

IBM Security Directory Server
Version 6.3.1

Troubleshooting Guide



IBM Security Directory Server
Version 6.3.1

Troubleshooting Guide



Note

Before using this information and the product it supports, read the general information under “Notices” on page 147.

Edition notice

Note: This edition applies to version 6.3.1 of *IBM Security Directory Server* licensed program (product number 5724-J39) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Figures	vii
--------------------------	------------

Tables	ix
-------------------------	-----------

About this publication	xi
---	-----------

Access to publications and terminology	xi
Accessibility	xii
Technical training	xiii
Support information	xiii
Statement of Good Security Practices	xiii

Chapter 1. Introduction to problem determination and troubleshooting	1
---	----------

Directory server overview	1
Built-in troubleshooting features	1
Tools for troubleshooting a directory server instance	1
Error Message Reference	2
Information about troubleshooting	2

Chapter 2. Utilities for logging	5
---	----------

Chapter 3. Other diagnostic tools	11
--	-----------

Core file generation	11
Server debug mode	12
Tracing and debugging LDAP client APIs	14
Collecting an ASCII server trace at startup	15
Collecting a binary server trace at startup	16
Collecting performance records dynamically	17
Collecting a dynamic ASCII server trace	17
Collecting trace information	18
Enabling tracing from the command-line interface	18
Enabling tracing with Web Administration Tool	20
Directory server log and configuration file locations	20

Chapter 4. Installation and uninstallation issues	21
--	-----------

Product installation overview	21
Prerequisite software	21
Failure during installation of IBM Security Directory Server with corequisite software	23
Mounting of DVD for installation	23
The <code>idsldap</code> user and group	24
Installation logs	25
Logs for embedded WebSphere Application Server	25
DB2 logs on Windows	26
DB2 logs on AIX, Linux, and Solaris systems	26
The log for <code>idslink</code> on AIX, Linux, Solaris, and HP-UX (Itanium) systems	27
GSKit logs on Windows operating systems	27
Log files for native packages on AIX, Linux, Solaris, and HP-UX operating systems	27
Information for troubleshooting installation and uninstallation	28

Installation with the installation wizard	28
Uninstallation with the installation wizard	32

Chapter 5. Migration issues	33
--	-----------

Migration log files	33
Kerberos service name	33
Database instance or database in configuration file but no longer on system	33
Format of backed-up schema files is incorrect	34
<code>ibm-slapdPlugin</code> entry in configuration file changed	34
Considerations for migration	34

Chapter 6. Instance creation and configuration issues	37
--	-----------

Instance creation overview and common errors	37
Instance creation overview	37
Cannot create another instance because of invalid IP address	38
<code>idssethost</code> command does not recognize second IP address	38
Two directory instances use the same port number	39
Instance creation fails on Windows 2003	39
Administration server fails to start after instance creation	40
Configuration overview and common errors	40
Configuration overview	40
Incorrect status is returned for files	41
Existing database instance and database configuration failure	41
Error occurs when Configuration Tool is started on AIX	41
Configuration programs terminate on AIX	42
DB2 is not configured properly	42
Server does not start after configuration file attributes are changed	42
Transaction log is full	43
Problems in Configuration Tool windows	43
Configuration issues	44

Chapter 7. DB2 issues	47
--	-----------

DB2 license file expired	47
Recovery from migration failure in DB2, version 9.1 or later	48
DB2 9.5 installation on Linux operating systems	48
DB2 diagnostic information in <code>db2diag.log</code>	49
SQL0964C error when large amount of data is loaded	49
Instance starts in config-only mode after DB2 fix pack	50

Chapter 8. Web Administration Tool and application server issues 51

Corruption of data entered in the Web Administration Tool	51
Migration of files before you patch or migrate Web Administration Tool	52
Additional login panels fail	52
Web Administration Tool in inconsistent state	53
Incorrect language is displayed in Web Administration Tool	53
Microsoft Internet Explorer browser problems	54
HTML special characters are not displayed correctly	54
Web Administration Tool requires IBM JDK on Domino server	54
Templates with object class that has no attributes	55
Non-editable fields are displayed as editable	55
Back and Forward buttons not supported	55
Log on issues in Internet Explorer	55
Web Administration Tool issues on Windows Server 2003	56
Web Administration Tool logon fails for new user	56
Web Administration Tool backup creates another folder	57
Embedded WebSphere Application Server - Express on AIX	57

Chapter 9. Replication issues 59

Replication overview	59
Diagnosis of replication errors	59
Sample replication topology	59
Replication status	60
Viewing replication errors with the Web Administration Tool	62
Viewing replication errors with the idsldapsearch command	63
Lost and found log	65
Write and replicated write messages	65
ibm-replicaSubentry object class in a replication topology	65
Command-line utilities to view replication status	66
IBMSLDAPD_REPL_UPDATE_EXTRA_SECS environment variable	67
Information for troubleshooting replication	68
Replicated suffix	68
Verify that suffixes and replication agreements exist	69
Peer-to-peer replication error	69
Insufficient access error	70
Replication topology extended operation fails with result code 80	70
Replication command-line interface error	71
Entries in LDIF file are not replicated	71
Problem with cn=ibmpolicies subtree	72
Master server becomes unstable or stops	73
Stopping a multithreaded replication supplier	74
Synchronizing directory servers in a replicated environment	77
Multimaster configurations	78
Options for replication filter and replication method are not available	79

Consumer server that does not support SHA-2 79

Chapter 10. Performance issues 81

Identification of performance problem areas	81
Adding memory after installation on Solaris systems	81
Setting the SLAPD_OCHANDLERS environment variable on Windows	82
DB2 rollbacks and isolation levels	82
Default value of LOGFILSIZ must be increased	83
Audits for performance profiling	83

Chapter 11. Information for troubleshooting in various scenarios . . . 87

Server is not responding	87
Memory leak is suspected	87
SSL communications return errors	88
Recovering data from a directory server instance where encryption seed value is lost	88
Attribute encryption to be avoided for older versions	89
Character sets larger than 7-bit ASCII in passwords	90
Premature expiry of user password	90
Troubleshooting the limitation in the idssethost command	90
Environment with SNMP agent configured	92
Configuration update to directory server instance after reinstallation	92
Tombstone entries in a directory server	93
Directory server instance backup	95
Configuration of preaudit records for serviceability	96
Entries that are displayed to root and anonymous users	97
Directory server instance is restored to latest consistent state	97
Online backup and restore limitation	98
Log management servers fails to stop	99
Instance does not start and returns error GLPCRY007E	100

Chapter 12. Interoperability 101

Interoperability with Novell eDirectory Server	101
Interoperability with Microsoft Active Directory	101

Chapter 13. Known limitations and general troubleshooting 103

Known limitations	103
IBM Installation Manager preinstallation summary page shows existing software for installation	103
Command-line utilities allow an option to be entered more than once	103
Invalid data entered on command-line utilities	103
No locking mechanism for conflicting commands	104
Unable to drop database	104
Partial replication	104
Replication is not initiated	105
Migration of an instance from version 6.1 or later to version 6.3.1	105
Alias dereferencing does not work	106

Operation times out	106
Instance stops when nCipher cryptographic hardware client is restarted.	106
Error with idsldapdiff tool query	106
Synchronization of entries with idsadsrun utility fails	107
idsadsrun utility fails when instance is run on a different port	107
Operations error during null base search	107
User account gets locked	107
Description attribute for groups does not sync from Active Directory	108
Possible memory leak with PKCS#11 support configured	108
LDIF files with SHA-2 encrypted password or attributes.	109
Multivalued attributes in a virtual list view search	109
Distributed directory environment search scope	110
Instance fails to start if system date is modified	110
Format of the DN gets changed	111
idsdbmaint tool error message.	111
Error opening filename.cat	111
The values TRUE and FALSE are not translated	111
Some schema-related keywords are not translated.	111
Date is not displayed properly for the Russian locale	111
Date and time are displayed in English in translated versions.	112
Error logo is not displayed with error messages	112
Mnemonics missing from tool panels	112
Attribute encryption in RDN of an entry	112
LDAP search filters that exceed 4K are not supported	113
Error during creation of a directory server instance from an existing instance	113
The idsideploy tool fails to restore a database	113
Creation of online backup image fails	113
Unable to connect from OpenLDAP client over DIGEST-MD5	114
Inconsistent data when transaction updates are replicated.	114
Directory server instance creation fails	114
Unable to log on to a system	114
Propagated schema updates rejected.	115
Accessibility tool is unable to read messages in the Configuration Tool	115
General troubleshooting	115
IBM Installation Manager generates an error when the GSKit repository contains multiple installable	116
Instance owner unable to access core file	116
Key labels do not match.	116
GSKit certificate error.	116
Server instance fails to start because of incorrect file permissions.	117
Server instance fails to start because host name is incorrect	117
Server instance cannot be started except by instance owner	117

