---|----------|
| Chapter 1. Introduction to problem determination | 1 |
| IBM Tivoli Directory Server overview | 1 |
| Built-in troubleshooting features | 1 |
| Using the Messages Guide to resolve errors | 2 |
| Troubleshooting topics | 2 |

| | |
|---|----------|
| Chapter 2. Logging utilities | 3 |
|---|----------|

| | |
|---|----------|
| Chapter 3. Other diagnostic tools. | 7 |
| Generating core files | 7 |
| For Windows operating systems (Dr. Watson debugger) | 7 |
| For Linux operating systems | 7 |
| For AIX operating systems | 8 |
| For Solaris operating systems | 8 |
| For HP-UX operating systems | 8 |
| IBM Tivoli Directory Server Support Tool | 9 |
| Data collected by the Support Tool | 9 |
| Using the Support Tool | 11 |
| Location of Support Tool log and collected data | 12 |
| Server debug mode | 12 |

| | |
|---|-----------|
| Chapter 4. Troubleshooting installation and uninstallation | 15 |
| Product installation overview | 15 |
| Prerequisite software | 15 |
| Installation logs | 16 |
| Logs for the embedded version of WebSphere | |
| Application Server - Express | 16 |
| DB2 on Windows logs | 16 |
| DB2 on AIX, Linux, Solaris, and HP-UX logs | 17 |
| idslink log on AIX, Linux, and Solaris operating systems | 17 |
| GSKit logs on Windows operating systems | 17 |
| Troubleshooting | 17 |
| InstallShield GUI installation | 17 |
| Operating system utility installation | 20 |
| InstallShield GUI uninstallation | 20 |

| | |
|---|-----------|
| Chapter 5. Troubleshooting migration | 21 |
| Migration log files | 21 |
| Kerberos service name change | 21 |
| Database instance or database in configuration file but no longer on system | 21 |
| Format of backed-up schema files incorrect | 22 |
| ibm-slapdPlugin entry in configuration file changed | 22 |

| | |
|---|-----------|
| Chapter 6. Troubleshooting instance creation and configuration | 23 |
| Instance creation overview and common errors | 23 |
| Instance creation overview | 23 |
| Common instance creation errors | 24 |
| Configuration overview and common errors | 24 |
| Overview | 24 |
| Common errors | 25 |

| | |
|--|-----------|
| Chapter 7. Troubleshooting DB2. | 29 |
| DB2 license file expired | 29 |
| Database performance is poor | 30 |

| | |
|--|-----------|
| Chapter 8. Troubleshooting the Web Administration Tool and the application server | 31 |
| Troubleshooting the Web Administration Tool | 31 |
| Corruption of data entered in the Web Administration Tool | 31 |
| Migrating files when patching or migrating the Web Administration Tool | 32 |
| Additional login panels fail | 32 |
| idsldapmodify command puts Web Administration Tool into inconsistent state | 32 |
| Web Administration Tool loses connections on HP-UX | 33 |
| Web Administration Tool tabs, table headers, and static list boxes are displayed in incorrect language | 33 |
| Microsoft Internet Explorer browser problems | 34 |
| HTML special characters are not displayed correctly | 35 |
| Web Administration Tool requires IBM JDK on a Domino server | 35 |
| Web Administration Tool does not save templates created with an object class that has no attributes | 35 |
| Using CTRL+L to view links makes non-editable fields appear editable | 35 |
| Internet browser Back and Forward buttons not supported for Web Administration Tool | 35 |
| Logging on to the Web Administration Tool console on Internet Explorer | 35 |
| Difficulties encountered using the Web Administration GUI console on the Windows Server 2003 platform | 36 |

| | |
|---|----|
| Troubleshooting the embedded version of WebSphere Application Server - Express | 37 |
| Error when starting the embedded version of WebSphere Application Server - Express on AIX | 37 |

Chapter 9. Troubleshooting replication 39

| | |
|--|----|
| Replication overview | 39 |
| Diagnosing replication errors | 39 |
| Sample replication topology | 39 |
| Monitoring replication status using idsldapsearch | 40 |
| Viewing replication errors using the Web Administration Tool | 42 |
| Viewing replication errors using the idsldapsearch command | 43 |
| Lost and found log | 44 |
| Replication Troubleshooting | 44 |
| Replicated suffix must have ibm-replicationcontext object class | 45 |
| Verify that suffixes and replication agreements exist using idsldapsearch | 45 |
| Peer to peer replication returns error "No such object occurred for replica" | 45 |
| Replication returns error "Insufficient access" | 46 |
| Replication topology extended operation returns result code 80. | 46 |
| Replication command-line interface error (Windows systems only) | 47 |
| Entries in LDIF file are not replicated. | 48 |

Chapter 10. Troubleshooting performance 49

| | |
|--|----|
| Identifying performance problem areas | 49 |
| Audit log | 49 |
| idslslapd trace. | 49 |
| Adding memory after installation on Solaris systems | 49 |
| Setting the SLAPD_OCHANDLERS environment variable on Windows | 50 |
| DB2 rollbacks and isolation levels | 50 |
| Default value of LOGFILSIZ needs to be increased | 50 |

Chapter 11. Troubleshooting scenarios 53

| | |
|--|----|
| Server is not responding | 53 |
| Memory leak suspected | 53 |
| SSL communications returning errors. | 54 |

Chapter 12. Known limitations and general troubleshooting 55

| | |
|--|----|
| Known limitations | 55 |
| Command line utilities allow an option to be entered more than once | 55 |
| Some types of invalid data entered on command line utilities do not produce an error | 55 |
| No locking mechanism for conflicting commands on the same directory instance | 55 |

| | |
|---|----|
| Using CTRL+C with a client command that takes passwords can cause an error | 55 |
| Unable to drop database | 55 |
| On AIX, Linux, Solaris, and HP-UX operating systems, error messages appear after command prompt when starting the administration daemon | 56 |
| General troubleshooting | 56 |
| Instance owner unable to access core file for core that occurred during server initialization | 56 |
| Key label in .kdb file and ibmslapd.conf file do not match. | 56 |
| GSKit certificate error | 56 |
| Server instance fails to start because of incorrect file permissions | 57 |
| Server instance fails to start because localhost hostname is set incorrectly | 57 |
| Server instance cannot be started except by instance owner | 57 |
| Error opening slapd.cat file on Windows systems | 57 |
| DSML file client produces error. | 58 |
| Nondefault log files need valid path | 58 |
| Null searches retrieve entries of deleted suffixes | 58 |
| Fixing an "SQL0964C Transaction log for database is full" error | 58 |
| idslldapsearch command with -h option gives error with the DIGEST-MD5 mechanism. | 59 |
| After enabling language tags, do not disable language tags. | 59 |
| Create the key database certificate before setting up SSL. | 59 |
| idslbulpload appears to hang during parsing phase | 60 |
| Platform specific problems | 60 |
| For AIX only | 60 |
| For Windows 2000, Windows Server 2003 Enterprise and Windows XP client only | 62 |

Appendix A. Support information 65

| | |
|---|----|
| Searching knowledge bases | 65 |
| Search the information center on your local system or network | 65 |
| Search the Internet | 65 |
| Obtaining fixes | 65 |
| Contacting IBM Software Support | 66 |
| Determine the business impact of your problem | 67 |
| Describe your problem and gather background information | 67 |
| Submit your problem to IBM Software Support | 67 |

Appendix B. Notices 69

| | |
|----------------------|----|
| Trademarks | 70 |
|----------------------|----|

Index 73

Preface

This book provides information about possible problems and corrective actions that can be tried before contacting IBM® Software Support. It also includes information about tools you can use for problem determination with IBM Tivoli® Directory Server 6.0.

Who should read 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

Read the descriptions of the IBM Tivoli Directory Server library, the prerequisite publications, and the related publications to determine which publications you might find helpful. After you determine the publications you need, see “Accessing publications online” on page vi for information about accessing publications online.

IBM Tivoli Directory Server library

The publications in the IBM Tivoli Directory Server library are:

IBM Tivoli Directory Server Version 6.0 Release Notes

Contains information about the new features in the IBM Tivoli Directory Server Version 6.0 release, as well as last-minute updates.

IBM Tivoli Directory Server Version 6.0 Installation and Configuration Guide

Contains complete information for installing, configuring, and uninstalling IBM Tivoli Directory Server. Includes information about migrating from a previous version of IBM Tivoli Directory Server or SecureWay® Directory.

IBM Tivoli Directory Server Version 6.0 Administration Guide

Contains instructions for performing administrator tasks through the Web Administration Tool and the command line.

IBM Tivoli Directory Server Version 6.0 Performance Tuning Guide

Contains information about tuning your server for better performance.

IBM Tivoli Directory Server Version 6.0 Server Plug-ins Reference

Contains information about writing server plug-ins.

IBM Tivoli Directory Server Version 6.0 C-Client SDK Programming Reference

Contains information about writing Lightweight Directory Access Protocol (LDAP) client applications.

IBM Tivoli Directory Server Version 6.0 Problem Determination Guide

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

IBM Tivoli Directory Server Version 6.0 Messages Guide

Contains a list of all informational, warning, and error messages associated with IBM Tivoli Directory Server 6.0.

Related publications

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

- IBM Tivoli Directory Server Version 6.0 uses the JNDI client from Sun Microsystems. For information about the JNDI client, refer to the *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>.
- 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/tdprodlist.html>
- The *Tivoli Software Glossary* includes definitions for many of the technical terms related to Tivoli software. The *Tivoli Software Glossary* is available at <http://publib.boulder.ibm.com/tividd/glossary/tivoliglossarymst.htm>
- The DB2® documentation library is located at <http://www.ibm.com/software/data/db2/library/>.

Accessing publications online

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

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

Product publications include release notes, installation guides, user's guides, administrator's guides, and developer's references.

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 the following Web site:

<http://www.elink.ibm.com/public/applications/publications/cgi-bin/pbi.cgi>

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, see the following Web site for a list of telephone numbers:

<http://www.ibm.com/software/tivoli/order-lit/>

Accessibility

Accessibility features help a user who has 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 also can use the keyboard instead of the mouse to operate all features of the graphical user interface.

Tivoli technical training

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

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:

- Searching knowledge bases: You can search across a large collection of known problems and workarounds, Technotes, and other information.
- Obtaining fixes: You can locate the latest fixes that are already available for your product.
- Contacting IBM Software Support: If you still cannot solve your problem, and you need to work with someone from IBM, you can use a variety of ways to contact IBM Software Support.

For more information about these three ways of resolving problems, see Appendix A, “Support information,” on page 65.

Conventions used in this book

This reference uses several conventions for special terms and actions and for operating system-dependent commands and paths.

Typeface conventions

The following typeface conventions are used in this reference:

Bold Lowercase commands or mixed case commands that are difficult to distinguish from surrounding text, keywords, parameters, options, names of Java classes, and objects are in **bold**.

Italic Variables, titles of publications, and special words or phrases that are emphasized are in *italic*.

<*Italic*>

Variables are set off with < > and are in <*italic*>.

Monospace

Code examples, command lines, screen output, file and directory names that are difficult to distinguish from surrounding text, system messages, text that the user must type, and values for arguments or command options are in monospace.

Operating system differences

This book uses the UNIX[®] convention for specifying environment variables and for directory notation. When you are using the Windows[®] command line, replace *\$variable* with *%variable%* for environment variables and replace each forward

slash (/) with a backslash (\) in directory paths. 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 6.0.

IBM Tivoli Directory Server overview

IBM Tivoli Directory Server is the IBM implementation of Lightweight Directory Access Protocol (LDAP) for supported Windows, AIX[®], xSeries[®] Linux[®], iSeries[™] Linux, pSeries[®] Linux, zSeries[®] Linux, Solaris, and HP-UX 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 7 for more information about core files and for instructions for enabling core file generation.

Support Tool (idssupport)

The Support Tool collects relevant data about the directory server and the system (such as logs, directory listing, schema files, and core files) and packages this information in a compressed file archive that you can send to IBM Software Support for help in diagnosing your problem. See "IBM Tivoli Directory Server Support Tool" on page 9 for more information.

To use the Support Tool, you must have the proxy server and IBM Tivoli Directory Integrator 6.0 with Fix Pack 1 installed. The *IBM Tivoli Directory Server Version 6.0 Installation and Configuration Guide* contains instructions for installing the proxy server.