Error opening slapd.cat file on Windows systems	118
DSML file client produces error	118
Non-default log files need valid path	118
Null searches retrieve entries of deleted suffixes	119
Error occurs with the idsldapsearch command	119
Server behavior when language tags are disabled	119
Key database certificate	119
idsbulkload hangs during parsing phase	120
Size of log file exceeds the system file size limit	120
Unable to connect to directory server over SSL	120
Directory server fails to start after running bulkload	121
The idsadsrun tool fails	121
Startup messages are displayed in two different locales.	121
Unable to open a new connection for an LDAP client	122
Error occurs when you deploy with idsideploy tool.	122
Error occurs because environment variable values contain spaces.	123
The idsadsrun tool stops during synchronization	123
The idsadsrun utility fails to synchronize	123
The idscfgdb command fails to configure a database	124
The idscfgdb command fails with error code GLPCTL028E	125
DMS cooked table space size is not extended exactly	126
Compatibility issue with Common Auditing and Reporting Service (CARS) 6.0.1 server	126
Problem with monitoring server instances on a Solaris system	126
The idsdbrestore utility displays error messages	126
The idsxinst tool fails	126
File path causes backup and restore to fail.	127
Directory server instance starts in config-only mode	127
The warning message GLPSRV147W is displayed	127
Platform-specific issues	127
AIX operating system	127
Windows operating system.	130
Solaris operating system.	136

Appendix A. Common Base Event features 139

Common Base Event scenarios	140
Log archiving and Common Base Event activity interference	141
Log activity of overlapping cycles	141

Appendix B. Support information . . . 143

Knowledge bases	143
Information center on your local system or network	143
Search the Internet	143

Product fixes	143
Contact IBM Software Support	144
Determine the business impact of your problem	145
Describe your problem and gather background	
information	145
Submit your problem to IBM Software Support	145

Notices	147
--------------------------	------------

Index	151
------------------------	------------

Figures

- | | | |
|----|---|-----|
| 1. | Unicode code point U+9DD7 displayed as a square | 131 |
| 2. | U+9DD7 displayed correctly | 131 |

Tables

1. Troubleshooting sections	2	8. Replication-related operational attributes	66
2. The default log file location of IBM Installation Manager	9	9. The results of different search base with search filter, objectclass=*.	94
3. Debug categories.	13	10. The results of different search base with different search filters	94
4. Debug levels	14	11. Entries and multivalued attribute of the entries	109
5. The prerequisite packages that are required on an AIX operating system	21	12. Severity level and their description	145
6. The default log file location of IBM Installation Manager	22		
7. The default log file location of IBM Installation Manager	29		

About this publication

IBM® Security Directory Server, previously known as IBM Tivoli® Directory Server, is an IBM implementation of Lightweight Directory Access Protocol for the following operating systems:

- Microsoft Windows
- AIX®
- Linux (System x®, System z®, System p®, and System i®)
- Solaris
- Hewlett-Packard UNIX (HP-UX) (Itanium)

IBM Security Directory Server Version 6.3.1 Troubleshooting Guide contains information about the possible limitations, problems, and corrective actions that can be attempted before you contact IBM Software Support. The guide also includes information about tools that you can use to determine problems.

Access to publications and terminology

This section provides:

- A list of publications in the “IBM Security Directory Server library.”
- Links to “Online publications” on page xii.
- A link to the “IBM Terminology website” on page xii.

IBM Security Directory Server library

The following documents are available in the IBM Security Directory Server library:

- *IBM Security Directory Server, Version 6.3.1 Product Overview*, GC27-6212-00
Provides information about the IBM Security Directory Server product, new features in the current release, and system requirements information.
- *IBM Security Directory Server, Version 6.3.1 Quick Start Guide*, GI11-9351-01
Provides help for getting started with IBM Security Directory Server. Includes a short product description and architecture diagram, and a pointer to the product documentation website and installation instructions.
- *IBM Security Directory Server, Version 6.3.1 Installation and Configuration Guide*, SC27-2747-01
Contains complete information for installing, configuring, and uninstalling IBM Security Directory Server. Includes information about upgrading from a previous version of IBM Security Directory Server.
- *IBM Security Directory Server, Version 6.3.1 Administration Guide*, SC27-2749-01
Contains instructions for administrative tasks through the Web Administration tool and the command line.
- *IBM Security Directory Server, Version 6.3.1 Command Reference*, SC27-2753-01
Describes the syntax and usage of the command-line utilities included with IBM Security Directory Server.
- *IBM Security Directory Server, Version 6.3.1 Server Plug-ins Reference*, SC27-2750-01
Contains information about writing server plug-ins.

- *IBM Security Directory Server, Version 6.3.1 Programming Reference, SC27-2754-01*
Contains information about writing Lightweight Directory Access Protocol (LDAP) client applications in C and Java™.
- *IBM Security Directory Server, Version 6.3.1 Performance Tuning and Capacity Planning Guide, SC27-2748-01*
Contains information about tuning the directory server for better performance. Describes disk requirements and other hardware requirements for directories of different sizes and with various read and write rates. Describes known working scenarios for each of these levels of directory and the disk and memory used; also suggests rules of thumb.
- *IBM Security Directory Server, Version 6.3.1 Troubleshooting Guide, GC27-2752-01*
Contains information about possible problems and corrective actions that can be taken before you contact IBM Software Support.
- *IBM Security Directory Server, Version 6.3.1 Error Message Reference, GC27-2751-01*
Contains a list of all warning and error messages associated with IBM Security Directory Server.

Online publications

IBM posts product publications when the product is released and when the publications are updated at the following locations:

IBM Security Directory Server documentation website

The <http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDS.doc/welcome.htm> site displays the documentation welcome page for this product.

IBM Security Systems Documentation Central and Welcome page

IBM Security Systems Documentation Central provides an alphabetical list of all IBM Security Systems product documentation. You can also find links to the product documentation for specific versions of each product.

Welcome to IBM Security Systems documentation provides and introduction to, links to, and general information about IBM Security Systems documentation.

IBM Publications Center

The <http://www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss> site offers customized search functions to help you find all the IBM publications you need.

IBM Terminology website

The IBM Terminology website consolidates terminology for product libraries in one location. You can access the Terminology website at <http://www.ibm.com/software/globalization/terminology>.

Accessibility

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

For more information, see the Accessibility Appendix in the *IBM Security Directory Server Product Overview*.

Technical training

For technical training information, see the following IBM Education website at <http://www.ibm.com/software/tivoli/education>.

Support information

IBM Support assists with code-related problems and routine, short duration installation or usage questions. You can directly access the IBM Software Support site at <http://www.ibm.com/software/support/probsub.html>.

IBM Security Directory Server Troubleshooting Guide provides details about:

- What information to collect before you contact IBM Support.
- The various methods for contacting IBM Support.
- How to use IBM Support Assistant.
- Instructions and problem-determination resources to isolate and fix the problem yourself.

Note: The **Community and Support** tab on the product documentation can provide additional support resources.

Statement of Good Security Practices

IT system security involves protecting systems and information through prevention, detection, and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed, misappropriated, or misused or can result in damage to or misuse of your systems, including for use in attacks on others. No IT system or product should be considered completely secure and no single product, service or security measure can be completely effective in preventing improper use or access. IBM systems, products and services are designed to be part of a comprehensive security approach, which will necessarily involve additional operational procedures, and may require other systems, products or services to be most effective. IBM DOES NOT WARRANT THAT ANY SYSTEMS, PRODUCTS OR SERVICES ARE IMMUNE FROM, OR WILL MAKE YOUR ENTERPRISE IMMUNE FROM, THE MALICIOUS OR ILLEGAL CONDUCT OF ANY PARTY.

Chapter 1. Introduction to problem determination and troubleshooting

Troubleshooting is the process of determining why a product is malfunctioning or not functioning as you expect it to. Use the problem determination process and guidelines to troubleshoot problems that are related to IBM Security Directory Server.

Directory server overview

IBM Security Directory Server, previously known as IBM Tivoli Directory Server, is the IBM implementation of Lightweight Directory Access Protocol (LDAP). IBM Security Directory Server provides a specialized in which to store, organize, and retrieve information about objects.

IBM Security Directory Server provides diagnostic tools that can be used to collect information and determine the exact cause of problems that occur. You can use the scenarios and workaround that deal with such topics as installation, configuration, and replication to troubleshoot and fix problems that you might encounter.

Built-in troubleshooting features

IBM Security 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 ensure that core file generation is enabled in order for core file information to be generated. For more information about core files and for instructions for enabling core file generation, see "Core file generation" on page 11.

Error logs

Error logs record error messages that occur during directory server processing. IBM Security Directory Server detects and saves these errors in a text file. For more information, see Chapter 2, "Utilities for logging," on page 5.

Server audit logs

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

Tools for troubleshooting a directory server instance

In addition to the built-in troubleshooting tools, you can use the IBM Support Assistant (ISA) Lite to troubleshoot IBM Security Directory Server.

IBM Support Assistant Lite

IBM Support Assistant Lite is a software support solution that helps to quickly collect diagnostic files. Some examples of diagnostic files that it collects are logs and configuration files, schema files, and traces and core files:

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

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

Error Message Reference

IBM Security Directory Server Error Message Reference contains a list of messages that you might encounter when you use IBM Security Directory Server. It includes messages that are displayed in the 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.

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

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

You can search for GLPL2D014E in the *IBM Security Directory Server Version 6.3.1 Error Message Reference* for information about why the error occurred and how to resolve it.