If you acquire IBM Tivoli Directory Server through Passport Advantage[®], it includes a copy of IBM Tivoli Directory Integrator 6.0 with Fix Pack 1 for limited use with IBM Tivoli Directory Server 6.0. (Installation instructions are included with the package.)

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 3 for more information.

Audit logs

Audit logs record suspicious patterns of activity in order to detect security violations. If security is violated, the 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 3 for more information.

Using the Messages Guide to resolve errors

The *IBM Tivoli Directory Server Version 6.0 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 6.0 Messages Guide*.

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

```
Sep 13 14:31:04 2004 GLPL2D014E Suffix entry has not been created for entry
cn=Robert Dean, ou=In Flight Systems, ou=Austin, o=IBM, c=US.
```

You can search for "GLPL2D014E" in the *IBM Tivoli Directory Server Version 6.0 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 6.0 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 15 for more information.
- Migration: See Chapter 5, "Troubleshooting migration," on page 21 for more information.
- Instance Creation: See Chapter 6, "Troubleshooting instance creation and configuration," on page 23 for more information.
- Configuration: See Chapter 6, "Troubleshooting instance creation and configuration," on page 23 for more information.
- DB2: See Chapter 7, "Troubleshooting DB2," on page 29 for more information.
- Web Administration Tool and application server: See Chapter 8, "Troubleshooting the Web Administration Tool and the application server," on page 31 for more information.
- Replication: See Chapter 9, "Troubleshooting replication," on page 39 for more information.
- Performance: See Chapter 10, "Troubleshooting performance," on page 49 for more information.
- Scenarios: See Chapter 11, "Troubleshooting scenarios," on page 53 for more information.
- General troubleshooting: See Chapter 12, "Known limitations and general troubleshooting," on page 55 for more information.

Chapter 2. Logging utilities

IBM Tivoli Directory Server Version 6.0 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 6.0 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 6.0 Administration Guide* for information. The IBM Tivoli Directory Server logs are:

Administration daemon error log (`ibmdiradm.log`)

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

A sample of the log looks like this:

```
Oct 18 09:56:29 2004 GLPCOM003I Non-SSL port initialized to 3538.
Oct 18 09:56:29 2004 GLPADM028I Admin Daemon audit logging is started.
Oct 18 09:56:34 2004 GLPADM004I IBM Tivoli Directory (SSL), Version 6.0
ibmdiradm started
```

Administration daemon audit log (`adminaudit.log`)

Administration daemon audit logging is used to improve the security of the administration daemon. 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.

A sample of the log looks like this:

```
2004-11-15-19:59:17.130-06:00GLPADM028I Admin Daemon audit logging is started.
AuditV3--2004-11-16-22:04:50.93986-06:00--V3 Bind--bindDN: CN=ROOT--client:
127.0.0.1:3665--connectionID: 0--received:
2004-11-16-22:04:50.93986-06:00--Success
AuditV3--2004-11-16-22:04:50.93986-06:00--V3 Search--bindDN: CN=ROOT--client:
127.0.0.1:3665--connectionID: 0--received:
2004-11-16-22:04:50.93986-06:00--Success
AuditV3--2004-11-16-22:04:50.93986-06:00--V3 Unbind--bindDN: CN=ROOT--client:
127.0.0.1:3665--connectionID: 0--received:
2004-11-16-22:04:50.93986-06:00--Success
AuditV3--2004-11-16-22:08:09.94185-06:00--V3 Bind--bindDN: CN=OT--client:
127.0.0.1:3678--connectionID: 1--received:
2004-11-16-22:08:09.94185-06:00--Invalid credentials
AuditV3--2004-11-16-22:08:09.94185-06:00--V3 Unbind--bindDN: --client:
127.0.0.1:3678--connectionID: 1--received:
2004-11-16-22:08:09.94185-06:00--Success
```

Audit log (`audit.log`)

Audit logging is used to improve the security of the directory server. The system administrator and administrative group members can use the activities 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. 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 following is a sample of the audit log:

```
2004-11-16-17:38:15.484-06:00--GLPSRV023I Audit logging started. The audit
configuration options are:
  ibm-slapdLog = C:\idsslapd-ldaptest\logs\audit.log,ibm-auditVersion =
    ,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.
2004-11-16-17:38:15.656-06:00--GLPSRV009I IBM Tivoli Directory (SSL),
Version 6.0 Server started.
AuditV3--2004-11-16-17:39:28.468-06:00--V3 anonymous Search--bindDN:
  <*CN=NULLDN*>--client:
  127.0.0.1:3792--connectionID: 1--received: 2004-11-16-17:39:28.453-06:00--
  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=ibm,c=us
objectclass:person
cncaa
sn:abc
```

The following bulkload error log resulted:

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

Configuration tools log (idstools.log)

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

The following is a sample of the configuration tools log:

```
Nov 09 16:41:02 2004 GLDPW009I Setting the directory server administrator DN.
Nov 09 16:41:02 2004 GLDPW010I Set the directory server administrator DN.
Nov 09 16:41:02 2004 GLDPW006I Setting the directory server administrator
password.
Nov 09 16:41:11 2004 GLDPW007I Set the directory server administrator
password.
Nov 09 16:41:17 2004 GLPCDB035I Adding database 'ldaptest' to directory server
instance: 'ldaptest'.
Nov 09 16:41:18 2004 GLPCTL017I Cataloging database instance node: 'ldaptest'.
Nov 09 16:41:19 2004 GLPCTL018I Cataloged database instance node: 'ldaptest'.
Nov 09 16:41:19 2004 GLPCTL008I Starting database manager for database
instance: 'ldaptest'.
Nov 09 16:41:22 2004 GLPCTL009I Started database manager for database
instance: 'ldaptest'.
Nov 09 16:41:22 2004 GLPCTL026I Creating database: 'ldaptest'.
Nov 09 16:43:11 2004 GLPCTL027I Created database: 'ldaptest'.
Nov 09 16:43:11 2004 GLPCTL034I Updating the database: 'ldaptest'
Nov 09 16:43:19 2004 GLPCTL035I Updated the database: 'ldaptest'
Nov 09 16:43:19 2004 GLPCTL020I Updating the database manager: 'ldaptest'.
Nov 09 16:43:22 2004 GLPCTL021I Updated the database manager: 'ldaptest'.
Nov 09 16:43:23 2004 GLPCTL023I Enabling multi-page file allocation:
'ldaptest'
Nov 09 16:43:37 2004 GLPCTL024I Enabled multi-page file allocation:
'ldaptest'
Nov 09 16:43:38 2004 GLPCDB005I Configuring database 'ldaptest' for
directory server instance:
'ldaptest'.
Nov 09 16:43:39 2004 GLPCDB006I Configured database 'ldaptest' for
directory server instance: 'ldaptest'.
Nov 09 16:43:39 2004 GLPCDB003I Added database 'ldaptest' to directory
server instance: 'ldaptest'.
```