The following messages are not contained in the *IBM Security Directory Server Version 6.3.1 Error Message Reference*:

- DB2® error log messages
- Lost and found log messages
- Admin audit log messages
- Server audit log messages
- Information messgaes

Information about troubleshooting

In addition to built-in troubleshooting tools, use the troubleshooting sections in the Troubleshooting Guide to resolve issues that occur when you use IBM Security Directory Server.

Table 1. Troubleshooting sections

Troubleshooting sections	For more information, see:
Installation and uninstallation	Chapter 4, "Installation and uninstallation issues," on page 21

Table 1. Troubleshooting sections (continued)

Troubleshooting sections	For more information, see:
Migration	Chapter 5, "Migration issues," on page 33
Instance Creation	Chapter 6, "Instance creation and configuration issues," on page 37
Configuration	Chapter 6, "Instance creation and configuration issues," on page 37
DB2	Chapter 7, "DB2 issues," on page 47
Web Administration Tool and application server	Chapter 8, "Web Administration Tool and application server issues," on page 51
Replication	Chapter 9, "Replication issues," on page 59
Performance	Chapter 10, "Performance issues," on page 81
Scenarios	Chapter 11, "Information for troubleshooting in various scenarios," on page 87
General troubleshooting	Chapter 13, "Known limitations and general troubleshooting," on page 103

Chapter 2. Utilities for logging

IBM Security Directory Server provides several logs that can be viewed either through Web Administration Tool or the system command line. Use these logs to identify the cause of a problem.

For information about viewing the logs, see the *IBM Security Directory Server Administration Guide*. For information about resolving error messages that you find in the logs, see “Error Message Reference” on page 2.

By default, all the logs that are listed here are in the `directory_server_instance_home/logs` (or `directory_server_instance_home\logs` on Windows) directory. The file names that are shown are the defaults, but you can change both the paths and the file names for the logs. For more information, see the *IBM Security Directory Server Administration Guide*. The IBM Security Directory Server logs are:

Administration server log (`ibmdiradm.log`)

An administration server is a limited LDAP server that accepts searches and extended operations to stop, start, and restart the LDAP server. You can view the status and errors that are encountered by the administration server in the administration server log.

A sample of the log is shown here:

```
05/06/2013 02:05:57 PM GLPADM056I Admin server starting.
05/06/2013 02:05:58 PM GLPCOM025I The audit plugin is successfully loaded from
libldapaudit.so.
05/06/2013 02:05:58 PM GLPCOM022I The database plugin is successfully loaded from
libback-config.so.
05/06/2013 02:05:58 PM GLPADM060I The admin server backup and restore server
configuration entry is not enabled.
05/06/2013 02:05:58 PM GLPCOM024I The extended Operation plugin is successfully
loaded from libloga.so.
05/06/2013 02:05:58 PM GLPCOM003I Non-SSL port initialized to 3546.
05/06/2013 02:05:58 PM GLPADM028I Admin server audit logging is started.
05/06/2013 02:05:58 PM GLPADM004I 6.3.1.0 ibmdiradm started
05/06/2013 02:05:58 PM GLPSRV048I Started 5 worker threads to handle client requests.
```

Administration server audit log (`adminaudit.log`)

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

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

Note: If any additional “MAY” attributes are specified, the server ignores the values and no error messages are written.

A sample of the log is shown here:

```

2013-01-15-19:59:17.130-06:00GLPADM028I Admin Server audit logging
is started.
AuditV3--2013-01-16-22:04:50.93986-06:00--V3 Bind--bindDN: CN=ROOT
--client: 127.0.0.1:3665--connectionID: 0--received:
2013-01-16-22:04:50.93986-06:00--Success
AuditV3--2013-01-16-22:04:50.93986-06:00--V3 Search--bindDN: CN=ROOT
--client: 127.0.0.1:3665--connectionID: 0--received:
2013-01-16-22:04:50.93986-06:00--Success
AuditV3--2013-01-16-22:04:50.93986-06:00--V3 Unbind--bindDN: CN=ROOT
--client: 127.0.0.1:3665--connectionID: 0--received:
2013-01-16-22:04:50.93986-06:00--Success
AuditV3--2013-01-16-22:08:09.94185-06:00--V3 Bind--bindDN: CN=ROOT
--client: 127.0.0.1:3678--connectionID: 1--received:
2013-01-16-22:08:09.94185-06:00--Invalid credentials
AuditV3--2013-01-16-22:08:09.94185-06:00--V3 Unbind--bindDN: --client:
127.0.0.1:3678--connectionID: 1--received:
2013-01-16-22:08:09.94185-06:00--Success

```

Server audit log (audit.log)

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

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

```

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

```

For more information about administrative roles and permissions that are required to access various objects, see *IBM Security Directory Server Administration Guide*. Search for the section, *Creating the administrative group*.

A sample of the Server audit log is shown here:

```

2013-01-16-17:38:15.484-06:00--GLPSRV023I Audit logging started.
The audit configuration options are:
ibm-slapdLog = C:\idsslapd-ldaptest\logs\audit.log,
ibm-auditVersion = true,ibm-audit = true,
ibm-auditFailedOPonly = true,ibm-auditBind = true,
ibm-auditUnbind = true,ibm-auditSearch = true,
ibm-auditAdd = true,ibm-auditModify = true,
ibm-auditDelete = true,ibm-auditModifyDN = true,
ibm-auditExtOPEvent = true,ibm-auditExtOp = true,
ibm-auditAttributesOnGroupEvalOp = true,ibm-auditCompare = true,
ibm-auditGroupsOnGroupControl = true.
2013-01-16-17:38:15.656-06:00--GLPSRV009I IBM Security Directory (SSL),
Version 6.3.1 Server started.
AuditV3--2009-01-16-17:39:28.468-06:00--V3 anonymous Search--bindDN:
<*CN=NULLDN*>--client: 127.0.0.1:3792--connectionID: 1
--received: 2009-01-16-17:39:28.453-06:00-- No such object
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: cn=monitor
scope: wholeSubtree
derefAliases: neverDerefAliases
typesOnly: false
filter: (objectclass=*)

```

Bulkload error log (bulkload.log)

The **idsbulkload** (or **bulkload**) command is used to load entries. Use the **bulkload** log to view status and errors that are related to **bulkload**.

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

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

The following **bulkload** error log resulted:

```
04/05/13 09:31:19 GLPCTL113I Largest core file size creation limit for
the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/13 09:31:19 GLPCTL114I Largest file size creation limit for
the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/13 09:31:19 GLPCTL115I Maximum data segment limit for
the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/13 09:31:19 GLPCTL116I Maximum physical memory limit for
the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/13 09:31:19 GLPBLK072I Bulkload started.
04/05/13 09:31:19 GLPBLK050I Extracting parent DNs ...
04/05/13 09:31:19 GLPBLK116E Invalid line detected: 3
04/05/13 09:31:19 GLPBLK044I 1 errors detected during parsing phase.
04/05/13 09:31:20 GLPBLK073I Bulkload completed.
```

Tools log (idstools.log)

The tools log contains status and error messages that are 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 tools log:

```
Aug 09 16:41:02 2013 GLPDPW009I Setting the directory server administrator DN.
Aug 09 16:41:02 2013 GLPDPW010I Set the directory server administrator DN.
Aug 09 16:41:02 2013 GLPDPW006I Setting the directory server administrator
password.
Aug 09 16:41:11 2013 GLPDPW007I Set the directory server administrator
password.
Aug 09 16:41:17 2013 GLPCDB035I Adding database 'ldaptest' to directory server
instance: 'ldaptest'.
Aug 09 16:41:18 2013 GLPCTL017I Cataloging database instance node: 'ldaptest'.
Aug 09 16:41:19 2013 GLPCTL018I Cataloged database instance node: 'ldaptest'.
Aug 09 16:41:19 2013 GLPCTL008I Starting database manager for database
instance: 'ldaptest'.
Aug 09 16:41:22 2013 GLPCTL009I Started database manager for database
instance: 'ldaptest'.
Aug 09 16:41:22 2013 GLPCTL026I Creating database: 'ldaptest'.
Aug 09 16:43:11 2013 GLPCTL027I Created database: 'ldaptest'.
Aug 09 16:43:11 2013 GLPCTL034I Updating the database: 'ldaptest'
Aug 09 16:43:19 2013 GLPCTL035I Updated the database: 'ldaptest'
Aug 09 16:43:19 2013 GLPCTL020I Updating the database manager: 'ldaptest'.
Aug 09 16:43:22 2013 GLPCTL021I Updated the database manager: 'ldaptest'.
Aug 09 16:43:23 2013 GLPCTL023I Enabling multi-page file allocation:
'ldaptest'
Aug 09 16:43:37 2013 GLPCTL024I Enabled multi-page file allocation:
'ldaptest'
Aug 09 16:43:38 2013 GLPCDB005I Configuring database 'ldaptest' for
directory server instance: 'ldaptest'.
Aug 09 16:43:39 2013 GLPCDB006I Configured database 'ldaptest' for
directory server instance: 'ldaptest'.
Aug 09 16:43:39 2013 GLPCDB003I Added database 'ldaptest' to directory
server instance: 'ldaptest'.
```

DB2 log (db2cli.log)

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

The following is a sample of the DB2 log:

```

2013-09-13-19:18:29.native retcode = -1031; state = "58031";
    message = "SQL1031N"
    The database directory cannot be found on the indicated file system.

SQLSTATE=58031

"
2013-09-13-19:18:29.native retcode = -1018; state = "E8";
    message = "SQL1018N"
    The node name "idsinode" specified in the CATALOG NODE command
    already exists.

"
2013-09-13-19:18:30.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 because of replication conflict resolution. Use the log of these entries to recover the data in the replaced entries if necessary.

The information that is 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 that is 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=utl,c=us
#Operation type:Add
#Corrective action:Replace
#Entry createTimeStamp: 20131106211242.000000Z
#Entry modifyTimeStamp: 20131030202533.000000Z
#Supplier address: 9.53.21.187
dn: cn=t6,o=utl,c=us
objectclass: person
objectclass: top
sn: aa
cn: aa
cn: t6
description: this should not be here
ibm-entryuuid: 0c4559de-0a76-4c91-96e4-5ae81d405466

```

Server log (ibmslapd.log)

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

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

```

Sep 13 14:31:04 2013 GLPL2D014E Suffix entry has not been created for
    entry cn=Robert Dean, ou=In Flight Systems, ou=Austin, o=sample.
Sep 13 14:31:04 2013 GLPRDB002W ldif2db: 0 entries have been successfully
    added out of 50 attempted.
Sep 13 14:39:41 2013 GLPCOM024I The extended Operation plugin is
    successfully loaded from libevent.dll.
Sep 13 14:39:41 2013 GLPCOM024I The extended Operation plugin is
    successfully loaded from libtranext.dll.

```


Installation and uninstallation logs

In addition, there are logs that are created during installation, modification, and uninstallation. The installation, modification, and uninstallation logs of the installation wizard are stored in the logs directory of IBM Installation Manager.

Table 2. The default log file location of IBM Installation Manager

Operating system	Default log path:
Microsoft Windows	C:\ProgramData\IBM\InstallationManager\logs
AIX and Linux	/var/ibm/InstallationManager/logs/

For more information about these logs, see “Failed installation recovery” on page 29.

Backup status file (dbback.dat)

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

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

If the entry is present and enabled, the Administration Server checks the backup location from the configuration for the date and time of current backup. Monitor searches can be used to fetch the latest snapshot of the data that pertains to back up or restore. The file `dbback.dat` is the prime source for monitor searches to fetch their data from. The `dbback.dat` file is created at a backup location that you specify when you configure backup, for example `backup_location/BACKUP_FILES`.

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

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

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

Chapter 3. Other diagnostic tools

Several diagnostic tools are built into IBM Security Directory Server and operating systems to help users and IBM Software Support determine why a problem is occurring. Configure and use these tools to gather information for troubleshooting.

Core file generation

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 enable 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 that is called Dr. Watson to generate a text file named `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 starts automatically. To start Dr. Watson manually by using the graphical user interface (GUI):

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 when Dr. Watson runs for the first time, 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 specifies the path. To determine whether the crash dump file is 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` can be found in the working directory of the process. The core file is generated 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 can see that a core file is 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. The following example shows the generated output:

```
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
```

To disable core file generation, use the following command:

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

Server debug mode

At times, the error logs do not provide enough information to resolve a problem. You can run IBM Security Directory Server in a special debug mode that generates more 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 effect on performance.

After you run 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 that is specified with the **-d** option overrides the debug mask that is specified in the environment variable.

Table 3. 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 back-end performance statistics
0x1000	8192	LDAP_DEBUG_RDBM	Relational back-end activities (RDBM)
0x4000	16384	LDAP_DEBUG_REFERRAL	Referral activities
0x8000	32768	LDAP_DEBUG_ERROR	Error conditions
0xffff	65535	LDAP_DEBUG_ANY	All levels of debug

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

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

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

```
ldtrc off
```

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

The generated debug output is displayed to standard error. You can place the output in a file in one of the following ways:

- Set the `LDAP_DEBUG_FILE` environment variable.
- On server commands (but not the `idsldapd` 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 that is 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 Security Directory Server Command Reference* for information about the `idsldaptrace` utility.

Tracing and debugging LDAP client APIs

You can enable tracing for LDAP client application programming interfaces (APIs) and use the information that is captured in the trace file to debug problems.

Before you begin

Before you enable tracing for LDAP client APIs, you must first stop the LDAP client application.

Procedure

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

On AIX, Linux, Solaris, and HP-UX (Itanium) platforms

```
$export LDAP_DEBUG=debug_level
```

On Windows platform

```
c:\>set LDAP_DEBUG=debug_level
```

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

Table 4. Debug levels

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

Table 4. Debug levels (continued)

Decimal	Value	Description
8192	LDAP_DEBUG_RDBM	Relational back-end activities (RDBM)
16384	LDAP_DEBUG_REFERRAL	Referral activities
32768	LDAP_DEBUG_ERROR	Error conditions
65535	LDAP_DEBUG_ANY	All levels of debug

For example, specifying a bit mask value of 65535 turns on full debug output and generates the most complete information. To know more about debug levels, see “Server debug mode” on page 12.

- Set the debug file name by using the `LDAP_DEBUG_FILE` environment variable.

On AIX, Linux, Solaris, and HP-UX (Itanium) platforms

```
$export LDAP_DEBUG_FILE=filename
```

ON Windows platform

```
c:\>set LDAP_DEBUG_FILE=filename
```

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

- Run the application from the same terminal where you have environment set. Re-create the problem that you want to debug.
- The debug information is captured in the file pointed by the `LDAP_DEBUG_FILE` environment variable. You can now use the information that is captured in the file to debug the problem. Or send this file to the IBM Support team for further analysis.

Collecting an ASCII server trace at startup

Collect ASCII server trace to determine and debug issues that are related to a failed directory server startup. You can also use it to trace a specific operation at directory server startup.

Procedure

- Stop the directory server instance, if it is running. Issue the command of the following format:
`ibmslapd -I instance_name -k`
- Determine whether tracing is enabled. Issue the following command:
`ldtrc info`
- If trace is disabled (in "off" mode), issue the following command to enable tracing:
`ldtrc on`
- Start the directory server in DEBUG mode and redirect the output to a file.

On AIX, Linux, and Solaris platforms

Issue the command of the following format:

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

On Windows platform

Issue the command of the following format:

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

- Re-create the problem. After the error or the condition you want to trace occurs and the screen no longer displays messages, press `Ctrl C` to stop the process. Now, you can analyze the trace file.

6. Disable tracing. Issue the following command:
`ldtrc off`

Collecting a binary server trace at startup

To debug issues that are related to a failed directory server startup, you must collect a binary server trace. You can also use it to trace a specific operation at directory server startup.

Procedure

1. Stop the directory server instance, if it is running. Issue the command of the following format: `ibmslapd -I instance_name -k`
2. Determine whether tracing is enabled or not. Issue the following command:
`ldtrc info`.
3. If trace is enabled, disable the trace. Issue the following command: `ldtrc off`.
4. Enable binary tracing. Issue the following command: `ldtrc on -l 50000000`
In the command, the value for the buffer size is set to 50 million bytes. This buffer size stores the latest 50 million bytes of trace record data in the shared memory. It flushes the oldest data when the 50-MB value is reached. If the command fails because of what might seem to be insufficient shared memory resources, you can scale the number down. However, less than 20 million might not provide the required information.
5. Start the directory server instance. Issue the command of the following format: `ibmslapd -I instance_name -n`.
6. Point the environment variable `TRCTFIDIR` to the `DS_INSTALL_HOME` directory. Use the following command:

On AIX, Linux, and Solaris platforms:

```
$export TRCTFIDIR=DS_INSTALL_HOME/etc
```

On Windows platform:

```
C:\> set TRCTFIDIR=DS_INSTALL_HOME\etc
```

where, `DS_INSTALL_HOME` on different operating system is as follows:

On AIX and Solaris platforms:

```
/opt/IBM/ldap/V6.3.1/etc
```

On Linux platform:

```
/opt/ibm/ldap/V6.3.1/etc
```

On Windows platform:

```
Install_Drive:\Program Files\IBM\LDAP\V6.3.1\etc
```

7. Re-create the problem to produce the error or condition that you want to trace.
8. Collect the trace records. After the error or the condition you want to trace occurs, issue the following command:
`ldtrc dump trace.raw`

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

9. Change to the `DS_INSTALL_HOME/etc` directory and then collect the format and flow of the binary trace. Issue the following commands: `ldtrc fmt trace.raw` `trace.fmt` `ldtrc flw trace.raw trace.flw`. Send the `trace.fmt` and `trace.flw` files to support.

10. Disable tracing. Issue the following command: `ldtrc off`

Collecting performance records dynamically

The performance profile information in trace is intended to help users diagnose performance problems. By using the independent trace facility, performance profiling is accomplished with minimum impact on server performance.

About this task

The independent trace facility profiles operation performance that consists of timestamps at key points that are traversed during an operation execution for a running server instance. The timestamps are profiled during different stages such as the following stages:

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

To activate tracing of performance records dynamically, complete the following steps.