DB2 error 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 error log:

```
2004-09-13-19:18:29.native retcode = -1031; state = "58031";
message = "SQL1031N
The database directory cannot be found on the indicated file system.

SQLSTATE=58031

"
2004-09-13-19:18:29.native retcode = -1018; state = "E8";
message = "SQL1018N
The node name "idsinode" specified in the CATALOG NODE command
already exists.

"
2004-09-13-19:18:30.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: 20041106211242.000000Z
#Entry modifyTimestamp: 20041030202533.000000Z
#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 error log (`ibmslapd.log`)

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

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

```
Sep 13 14:31:04 2004 GLPL2D014E Suffix entry has not been created for entry
  cn=Robert Dean, ou=In Flight Systems, ou=Austin, o=IBM, c=US.
Sep 13 14:31:04 2004 GLPRDB002W ldif2db: 0 entries have been successfully added
  out of 50 attempted.
Sep 13 14:39:41 2004 GLPCOM024I The extended Operation plugin is successfully
  loaded from libevent.dll.
Sep 13 14:39:41 2004 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 15.

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
```

For HP-UX 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
```

IBM Tivoli Directory Server Support Tool

The Support Tool collects relevant data such as logs, directory listings, schema files, and core files, about the directory server. The Support Tool then packages the information into a compressed file archive that you can send to IBM Software Support for help in diagnosing your problem. This section describes the information collected by the Support Tool, and contains instructions for generating the idssupport file.

To use the Support Tool, you must have the proxy server and IBM Tivoli Directory Integrator 6.0 with Fix Pack 1 installed. The *IBM Tivoli Directory Server Version 6.0 Installation and Configuration Guide* contains instructions for installing the proxy server.

If you acquire IBM Tivoli Directory Server through Passport Advantage, it includes a copy of IBM Tivoli Directory Integrator 6.0 with Fix Pack 1 for limited use with IBM Tivoli Directory Server 6.0. (Installation instructions are included with the package.)

Data collected by the Support Tool

The Support Tool collects the following data:

Configuration commands

The following configuration command data is collected:

- The configuration tools log (idstools.log). This log contains information about the following configuration commands:
 - idscfgdb
 - idsucfgdb
 - idscfgchglog
 - idsucfgchglog
 - idscfgsuf
 - idsucfgsuf
 - idscfgsch
 - idsucfgsch
 - idsdnpw
 - idsxcfg
- Bulkload data. The following bulkload data is collected:
 - Bulkload log (bulkload.log), including archived bulkload logs
 - db2load.log
 - Recovery information
- Idif2db and db2ldif data

ibmdiradm

The following ibmdiradm information is collected:

- Administration daemon error log (ibmdiradm.log)
- Archived administration daemon error logs
- Administration daemon audit log (adminAudit.log)
- Archived administration daemon audit logs

ibmslapd

The following ibmslapd information is collected:

- The ibmslapd version

- Server error log (ibmslapd.log)
- Archived server error logs.
- Audit log (audit.log)
- Archived audit logs
- DB2 log (db2cli.log)
- Archived DB2 logs
- Trace log (traceibmslapd.log)
- ibmslapd.conf
- Lost and found log (lostandfound.log)
- Schema files
- Core files. Core file generation must be turned before the Support Tool can collect core files
- Stack traces of previous and latest core files

Note: To enable stack tracing on Linux operating systems, the gdb debugger is required. To enable stack tracing on HP operating systems, the adb debugger is required.

Administrative commands

The following instance administration information is collected:

- idsadm.log. This log contains data about the following commands:
 - idsilist
 - idsicrt
 - idsidrop
 - idssetip
 - idssetport
 - idsimigr
 - idxinst
 - idslink
 - idsadmdb2
- idsadmdb2.log
- idsadmdb2cmds.log
- idsinstances.ldif
- idslink.log (non-Windows operating systems only)
- idslink.preview (non-Windows operating systems only)

Serviceability commands

The following serviceability information is collected:

- idssupport.log
- idslogmgmt.log

DB2 data

The following DB2 information is collected:

- db2diag.log
- db2support
- DB2 level
- DB2 license

Operating system data

The following operating system information is collected:

- Platform name
- Kernel mode
- Home directory permission of the DB2 user
- Directory listing of the installation directory
- Directory listing of the directory server instance owner home directory
- Installed package listing of all installed products on UNIX and Linux platforms

System hardware data

The following system hardware information is collected:

- Amount of memory installed
- Hard disk information
- Amount of disk space in the /var and /tmp directories
- Processor information

GSKit data

The following GSKit information is collected:

- GSKit version

Embedded version of WebSphere® Application Server - Express and WebSphere Application Server data

The following embedded version of WebSphere Application Server - Express or WebSphere Application Server information is collected:

- the embedded version of WebSphere Application Server - Express version or WebSphere Application Server version

Installation data

The following installation information is collected:

- InstallShield GUI installation and uninstallation logs: ldapinst.log, ldapuninst.log and ldaplp_inst.log

Using the Support Tool

To use this tool you must have a server and IBM Tivoli Directory Integrator 6.0 with Fix Pack 1 installed. The ibmslapd.conf file must also be located in its default directory. If these conditions are not met, the Support Tool will not start.

Note: You might receive an error if the directory where the idssupport zip file resides runs out of disk space. To avoid this error, be sure that you have plenty of disk space available in the directory where the idssupport file resides. Core files can cause the file to be rather large, so if you plan to include core files in the file generated by the Support Tool, be sure that you have enough disk space.

To enable the Support Tool, run the following command at the command prompt:

Note: Use only absolute paths with the **-I** parameter.

```
<install_home>\sbin\idssupport [-I <directory_server_instance_name>]
[-l <data_collection_path>] [-n core] [-q] [-v] [-?]
```

Command options:

All options are case sensitive and optional.