Procedure

1. Activate tracing for performance records. Issue the following command:

```
ldaptrace -h hostname -p port number -D adminDN -w adminPW -l on \
-t start -- -perf
```
2. Dump the trace to a binary trace file. Issue the following command:

```
ldtrc dump trace.bin
```
3. Format the trace. Issue the following command:

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

What to do next

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

```
ldtrc off
```

For more information about performance profiling, see the *IBM Security Directory Server Administration Guide*.

Collecting a dynamic ASCII server trace

Collecting a dynamic ASCII server trace helps in debugging issues that are related to a specific operation of a server. You can collect a dynamic server trace only if the IBM Security Directory Server instance that you want to debug is running.

Procedure

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

```
idsilist -I instance_name -a
```
2. Start the dynamic ASCII server trace for your directory server instance.

On AIX, Linux, and Solaris platforms

Issue the command of the following format: `idsldaptrace -p port -a admin_port -D adminDN -w adminPW\ -h hostname -l on -t start -m 65535 -o /tmp/ibmslapd.dbg`

On Windows platforms

Issue the command of the following format: `idsldaptrace -p port -a admin_port -D adminDN -w adminPW\ -h hostname -l on -t start -m 65535 -o C:\Temp\ibmslapd.dbg`

3. Re-create the problem and issue the specific operation that is failing.
4. Disable the dynamic ASCII server trace. Issue the command of the following format: `idsldaptrace -p port -a adminPort -D adminDN -w adminPW \ -h hostname -l off -t stop`

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

Collecting trace information

Enabling tracing is a multistep process that involves starting the trace facility and printing trace information. The trace facility enables tracing of directory server and other commands. You can use the command line interface or the Web Administration Tool to enable tracing.

Enabling tracing from the command-line interface

An administrator can enable the trace facility and request for specific processes like directory server or commands like **ldif2db** to print trace information. Trace information can be sent to the command line or to a file.

Procedure

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

OR,

```
idsldaptrace -p adminServerPort -h host_name -D cn=adminDN \
-w adminPW -l on
```

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

2. Enable the tracing for a specific process or a command. Select a debug level for the trace. For example, specifying a bit mask value of 65535 turns on full debug output and generates the most complete information. For more information about debug levels, see “Server debug mode” on page 12. You can use one of the following options to set the debug level that is base on the process or command that you want to trace.
 - Set the `LDAP_DEBUG` environment variable to specify the debug level. Set this environment variable with a value that you want to use for `debug_mask`. If the `LDAP_DEBUG` environment variable is set and you use the **-d** option with a different debug mask, the debug mask that is specified with the **-d** option overrides the debug mask that is specified in the environment variable.

On AIX, Linux, and Solaris platforms:

```
$export LDAP_DEBUG=debug_level
```

On Windows platform:

```
C:\> set LDAP_DEBUG=debug_level
```

To disable the environment variable, use the **unset LDAP_DEBUG** command.

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

```
idsldapmodify -p port -D cn=adminDN -w adminPW
dn: cn=Configuration
changetype: modify
replace: ibm-slapdStartupTraceEnabled
ibm-slapdStartupTraceEnabled: TRUE
f-
replace: ibm-slapdTraceMessageLevel
ibm-slapdTraceMessageLevel: 0xFFFF
-
replace: ibm-slapdTraceMessageLog
ibm-slapdTraceMessageLog: /var/ibmslapd.trace.log
```

Operation 0 modifying entry cn=Configuration

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

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

- You can also dynamically enable tracing after a directory server instance starts by using the **idsldaptrace** command.

To start tracing a directory server, issue the **idsldaptrace** command of the following format:

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

To stop tracing of a directory server, issue the **idsldaptrace** command of the following format:

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

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

- To stop tracing, issue the following command:
ldtrc off

OR

```
idsldaptrace -p adminServerPort -h host_name -D cn=adminDN \
-w adminPW -l off
```

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

Enabling tracing with Web Administration Tool

You can use the Web Administration Tool to enable tracing.

About this task

If you use the Web Administration Tool, it takes care of starting and stopping the trace facility. To know more about logging utilities, see *IBM Security Directory Server Administration Guide*.

Procedure

1. In the Web Administration Tool navigation area, under **Server administration**, select **Logs**.
2. On the expanded list, select **Start/Stop server trace** to enable or disable server tracing.
3. In the **Trace debug levels** field, you can specify the debug level. For more information about debug levels, see “Server debug mode” on page 12.
4. In the **Trace debug file** field, you can specify the output file to store trace information.

Directory server log and configuration file locations

To diagnose any issue that is related to IBM Security Directory Server, it is important to collect the log and configuration file.

You can find the log and configuration file that you usually check to determine issues that are related to IBM Security Directory Server at the following locations for version 6.3.1.

On AIX, Linux, and Solaris platforms

- Configuration file: *instance_home/idsslapd-instance_name/etc/ibmslapd.conf*
- Administration Server log file: *instance_home/idsslapd-instance_name/logs/ibmslapd.log*
- DB2 error log: *instance_home/idsslapd-instance_name/logs/db2cli.log*
- Audit log file: *instance_home/idsslapd-instance_name/logs/audit.log*

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

On Windows platforms

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

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

Chapter 4. Installation and uninstallation issues

There are many points during the installation of a product and its prerequisite software where problems might be encountered. Use this information to troubleshoot problems during the installation process and proceed with recovery actions.

Product installation overview

IBM Security Directory Server consists of several components that you can install by using the installation wizard or the command-line interface.

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

- C Client
- Java Client
- IBM Java Development Kit
- Server
- Proxy Server
- Web Administration Tool
- IBM DB2 (Enterprise Server Edition / Workgroup Server Edition)
- IBM Global Security Kit (GSKit)

You can install these components by using the installation wizard or operating-system-specific installation methods, such as the command line or installation tools for the operating system.

Prerequisite software

To start IBM Security Directory Server installation with IBM Installation Manager, your computer must contain IBM Installation Manager, version 1.7.0 or later. If you are using the operating system utilities for the installation, your computer must contain the prerequisite software.

To use IBM Installation Manager, your computer must contain the prerequisite packages that are required by IBM Installation Manager. The following prerequisite packages are required for installation of IBM Security Directory Server with IBM Installation Manager on the following operating systems:

AIX For installation of rpm packages on AIX, download the rpm package manager for AIX systems from the <ftp://public.dhe.ibm.com/aix/freeSoftware/aixtoolbox/INSTALLP/ppc/rpm.rte> website.

Table 5. The prerequisite packages that are required on an AIX operating system

Packages	Reason	Download address
Mozilla Firefox web browser for AIX	To open the launchpad on AIX, a supported version of browser must exist.	For more information about the web browsers for AIX, see the http://www.ibm.com/systems/power/software/aix/browsers/ website.

Table 5. The prerequisite packages that are required on an AIX operating system (continued)

Packages	Reason	Download address
gtk+ RPM (gtk2-2.10.6-4.aix5.2.ppc.rpm or later)	Eclipse changed the window system requirement from motif to gtk on UNIX operating systems. For AIX, this Eclipse window system change requires the gtk libraries to be installed to support the GUI. For IBM Installation Manager, the GUI is the wizard mode of operation.	For more information about installation of the gtk libraries, see the Required gtk libraries for Installation Manager on AIX technote at the http://www.ibm.com/support/docview.wss?uid=swg21631478 website.
GNU tar	To uncompress archive files that are provided with IBM Security Directory Server on AIX systems, the GNU file archive program is required. You must set the path of the GNU tar program before the tar program provided with the operating system. The GNU tar program is installed in the /opt/freeware/bin directory, and the tar program that is provided with the operating system in the /usr/bin directory. To set the /opt/freeware/bin path, run the following command: export PATH=/opt/freeware/bin:\$PATH.	To download the GNU tar archive file (tar), see the http://www.ibm.com/systems/power/software/aix/linux/toolbox/alpha.html website.

To obtain a detailed list of all the prerequisite software for your operating system, see the IBM Security Directory Server System Requirements documentation at <http://pic.dhe.ibm.com/infocenter/prodguid/v1r0/clarity/index.html>.

If installation does not complete successfully with IBM Installation Manager, you must check the logs. IBM Installation Manager creates the log files related to installation, modification, uninstallation in the following locations:

Table 6. The default log file location of IBM Installation Manager

Operating system	Default log path:
Microsoft Windows	C:\ProgramData\IBM\InstallationManager\logs
AIX and Linux	/var/ibm/InstallationManager/logs/

IBM Installation Manager creates an xml log file in following format: `yyyymmdd_hhmm`. The following is an example of log file that is created during the installation of IBM Security Directory Server, `20131010_1146.xml`.

The `yyyymmdd_hhmm` format represents the following values:

- `yyyy` represents year
- `mm` represents month

- dd represents day
- HH represents hour in 24-hour format
- MM represents minutes

Failure during installation of IBM Security Directory Server with corequisite software

Use the installation log to identify the cause of the failure when you are installing IBM Security Directory Server with corequisite software.

IBM Security Directory Server uses corequisite software, such as IBM DB2, IBM Global Security Kit, IBM Java Development Kit, and IBM Embedded WebSphere Application Server to run various features.