-I <directory_server_instance_name>

Specifies the directory server instance to collect data from. If this

name is not specified then data is collected from all directory server instances. If no directory server instances exist then only global data is collected.

- l** *<data_collection_path>*
Specifies the location to save the collected data. */<timestamp>* is appended to this path. The data collection path specified must be a full path and not a relative path. This value overrides the default value. For the default location, see “Location of Support Tool log and collected data.”
- n** *core*
The latest IBM Tivoli Directory Server core file will not be collected unless this option and value are specified.
- q**
Prevents log messages from displaying on the screen. This is an optional parameter.
- v**
Displays the version number of the Support Tool.
- ?**
Displays the command usage.

Location of Support Tool log and collected data

You can send the zip file containing the captured information to IBM Software Support to help you diagnose your problem. If the **-l** parameter is not specified, the zip file is saved in one of the following default locations.

- On Windows operating systems:
<install_home>\var\idssupport\<timestamp>
- On AIX, Linux, Solaris, and HP-UX operating systems:
/var/idsldap/V6.0/idssupport/<timestamp>/

where *timestamp* is the time at which the file was generated and *install_home* is the directory where IBM Tivoli Directory Server is installed.

The Support Tool log is saved in one of the following locations by default:

- On Windows operating systems:
<install_home>\var\idssupport\<timestamp>\idssupport.log
- On AIX, Linux, Solaris, and HP-UX operating systems:
/var/idsldap/V6.0/idssupport/<timestamp>/idssupport.log

where *timestamp* is the time at which file was generated and *install_home* is the directory where you installed IBM Tivoli Directory Server.

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 bitmask 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 6.0 Administration Guide* or the *IBM Tivoli Directory Server Version 6.0 C-Client SDK Programming Reference* for information about the **idsldaptrace** utility.

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
 - Proxy server
 - Full server
- Web Administration Tool
- The embedded version of WebSphere Application Server - Express
- IBM DB2 Enterprise 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. InstallShield GUI installation is not available for HP-UX operating systems.

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 6.0 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 is, in `c:\Program Files\IBM\LDAP\V6.0\var`.
- On AIX, Linux, Solaris, and HP-UX systems, the `ldapinst.log` file is in `/var/idsldap/V6.0`.

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, Solaris, and HP-UX 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, if the full server is being installed but the DB2 installation fails, the full server cannot be installed.

Installation logs

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

Logs for the embedded version of WebSphere Application Server - Express

Logs used by the InstallShield GUI when installing the embedded version of WebSphere Application Server - Express 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*

On AIX, Linux, and Solaris platforms

- */var/idsldap/V6.0/installApp.log*
- */var/idsldap/V6.0/installAppErr.log*
- */var/idsldap/V6.0/configApp.log*
- */var/idsldap/V6.0/configAppErr.log*
- */var/idsldap/V6.0/migrateApp.log*
- */var/idsldap/V6.0/migrateAppErr.log*

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

DB2 on Windows logs

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

When installing

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

When uninstalling

- *Documents and Settings\<userID>\Local Settings\temp\DB2remove.log*
- *Documents and Settings\<userID>\Local Settings\temp\db2uninst.log*

- Documents and Settings\<<userID>\Local Settings\temp\db2uninsterr.log
- Documents and Settings\<<userID>\Local Settings\temp\DB2UninstTrc.log

DB2 on AIX, Linux, Solaris, and HP-UX logs

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

When installing using the InstallShield GUI

- /var/idsldap/V6.0/db2licmAddstdout.log
- /var/idsldap/V6.0/db2licmAddstderr.log
- /var/idsldap/V6.0/db2licmstdout
- /var/idsldap/V6.0/db2licmstderr.log

When installing using the db2_install utility

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

idslink log on AIX, Linux, and Solaris operating systems

The **idslink** script runs automatically during InstallShield GUI and operating system utility installation of the client, the proxy server, and full server. The **idslink.log** and **idslink.preview** files are located in the `/var/idsldap/V6.0/` 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

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 recommended free disk space specified in the *IBM Tivoli Directory Server Version 6.0 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 6.0 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 15 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 6.0 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\V6.0`.
4. Use **regedit** to remove the LDAP entry in the registry:
`HKEY_LOCAL_MACHINE\SOFTWARE\IBM\IDSLDAP\6.0`
5. Remove the following paths from the PATH environment variable:
 - `<install_home>\sbin`
 - `<install_home>\bin`
 - `<install_home>\lib`

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 15 for more information about the `ldapinst.log` file.

Note: AIX operating system installation generates an additional log named `/opt/IBM/ldap/V6.0/installp_isje.log`. You must review this log to determine if there were failures in the **installp** commands issued by the InstallShield GUI.

2. Uninstall IBM Tivoli Directory Server using the InstallShield GUI. See the *IBM Tivoli Directory Server Version 6.0 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 ids1`
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 6.0 Installation and Configuration Guide* for information about package names for IBM Tivoli Directory Server.
5. Remove the `/opt/IBM/ldap/V6.0` 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 15 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 6.0 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 ids1
```

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 6.0 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/V6.0` 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 15 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 6.0 Installation and Configuration Guide* for more information about uninstalling IBM Tivoli Directory Server.
3. Type the following at a command prompt:

```
pkginfo | grep -i ids1
```
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 6.0 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/V6.0` 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**, the 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 idldap user and the idldap group and then do one of the following:
 - Recreate the idldap user and group as described in the *IBM Tivoli Directory Server Version 6.0 Installation and Configuration Guide*.
 - Do not recreate the idldap user and group, but let the proxy server installation recreate them (when you do step 3.)
- 3. Reinstall the proxy server feature.
- 4. Reinstall the full server feature if it was not installed and you want the full server.

This problem is most likely to occur if you had an idldap user on the system before installing IBM Tivoli Directory Server, but it did not meet the requirements for this user.

Note: You do not need to define the idldap user and group before installation. If they do not exist, the proxy server installation creates them correctly for you.

Operating system utility uninstallation after InstallShield GUI installation

The *IBM Tivoli Directory Server Version 6.0 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 17.

Operating system utility installation

The following items relate to operating system utility installation.

DB2 installation fails

To install DB2, using the `db2_install` utility on AIX, Linux, Solaris, and HP-UX systems, the locale setting must be `en_US` or `C` (that is, you must call `export LANG=C` and `export LC_ALL=C`); otherwise `db2_install` fails.

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 6.0 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, Solaris, and HP-UX systems:

Errors that occurred during migration are logged in the `/var/idsldap/V6.0/idsadm.log` file.

On Windows platforms:

Errors that occur during migration are logged in the `install_directory\var\idsadm.log` file.

Kerberos service name change

Before IBM Directory Server 4.1, the LDAP server uses **LDAP** as its Kerberos service name, (**LDAP/ldaphost.austin.ibm.com**; ldaphost is the hostname of the computer where the LDAP server is located) to communicate with its client and the Kerberos KDC. For Version 4.1 and above, a lower case service name is used (**ldap/ldapname.austin.ibm.com**). Because of this change, a 6.0 server might not start after migrating from a 3.2.2 server that is configured to use Kerberos. This is because the 6.0 server is looking for **ldap** in the keytab file in which an **LDAP** service name was located and used by the previous 3.2.2 server. To correct this situation you 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` to **LDAP** before starting the server. This environment variable causes the LDAP server to continue using the upper case LDAP server service name in the keytab file and to communicate with its clients. In the latter case, 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.

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 for 6.0.

Chapter 6. Troubleshooting instance creation and configuration

If you install the proxy server or full server, IBM Tivoli Directory Server requires instance creation and configuration (for the 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 create a directory server instance. You can use either the Instance Administration Tool (**idsxinst**), which has a GUI, or the **idsicrt** command-line utility to create this instance. When you create a directory server instance, a database instance is also created if the 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 6.0 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 a UNIX system.
- 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 6.0 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, Solaris, and HP-UX 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.

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.

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, Solaris, or HP-UX 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 .
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.

NullPointerException when exiting the Configuration Tool: If you exit the Configuration Tool after entering an invalid database name, a NullPointerException 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, Solaris, and HP-UX platforms). If you have a database configuration problem on an AIX, Linux, Solaris, or HP-UX 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 you 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 DB2 CD 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 CD: /db2/license/ldap-custom-db2ese.lic (or \db2\license\ldap-custom-db2ese.lic for Windows)
- If you downloaded a zip file for InstallShield GUI installation: *directory_where_file_was_unzipped*\itdsv60\db2\license\ldap-custom-db2ese.lic
- If you downloaded a tar file for InstallShield GUI installation: *directory_where_file_was_untarred*/itdsv60/ismp/db2/license/ldap-custom-db2ese.lic
- If you downloaded a tar file for operating system utility installation: *directory_where_file_was_untarred*/itdsv60/db2/license/ldap-custom-db2ese.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.

Database performance is poor

For detailed information about improving performance (including information about buffer pools), see the *IBM Tivoli Directory Server Version 6.0 Performance Tuning Guide*.

The BUFFPAGE and DBHEAP database configuration parameters can affect performance. The default BUFFPAGE included with DB2 is 1000 (4 KB pages), which might not be big enough for a large database. Also, if you increase the BUFFPAGE parameter, you must also increase the DBHEAP size by 1 for every 30 incremented in the BUFFPAGE.

DB2 database supports multiple buffer pools. However, unless you know how to do specialized tuning on DB2, use a single buffer pool. This single buffer pool can be specified using the command:

```
db2 alter bufferpool ibmdefaultbp size -1
```

To update the database configuration parameters for a database, use the command:

```
db2 update database configuration for databasename using  
    param value
```

For example, to increase the BUFFPAGE and DBHEAP size, use the command:

```
db2 update database configuration for databasename using  
    BUFFPAGE 20000 DBHEAP 1866
```

Note: For more detailed performance information, see the *IBM Directory Server Version 6.0 Performance Tuning Guide*.

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

The IBM Tivoli Directory Server Version 6.0 Web Administration Tool is installed on an application server, such as the embedded version of IBM WebSphere Application Server - Express, 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 contains 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 adminstartor login and password settings
`${WASHome}/installedApps/DefaultNode/IDSWebApp.war.ear/IDSWebApp.war/WEB-INF/classes/security/console_passwd`
- # console servers & console properties / SSL key database settings
`${WASHome}/installedApps/DefaultNode/IDSWebApp.war.ear/IDSWebApp.war/WEB-INF/classes/IDSConfig/IDSConfig/IDSServersConfig/IDSServersInfo.xml`
- # console properties / component management settings
`${WASHome}/installedApps/DefaultNode/IDSWebApp.war.ear/IDSWebApp.war/WEB-INF/classes/IDSConfig/IDSAppReg/IDSAppReg.xml`
- # console properties / session properties settings
`${WASHome}/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, Solaris, and HP-UX 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 loses connections on HP-UX

If you are using the Web Administration Tool on the HP-UX operating system, you must set the parameters listed in the table below. If you do not set the parameters, the kernel might not allocate enough threads and the system might run out of memory.

The following table contains the parameters and values that must be set before installing Web Administration Tool.

Table 2. HP-UX operating system kernel configuration parameters

| Kernel parameter | Value 256MB+ physical memory |
|------------------|------------------------------|
| max_thread_proc | 1024 |
| maxusers | 256 |
| nproc | 2068(+) |
| nkthread | 3635(+) |

Note: After you update the max_thread_proc and maxusers parameters, be sure that the nproc parameter is set to 2068 or more, and the nkthread parameter to 3635 or more.

Use this procedure to set the kernel configuration parameters:

1. At a command prompt, type: sam
The System Administration Manager opens.
2. Double-click **Kernel Configuration**.
3. Double-click **Configurable Parameters**.
4. Double-click the parameter you want to edit and specify the new value in the **Enter New Formula/Value** field. Click **OK**.
5. Repeat step 4 for each parameter that needs to be set.
6. Click **Actions-->Process New Kernel**.
7. To process the modifications, click **Yes**.
8. Select **Move Kernel Into Place and Shutdown/Reboot Now** and click **OK**.

See the *IBM Tivoli Directory Server Version 6.0 Installation and Configuration Guide* for additional parameter settings.

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

The following problem has been encountered only on the HP-UX and 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. 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 6.0 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 6.0 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 version of WebSphere Application Server - Express.
- 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 version of WebSphere Application Server - Express, V5.1.1.
- An IP address or hostname is part of the URL used to access the Web Administration Tool.

To avoid these errors:

1. If the embedded version of WebSphere Application Server - Express, V5.1.1 is running locally, add **http://localhost** to the list of trusted sites.
2. If the embedded version of WebSphere Application Server - Express, V5.1.1 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`

Troubleshooting the embedded version of WebSphere Application Server - Express

The following sections contains 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 6.0 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=ibm,c=us`.

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=IBM,c=US
objectclass: organization
objectclass: ibm-replicationContext

dn: ibm-replicaGroup=default,o=ibm,c=us
objectclass: ibm-replicaGroup

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