If a failure occurs while you are installing corequisite software, you might observe various error messages. The results that you see can vary according to the software that you are installing and other related components when the failure occurs.

To avoid installation failures, you must meet the minimum software and hardware requirements that are required for IBM Security Directory Server. However, it is always advisable to provide more resources than the minimum required.

If you are installing IBM Security Directory Server with corequisite software, you must provide the appropriate installable path for the corequisite software. Installation might not continue if the corequisite software cannot be located that you selected. If the installation of IBM Security Directory Server fails with IBM Installation Manager, no features are installed and Installation Manager rolls back the installation.

- If you are installing from downloaded and uncompressed .zip or .tar files:
 - Uncompress the .zip files to your computer in a path that does not contain a space in the path name. Uncompress all .zip files in the same directory.
 - Uncompress all .tar files in the same directory that does not contain a space in the path name. Uncompress all .tar files in the same directory.
- The .iso file version of the product are used to burn installation DVDs that can then be used in the installation process. The .iso files are images that must be processed through a DVD burner program to create DVDs. When you create the DVDs, be sure that you do not make data DVDs of the .iso files. Select the option that unencapsulates the data from the .iso files and burns the files on the DVD.

Mounting of DVD for installation

You must mount DVD that contains the .iso files of IBM Security Directory Server for the installation of the product and its corequisite software.

About this task

On a Linux operating system, you might see the following error message when you run the installation program:

```
Bad interpreter: /bin/sh: Permission denied
```

The mount issue might occur if the default automount settings contain -noexec permission. You must use the valid parameters before you mount the DVD and run the installation program.

Procedure

1. Run the following command to unmount the DVD:
`umount /media/DirectoryServerV6.3.1`
2. Run the following command to mount the DVD again:
`mount -o loop,ro /dev/hda /mnt/cdrom/`

The idsldap user and group

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

The requirements are:

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

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

On AIX systems:

Use the following commands.

To create the idsldap group:

```
mkgroup idsldap
```

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

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

To set the password for user idsldap:

```
passwd idsldap
```

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

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

On Linux systems:

Use the following commands.

To create the idsldap group:

```
groupadd idsldap
```

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

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

To set the password for user idsldap:

```
passwd idsldap
```

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

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


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

On Solaris systems:

Use the following commands.

To create the `idsldap` group:

```
groupadd idsldap
```

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

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

To set the password for user `idsldap`:

```
passwd idsldap
```

To modify the root user ID so that root is a member of the group `idsldap`, use an appropriate tool.

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

Installation logs

The log files that are used during various stages of installation can help you to identify problems that you encounter when you use the installation wizard.

Logs for embedded WebSphere Application Server

The logs that are used by the installation wizard when you install embedded WebSphere® Application Server can help you to identify the cause of related installation problems.

On Windows platforms

- `install_home\var\installApp.log`
- `install_home\var\installAppErr.log`
- `install_home\var\configApp.log`
- `install_home\var\configAppErr.log`
- `install_home\var\migrateApp.log`
- `install_home\var\migrateAppErr.log`

The following logs are used when you add embedded WebSphere Application Server Web Administration Tool as a Windows service.

- `addWebAdminSrv.log`
- `addWebAdminSrvErr.log`

The following logs are used when you start embedded WebSphere Application Server Web Administration Tool as a Windows service.

- `startWebAdminSrv.log`
- `startWebAdminSrvErr.log`

On AIX, Linux, and Solaris platforms

- `/var/idsldap/V6.3.1/installApp.log`
- `/var/idsldap/V6.3.1/installAppErr.log`
- `/var/idsldap/V6.3.1/configApp.log`

- /var/idsldap/V6.3.1/configAppErr.log
- /var/idsldap/V6.3.1/migrateApp.log
- /var/idsldap/V6.3.1/migrateAppErr.log

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

DB2 logs on Windows

The logs that are used by the installation wizard when you install and uninstall DB2 on Windows can help you identify the cause of related problems.

When you install

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

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

When you uninstall

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

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

DB2 logs on AIX, Linux, and Solaris systems

The logs that are used when you install DB2 on AIX, Linux, and Solaris systems can help you identify the cause of related installation problems.

When you install with the installation wizard

- /var/idsldap/V6.3.1/db2inst.log
- /var/idsldap/V6.3.1/db2insterr.log
- /var/idsldap/V6.3.1/DB2setup.log

When you uninstall

- /var/idsldap/V6.3.1/db2uninst.log
- /var/idsldap/V6.3.1/db2uninsterr.log

When you install by using the db2_install utility

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

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

You must run the `idslink` script manually when you install with the installation wizard and the operating system utility of the client, IBM Security Directory Server Proxy Server and IBM Security Directory Server Full Server.

The `idslink.log` and `idslink.preview` files are in the `/var/idsldap/V6.3.1/` directory.

GSKit logs on Windows operating systems

The logs that are used by the installation wizard when you install and uninstall the GSKit on Windows systems can help you troubleshoot related problems.

The logs that are used by the installation wizard when you install and uninstall GSKit on Windows systems are:

- `install_home\var\gsksetup.log`
- `install_home\var\gskitinst.log`
- `install_home\var\gskitinsterr.log`

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

On AIX, Linux, Solaris, and HP-UX (Itanium) platforms, two logs are generated for each native package that is installed. These logs give information about the native packages. You can refer to these logs to determine the reason why an installation failed.

The log files are created in the `/var/idsldap/V6.3.1` directory. These log files are of importance since the installation wizard installs the native packages in the background.

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

The various log files that are created during installation are:

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

Note: Here, XX can be either 64 or 32 depending on whether the hardware is 64-bit or 32-bit.

- `clientBaseErr.log`, `clientBase.log`
- `engMsgErr.log`, `engMsg.log`
- `gsKitErr.log`, `gsKit.log`
- `javaClientErr.log`, `javaClient.log`
- `proxyErr.log`, `proxy.log`
- `serverErr.log`, `server.log`
- `srvBaseErr.log`, `srvBase.log`
- `webAdminErr.log`, `webAdmin.log`

On AIX systems, some more logs are generated for SSL packages. The log files are created in the `/var/idsldap/V6.3.1` directory.

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

On AIX, Linux, Solaris, and HP-UX (Itanium) platforms, when native packages are uninstalled log files are created in the `/var/idsldap/V6.3.1/uninstall` directory. The various log files that are created depending on the native packages that you install are:

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

On AIX systems, in addition to these log files that are created for native packages during uninstallation, the following log files are also created:

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

Information for troubleshooting installation and uninstallation

Use the fixes provided for troubleshooting installation problems with IBM Security Directory Server.

Installation with the installation wizard

When you use the installation wizard for installation, you might face certain issues with the installation. Use the explanations and solutions to understand the cause of the problems and to resolve them.

IBM Installation Manager might remove an existing version of IBM Global Security Kit

IBM Security Directory Server uses IBM Global Security Kit for securing communication with its clients by using the secure communication protocol.

IBM Installation Manager might remove an existing IBM Global Security Kit when you meet all the following conditions:

1. IBM Global Security Kit, version 8.0 or later exists on your computer that is earlier than the IBM Global Security Kit version that is provided with IBM Security Directory Server.
2. The IBM Global Security Kit feature is selected that is provided with IBM Security Directory Server for installation or modification.

3. You clicked the Cancel button during the installation or modification operation.

Installation failure because of lack of disk space

Understand the disk space requirements for installation to resolve issues that are related to installation failure because of lack of disk space.

One of the reasons for an installation failure is lack of disk space. IBM Security Directory Server attempts to verify that there is enough space and generates messages if the required disk space is not found. However, sometimes the installation wizard cannot progress far enough to issue a message.

Before you install, make sure that you have the required free disk space available as specified in the *IBM Security Directory Server 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.

Failed installation recovery

If your IBM Security Directory Server installation fails, IBM Installation Manager rolls back the installation. You must see the installation log file to verify the reason for failure.

After IBM Installation Manager rolls back the installation, you must check whether the IBM Security Directory Server installation location is empty. If there are files or subdirectories in the IBM Security Directory Server installation location, you must take a backup of the files and subdirectories. You must remove the files and subdirectories in the IBM Security Directory Server installation location before you run the installation again.

For information about uninstalling by using the installation wizard, see the *IBM Security Directory Server Installation and Configuration Guide*.

IBM Installation Manager creates the log files related to installation, modification, uninstallation in the following locations:

Table 7. The default log file location of IBM Installation Manager

Operating system	Default log path:
Microsoft Windows	C:\ProgramData\IBM\InstallationManager\logs
AIX and Linux	/var/ibm/InstallationManager/logs/

IBM Installation Manager creates an xml log file in following format: `yyyymmdd_HHMM`. The following is an example of log file that is created during the installation of IBM Security Directory Server, `20131010_1146.xml`.

The `yyyymmdd_hhmm` format represents the following values:

- `yyyy` represents year
- `mm` represents month
- `dd` represents day
- `HH` represents hour in 24-hour format
- `MM` represents minutes

Recovering from a failed installation on Windows operating systems:

On Windows systems, if the installation with the installation wizard fails, you must follow the steps for uninstallation and manual cleanup before you reinstall.

Procedure

1. Correct any problems that are listed in the C:\Program Files\IBM\InstallationManager\eclipse\yyyymmdd_HHMM.xml file.
2. Remove the IBM Security Directory Server installation directory. The default directory is C:\Program Files\IBM\ldap\V6.3.1.
3. Use **regedit** to remove the LDAP entry in the registry: HKEY_LOCAL_MACHINE\SOFTWARE\IBM\IDSLDAP\6.3.1.

Recovering from a failed installation on AIX operating systems:

On AIX systems, if the installation with the installation wizard fails, you must follow the steps for uninstallation and manual cleanup before you reinstall.

Procedure

1. Correct any problems that are listed in the /opt/IBM/InstallationManager/eclipse/yyyymmdd_HHMM.xml file.
2. Type the following at a command prompt to verify if packages are installed:
`lslpp -l |grep -i ids1`
3. Remove the /opt/IBM/ldap/V6.3.1 directory.

Recovering from a failed installation on Linux operating systems:

On Linux systems, if the installation with the installation wizard fails, you must follow the steps for uninstallation and manual cleanup before you reinstall.

Procedure

1. Correct any problems that are listed in the /opt/IBM/InstallationManager/eclipse/yyyymmdd_HHMM.xml file.
2. Type the following at a command prompt to verify if the packages are installed:
`rpm -qa | grep -i ids1`
3. If an **rpm** command hangs, try running the command with the **noscripts** option: `rpm -ev --noscripts packagenames`
4. Remove the /opt/ibm/ldap/V6.3.1 directory.

Missing files after server installation

After you use installation wizard to install on AIX, Linux, or Solaris systems, files such as `idsxinst`, `idsicrt`, or `idsilist` might be missing. When these files are missing, it indicates that IBM Security Directory Server Proxy Server feature might not be installed correctly. Follow the steps to resolve problems that are related to missing files.

About this task

You might notice this problem when instance creation begins because the Instance Administration Tool is not available.

Procedure

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 steps
 - Modify the `idsldap` user so that it belongs to the `idsldap` group.
 - Delete the `idsldap` user and the `idsldap` group and then do one of the following steps:
 - Re-create the `idsldap` user and group as described in the section The `idsldap` user and group.
 - Do not re-create the `idsldap` user and group. IBM Security Directory Server installation re-creates them when you do step 3.)
3. Reinstall the base server package, and the server or proxy server packages according to the type of server that is required.

Results

The base server package gets installed with IBM Security Directory Server Proxy Server or with IBM Security Directory Server Full Server package. If you are using the installation wizard, this hidden feature is installed when you choose IBM Security Directory Server Proxy Server or IBM Security Directory Server Full Server.

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

Operating system utility uninstallation after installation with the installation wizard

If you installed by using the installation wizard and then do an operating system uninstallation, you must manually clean up the registry entries.

If you used the installation wizard to install, the *IBM Security Directory Server Installation and Configuration Guide* instructs you to use the installation wizard to uninstall. However, if you do an operating system uninstallation instead, you must clean up any registry entries that might be made by the installation process. For instructions for cleaning up the registry entries, see “Failed installation recovery” on page 29.

Installation fails on Windows operating system

On a Microsoft Windows operating system, the installation might fail and display a message such as “DB2 Installation was NOT successful.” Understand the cause of the problem and follow the steps to ensure a successful installation.

If you are creating a user ID and your system has “Password must meet complexity requirements” enabled, ensure that the password meets the complexity requirements. If it does not, the installation fails. See the Windows documentation for information about password complexity requirements.

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

Password must meet complexity requirements

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

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

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

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

If the password complexity requirements are not met, the installation fails and logs messages to `DB2setup.log` in the `install_location\var` directory. Here is an extract from the log:

```
Found echo string in C:\Program Files\IBM\LDAP\V6.3.1\var\db2inst.log file.
returnCode from DB2 Install is set to: 87
Return Code from DB2 install is: 87
Found failing return code in db2inst.log file.
DB2 Install was NOT successful.
```

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

Uninstallation with the installation wizard

When you use the installation wizard for uninstallation, you might face certain issues with the uninstallation. Use the explanations and solutions to understand the cause of the problems and to resolve them.

Product directories still exist after uninstallation

When you uninstall with the installation wizard, you must ensure that the product is removed completely after uninstallation.

The `installation_path/config` and `installation_path/var` directories might exist even after the successful uninstallation of features.

You must manually remove the files and subdirectories that remains in the IBM Security Directory Server installation location after the uninstallation. The removal files and subdirectories are required because during the subsequent installation with IBM Installation Manger, the program generates an error messages if files or subdirectories exist in the IBM Security Directory Server installation location.

Chapter 5. Migration issues

Migration is the process of installing IBM Security Directory Server, Version 6.3.1 to replace an earlier version. During migration, the changes that were made to the data, schema definitions, and directory server configuration are preserved from the earlier version. Use the troubleshooting information to resolve issues that are related to migration.

Migration log files

The migration process is recorded in log files that you can use to trace errors.

Check the following log files for information about migration processes:

On AIX, Linux, and Solaris platforms:

Errors that occurred during migration are logged in the `/var/idsldap/V6.3.1/idsadm.log` file.

On Windows platforms:

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

Kerberos service name

In IBM Security Directory Server, Version 6.1 and later, you must find the service name that was used by the previous version of the server.

You must look for `ldap` in the keytab file in which an LDAP service name was located.

To avoid any problem that might arise because of this issue, you can take one of the following actions:

- Generate a keytab file by adding a lowercase LDAP Kerberos service name and use the new keytab file to communicate.
- Set the environment variable `LDAP_KRB_SERVICE_NAME`. If this environment variable is set, then the Kerberos service name is used with LDAP server service name in the keytab file and to communicate with its clients. You must also set the environment variable on the client side to continue to use the uppercase LDAP service name to communicate with its server. If the environment variable `LDAP_KRB_SERVICE_NAME` is not set, then the Kerberos service name is used with `ldap`.

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

If you are using the Instance Administration Tool to migrate, you might encounter an error that is related to database migration. Follow the steps to resolve this error.

If 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 that the database instance or database is not present and migration cannot continue.

To recover from this problem, take one of the following actions:

- Comment out the database information from the configuration file and migrate by 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 cannot 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 is 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 the **ibm-slapdPlugin** configuration files were changed, you might encounter an error during migration. Follow the steps to resolve the error.

If a line in the `ibmslapd.conf` or `slapd32.conf` file for the **ibm-slapdPlugin** was 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 is not able to load `/usr/ldap/lib/libback-rdbm.so` at startup because the path is not a valid path.

Considerations for migration

Before you plan to migrate from a previous version of IBM Security Directory Server, ensure that you consider the prerequisite conditions.

When you are migrating from a previous version of IBM Security Directory Server to a current version, you must consider the following conditions before migration:

- Specify a valid backup directory, encryption key, and encryption salt for an instance that you intend to migrate.
- Ensure that the required files are available in the backup directory.
- The source and target versions of IBM Security Directory Server are supported versions for migration.
- Ensure that the `ldapdb.properties` file in the `DS_instancehome/etc` is present and valid.

If error conditions are encountered during migration, the tool attempts to restore schema and configuration files to their original state and exits with error messages. Some of the likely causes for exit of migration are:

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

To identify the reasons for migration failure, you must examine the following messages and trace output:

- Use the error messages that are displayed on the console when the migration fails to determine the cause of failure.
- Run migration with trace mode ON and redirect the trace output to a file. This trace output can be used to determine the reason for migration failure.
- If an `ibm-slapdInvalidLine` message is displayed during server startup, then you must check whether the configuration file is corrupted.
- The `idsdbmigr.log` file must be checked for any failure in database migration.
- If the server fails to start after migration, you must set the trace mode ON. The following files must be examined to determine the cause of failure.
 - trace file (output from trace)
 - `db2diag.log`
 - `db2cli.log`

Chapter 6. Instance creation and configuration issues

Use the descriptions of instance creation and configuration options and instructions for avoiding common problems to identify and resolve related issues.

If you install IBM Security Directory Server Proxy Server or IBM Security Directory Server Full Server, no directory server instance is created by default. IBM Security Directory Server Full Server requires instance creation and configuration after the installation.

The troubleshooting steps for instance creation and configuration-related errors are provided in the following topics.

Instance creation overview and common errors

Use the overview, descriptions, and instructions to identify and resolve the possible errors that you might encounter with instance administration.

Instance creation overview

After you install a server, you must have a user account on the system and then create a directory server instance. Use the Instance Administration Tool or the command-line interface to create a user and the directory server instance.

You can create a user on a system either by using the Instance Administration Tool (**idsxinst**), which has a graphical user interface (GUI), or the **idsadduser** command-line utility. You can then create an instance by using either the Instance Administration Tool (**idsxinst**) or the **idsicrt** command-line utility.

When you create a directory server instance, a database instance is also created if the IBM Security Directory Server 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 Security Directory Server 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 installation wizard to install, the Instance Administration Tool runs. If you did not use the installation wizard to install, you must run the Instance Administration Tool or use the **idsicrt** command-line utility.