```

objectclass: ibm-replicaSubentry

dn: cn=replicahost,ibm-replicaServerId=masterhost-389,ibm-replicaGroup=default,o=ibm,c=us
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=ibm, c=us

```

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 39 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=ibm,c=us" -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=ibm,c=us" -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 from 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 base  
objectclass=ibm*nt ibm-replicationfailedchanges ibm-replicationfailedchangeount
```

This command can return output similar to the following:

```
cn=<server>-1389,ibm-replicaServerId=<server>-389,ibm-replicaGroup=default,o=ibm,c=us  
ibm-replicationfailedchanges=1 20050407202221Z 68 1  
170814 add cn=entry-85,o=IBM,c=US  
ibm-replicationfailedchangeount=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=ibm,c=us
```

This command can return output similar to the following:

```
dn: entry-85,o=IBM,c=US  
cn: entry-85  
objectclass: person  
objectclass: eperson  
objectclass: organizationalperson  
objectclass: inetorgperson  
objectclass: top  
userpassword: {AES256}tD09yQT540xpp7ZMIg95mA==  
sn: user  
ibm-entryuid: 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::  
MIIQAADJMIQAAAAnCgEAMIQAAAAeBAXjcmVhdG9yc05hbWUxhAAAAoECENOPUFETU10MIQAAA  
AxCGEAMIQAAAAoBA9jcmVhdGVUaW1lc3RhbXAxhAAAAEEDzIwMDUwMzMwMjMyNzQ3wjcEAAAAKA  
oBADCEAAAAHwQNbW9kaWZpZXJzTmFtZTZEAAAAACgQIQ049QURNSU4whAAAAEKAQAwhAAACgED2  
1vZGlmeVRpbWVzdGFtcDGEAAAAEQPMjAwNTAzMzAyMzI3NDda
```

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=ibm,c=us
```

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 base
objectclass=ibm*nt ibm-replicationfailedchanges ibm-replicationfailedchangeount
```

the search can return output similar to the following:

```
cn=<server>-1389,ibm-replicaServerId=<server>-389,ibm-replicaGroup=default,o=ibm,c=us
ibm-replicationfailedchanges=2 20050407214939Z 68 2
170814 add cn=entry-85,o=IBM,c=US
ibm-replicationfailedchangeount=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=ibm,c=us
```

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 base
objectclass=ibm*nt ibm-replicationfailedchanges ibm-replicationfailedchangeount
```

the search can return output similar to the following:

```
cn=<server>-1389,ibm-replicaServerId=<server>-389,ibm-replicaGroup=default,o=ibm,c=us
ibm-replicationfailedchangeount=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`.

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=ibm,c=us, 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=ibm,c=us.
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=ibm,c=us.
```

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=ibm,c=us
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=ibm,c=us objectclass=ibm-repl*
```

where `o=ibm,c=us` 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 6.0 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 <replctx> objectclass=ibm-replicationAgreement
```

where `replctx` is the location where the replication agreements for a replication context are stored; for example, `o=ibm,c=us`. 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 6.0 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>,<CN=SERVER3>,IBM-REPLICAGROUP=DEFAULT,0=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 error 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 server and the server's consumer is a Release 5.1 or 5.2 (with a Fix Pack lower than 3) server, the operation fails. In the server trace, Insufficient Access can be identified as the cause of the failure.

In release 6.0, 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 either Release 5.2 with at least Fix Pack 3 or Release 6.0 consumers. Consumers from either of the following releases are not supported for this extended operation:

- Release 5.1
- Release 5.2 with a Fix Pack level lower than 3

For a consumer of either of these levels 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.
- One of the consumer servers (either the first level or further downstream) is a server from a release prior to version 5.1.
- 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=ibm,c=us` is the context for the topology you want to replicate. If `o=ibm,c=us` is the suffix on the consumer and `cn=people,o=ibm,c=us` does not exist, the operation will fail.
- The `repltopology` extended operation timed out on a heavily loaded consumer. (This results in message `GLPRPL098E`.)
- If the consumer is from release 5.1 or 5.2 and it has no suffix to which the context can "belong", then the `repltopology` extended operation will fail.
In the previous example, if neither `o=ibm,c=us` nor `cn=people,o=ibm,c=us` are suffixes, the `repltopology` extended operation actually creates a suffix, `cn=darshan,cn=people,o=ibm,c=us`. This is true **only** for release 6.0 consumers. Release 5.1 and 5.2 consumers do not have this capability and so the extended operation fails.
- 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 `\sqlib\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.

Chapter 10. Troubleshooting performance

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

Identifying performance problem areas

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

Audit log

The audit log shows what searches are being performed and the parameters used in each search. The 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 5 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=5
```
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 recommended 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 6.0 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, Solaris, and HP-UX 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.

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 6.0 Performance Tuning 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 **idssupport** tool. This tool gathers information that you can provide to IBM Software Support to help identify the problem. See “IBM Tivoli Directory Server Support Tool” on page 9 for information about the **idssupport** 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/V6.0/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 aux command: ' | tee -a /tmp/monitor.out
  ps aux | 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=ibm,c=us

Record and send any errors to IBM Software Support.

Chapter 12. 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 6.0.

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.

Using CTRL+C with a client command that takes passwords can cause an error

On AIX, Solaris, and HP-UX systems only, using the **CTRL+C** keystroke at a client password prompt results in a new command prompt. Any data entered at the new command prompt does not display. To work around this problem, do not use the **CTRL+C** keystroke with a client password prompt

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.

On AIX, Linux, Solaris, and HP-UX operating systems, error messages appear after command prompt when starting the administration daemon

When you start the administration daemon using the command line, an error message might appear after the process has returned. This is a shell limitation on AIX, Linux, Solaris, and HP-UX operating systems.

General troubleshooting

The following sections describe general problems and solutions in IBM Tivoli Directory Server 6.0.

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, Solaris, and HP-UX 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 6.0 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, HP-UX, 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, HP-UX, 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 6.0 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/V6.0/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.

Nondefault 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.

Fixing an "SQL0964C Transaction log for database is full" error

If you are loading a file that contains a large number of entries, you might receive the following error message:

```
SQL0964C The transaction log for the database is full.
SQLSTATE=57011
```

Use the following procedure to increase the size of the transaction log:

1. Determine the current log file size setting by issuing the command:
`db2 get db config for ldapdb2 | grep -i logfilsiz`

2. Increase the size of the log file size setting by issuing the command:
`db2 update db cfg for ldapdb2 using LOGFILSIZ <new value>`
3. Stop the **idsslapd** process.
4. Issue the command:
`db2 force applications all`
5. Restart the **idsslapd** process.

Alternately, you can use the **bulkload** utility to load files with large amounts of entries.

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 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, **gsk7ikm**, to create the necessary certificates. See "Using gsk7ikm" and "Secure Sockets Layer" in the *IBM Tivoli Directory Server Version 6.0 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 frequency of progress reporting.

Set the variable to a positive integer value; for example, 100. Use the following commands:

- On AIX, Linux, HP-UX, 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.
..
```

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. If you have any of the following, you have 64-bit hardware: RS64 I, II, III, IV, POWER3™, POWER3 II, POWER4™, or POWER5™..

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. If you have any of the following, you have 64-bit hardware: RS64 I, II, III, IV, POWER3, POWER3 II, POWER4, or POWER5.

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 5.1, 5.2, or 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.

For Windows 2000, Windows Server 2003 Enterprise and Windows XP client only

The following sections apply only to the Windows 2000, Windows Server 2003 Enterprise and Windows XP 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 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 Admin Daemon** entry in the Services window is set to **Automatic**. If it is not, the administration daemon control program (`ibmdirctl`) will not work.

Appendix A. 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 66

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 or UNIX 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 zSeries, pSeries, and iSeries 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:

| | |
|-------------------|--|
| 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 B. Notices

This information was developed for products and services offered in the U.S.A. IBM might not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the information. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this information at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Department MU5A46
11301 Burnet Road
Austin, TX 78758
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, or other countries, or both:

AIX
DB2
developerWorks
Domino
eServer
IBM
iSeries
Lotus
Passport Advantage

pSeries
Rational
Redbooks
RS/6000
S/390
SecureWay
Tivoli
WebSphere
xSeries
z/OS
zSeries

Intel™, Intel Inside (logos), MMX™ and Pentium® are trademarks of Intel Corporation in the United States, other countries, or both.

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, MS-DOS, Windows, and Windows NT® are registered trademarks of Microsoft Corporation

UNIX is a registered trademark in the United States and/or other countries licensed exclusively through X/Open Company Limited.

Linux is a 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

A

- adminaudit.log 3
- administration daemon
 - audit log 3
 - error log 3
- audit log 3, 49
- audit.log 3

B

- bulkload error log 4
- bulkload.log 4

C

- configuration
 - Configuration Tool 24
 - troubleshooting 25
- configuration tools log 4
- core files
 - AIX operating systems 8
 - description 7
 - HP-UX operating systems 8
 - Linux operating systems 7
 - Solaris operating systems 8
 - Windows operating systems 7
- customer support
 - see Software Support 66

D

- DB2 error log 5
- DB2 rollbacks 50
- DB2 troubleshooting 29
- db2cli.log 5
- debugging
 - advanced output 28
 - description 12
 - ldtrc command 12
 - server debug mode 12

F

- fixes, obtaining 65

I

- ibmdiradm.log 3
- ibmslapd.log 6
- idsldaptrace utility 14
- idslink log 17
- idslink.log 17
- idslink.preview 17
- idsslapd trace 49
- idstools.log 4
- information centers, searching to find
 - software problem resolution 65

- installation
 - overview 15
 - prerequisite software 15
- installation logs 6, 15, 16, 17
- installation troubleshooting
 - InstallShield GUI 17
 - operating system utility 20
- instance creation
 - idsicrt 23
 - Instance Administration Tool 23
 - troubleshooting 24
- Internet, searching to find software
 - problem resolution 65
- isolation levels 50

K

- knowledge bases, searching to find
 - software problem resolution 65

L

- LDAP_DEBUG 12, 13
- LDAP_DEBUG_FILE 13
- ldapinst.log 15
- ldaplp_inst.log 15
- ldtrc command 12
- LOGFILSIZ, modifying 50
- logs
 - administration daemon audit log 3
 - administration daemon error log 3
 - audit log 3
 - bulkload error log 4
 - configuration tools log 4
 - DB2 error log 5
 - DB2 installation
 - AIX 17
 - Linux 17
 - Solaris 17
 - Windows 16
 - DB2 uninstallation
 - Windows 16
 - GSKit installation
 - Windows 17
 - idslink 17
 - installation 6, 15
 - AIX 16
 - Linux 16
 - Solaris 16
 - Windows 16
 - lost and found log 5
 - overview 3
 - server error log 6
 - lost and found log 5
 - lostandfound.log 5

M

- memory leak 53
- memory, adding on Solaris 49

- messages, resolving 2
- migration troubleshooting 21

P

- performance troubleshooting 49
- problem determination
 - describing problem for IBM Software Support 67
 - determining business impact for IBM Software Support 67
 - submitting problem to IBM Software Support 67
- publications
 - accessing online vi
 - IBM Tivoli Directory Server 6.0 v
 - ordering vi
 - related vi

R

- replication
 - overview 39
 - troubleshooting 39

S

- Secure Sockets Layer (SSL) 54
- server error log 6
- SLAPD_OCHANDLERS environment variable 50
- Software Support
 - contacting 66
 - describing problem for IBM Software Support 67
 - determining business impact for IBM Software Support 67
 - submitting problem to IBM Software Support 67
- Support Tool
 - data collected 9
 - location of data collected 12
 - location of log 12
 - overview 9
 - using 11

T

- thread stacks 53
- trace, idsslapd 49
- troubleshooting features, overview 1

U

- uninstallation logs 16
- uninstallation troubleshooting 20

W

Web Administration Tool
troubleshooting 31



Printed in USA

SC32-1679-00