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

- Create the directory server instance.
- Set the IBM Security Directory Server administrator distinguished name (DN) and password. This operation can be compared to defining the root user ID and password on AIX, Linux, and Solaris systems.
- Configure the database, unless the server is a proxy server. Ensure that you first create the user ID for the database owner. You are not required 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 specific instance
- Delete an instance
- Migrate a server from a previous release to an IBM Security Directory Server, version 6.1 instance or later versions of directory server instance

Cannot create another instance because of invalid IP address

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

For example, assume that you have IP addresses 9.42.40.67 and 9.42.40.125 configured. 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 caused by one of the following reasons:

- 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 doing a DNS lookup. The setting in the operating system Name service switch file must be changed to do Host IP resolution lookup before it goes 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.

idssethost command does not recognize second IP address

On a system with `sles10`, the **idssethost** command fails to recognize the second IP address. Follow the instructions to resolve this problem.

On a system with `sles10`, you must add the configuration for the directory server to support multiple IP addresses. Add the IP addresses in the configuration file under the entry `cn=Configuration` as:

```
ibm-slapdIPAddress: IP_address1
ibm-slapdIPAddress: IP_address2
```

Restart the directory server. The server now listens at the IP addresses specified in the configuration file.

Note: You can provide any number of IP addresses.

Two directory instances 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 that is configured for "all" and another that is configured for a specific IP address can use the same port.

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

Instance creation fails on Windows 2003

On Windows 2003, instance creation might fail if the user password does not meet the operating system password requirements. Understand the cause of this problem and follow the steps to resolve it.

Instance creation fails on Windows 2003

Reason:

After the installation of IBM Security Directory Server, you can create an instance with default options. Instance creation and configuration utilities of IBM Security Directory Server might take a password that does not meet the operating system password requirements.

On Windows 2003, the configuration fails during the instance owner creation stage with the following messages. These messages are seen on the console and are also available in the `idsadm.log` file of directory server:

```
Jul 16 10:21:42 2008 You have chosen to perform the following actions:
```

```
Jul 16 10:21:42 2008 GLPGRP019I System user will be created for directory server instance.
```

```
Jul 16 10:21:42 2008 GLPGRP020I The system user 'idsinst' will be created.
```

```
Jul 16 10:21:42 2008 GLPGRP026I The user 'idsinst' will be a member of the 'administrators' group.
```

```
Jul 16 10:21:42 2008 GLPGRP005I The password for user 'idsinst' will be set.
```

```
Jul 16 10:21:42 2008 GLPGRP002I Creating system user 'idsinst'.
```

```
Jul 16 10:21:42 2008 GLPGRP006I Setting the password for user 'idsinst'
```

```
Jul 16 10:21:43 2008 GLPGRP043E Failed to create user 'idsinst'. Error code return by 'NetUserAdd API' is 2245.
```

```
Jul 16 10:21:43 2008 GLPGRP010W The program did not complete successfully. View earlier error messages for information about the exact error.
```

Solution:

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

The instance creation can also fail for the following reasons:

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

Administration server fails to start after instance creation

The Administration server might fail to start after the creation of a proxy server instance. This problem is specific to a 32-bit Windows 2008 operating system that is installed on 64-bit hardware.

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

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

This behavior is observed when a proxy server instance is created with the command **idsicrt -x** on a 32-bit Windows 2008 operating system, which is installed on a 64-bit hardware platform. In addition, the following behaviors were observed for the administration server with this configuration:

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

Configuration overview and common errors

Use the overview, descriptions, and troubleshooting instructions to identify and resolve the possible errors that you might encounter during configuration.

Configuration 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 configuration tasks.

The Configuration Tool has a graphical user interface (GUI) and it can be used for the following tasks:

- Setting or changing the IBM Security 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

Incorrect status is returned for files

Interruption of the Configuration Tool database tasks causes an incorrect status for the files. Follow the steps to recover from this problem.

When you are using the Configuration Tool to configure, unconfigure, import, export, back up, restore, or optimize a database, the process might be interrupted. For example, the interruption might be caused by a segmentation fault. When the process is interrupted, the status of the files is returned incorrectly. When you try to restart the process, the following message is displayed:

Task is already running.

This error is because the status output for the process is monitored through files in the `idsldapd-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 `idsldapd-instance_name/tmp` directory (where *instance_name* is the instance name).

Existing database instance and database configuration failure

When you configure an existing database and database instance with the **idscfgdb** command, a core dump might occur after the configuration is completed. This problem is specific to AIX, Linux, or Solaris operating systems. You can ignore this failure because the database is successfully configured.

Error occurs when Configuration Tool is started on AIX

An error might occur when you start the Configuration Tool on AIX. Follow the steps to resolve this error.

For example, the following error might occur when you start the Configuration Tool on AIX:

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

If this error occurs, ensure that you meet the following requirements:

- You have a supported version of DB2. For more information about supported versions of DB2, see the *IBM Security Directory Server Installation and Configuration Guide*.
- You have 64-bit hardware.
- You are running a 64-bit kernel.
- Your database migration to 64-bit is completed.

Configuration programs terminate on AIX

On an AIX system, the Configuration Tool might terminate as soon as you start it. Follow the steps to resolve this problem.

If the Configuration Tool terminates immediately after you start it, check the *LIBPATH*.

If the *jre/bin/classic* directory of a JVM other than the one provided with IBM Security Directory Server comes before the *\$LDAPHOME/java/bin/classic* directory, take one of the following actions:

- 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 is not configured properly

A failure might occur during DB2 database configuration. Understand the probable causes of the problem and follow the steps to troubleshoot.

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

One of the following reasons might be the cause of a failure that occurs during database configuration:

- 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 whether 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. This cleanup 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 configuration file attributes are changed

After you change configuration file attributes, the server might not start. Understand the reason and follow the steps to resolve this issue.

The attributes that are defined in IBM Security 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 more information, see the Web Administration Tool help documentation under **Reference > Directory Schema**.

Transaction log is full

If the schema defines too many attributes, you might get an error that the transaction log is full. Follow the steps to resolve this issue.

The following messages might be displayed at IBM Security Directory Server startup:

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

You might be required to increase the DB2 transaction log sizes. Type the following commands:

```
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

You might encounter various problems when you use the Configuration Tool, such as truncated titles and error messages. Follow the steps to resolve these issues.

The following sections describe the different problems that might occur on the Configuration Tool windows.

Translated titles might be truncated 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 according to 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 **Browse**, 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 you use keyboard

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

NullPointerException exception when you exit the Configuration Tool

If you exit the Configuration Tool after you enter an invalid database name, a `NullPointerException` exception occurs in the command window where the `idsxcfg` command was run. 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.

Configuration issues

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 are installed for it.

Note: DB2 comes in a wide variety of packages: Personal Edition, Enterprise Edition, Extended Enterprise Edition, and others. Many of these packages are supported across several versions of DB2, and each version can have several available fix packs.

- Amount of disk space available in affected drives and partitions.
- Third-party software that alters commonly used environment variables.

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

Standard sources of output

There are several "standard" sources of information available:

- The output on the screen

All of the configuration programs are either started from a console command-line prompt or open a background console. As the database configuration progresses, status messages (and limited error messages) are displayed in the associated console window. If a problem occurs, copy these messages to the system clipboard and then save them in a file for the IBM Software Support teams.

- DB2 log files

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

On Windows systems, examine any DB2 error logs in your DB2 installation directory. The error logs are under the directory that is 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:\sql11b`, examine the files in the `D:\sql11b\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. DB2 issues

When you use IBM Security Directory Server, you might face issues related to installation, configuration, and migration of DB2 database. Use the troubleshooting information to resolve issues that are related to DB2.

DB2 license file expired

Understand the cause of problems that are related to your electronic DB2 license and follow the steps to fix them.

If you see the following message during DB2 or server startup, there might be a problem with your electronic DB2 license:

GLPCTL010E Failed to start database manager for database instance: *instance name*.

To verify the cause of the problem, type the following command at the command prompt:

```
db2start
```

If your license is correct, you see the message:

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

Otherwise, you see a message that your license is expired or is going to 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. You must upgrade your DB2 product from a demonstration license to a product license. Copy the license file from the DVD to the system where DB2 is installed. You are not required to reinstall DB2.

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

- If you have a DVD: *mount_point/db2_installable_directory/license/db2ese_o.lic* (or *cdrom_drive:\db2_installable_directory\license\db2ese_o.lic* for Windows)
- If you downloaded a .zip file for installation:
directory_where_file_was_unzipped\sdsV6.3.1\db2_installable_directory\license\db2ese_o.lic
- If you downloaded a .tar file for installation:
directory_where_file_was_untarred\sdsV6.3.1/db2_installable_directory/license/db2ese_o.lic

Note: Your Proof of Entitlement and License Information booklets identify the products for which you are licensed.

2. After you have a valid license file on the system, run the following command to activate the license:

```
db2licm -a license_filename
```

- You 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 first. Then, install the new one that you 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.

Recovery from migration failure in DB2, version 9.1 or later

Direct recovery from migration failure is not available with DB2, version 9.1 or later. You can use the **idsdbmigr** tool, which uses DB2 backup and DB2 restore mechanism to recover from migration failure.

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

See the section “Overview of online backup and restore procedures for IBM Security Directory Server” in *IBM Security Directory Server Administration Guide* to know more about DB2 backup and restore.

DB2 9.5 installation on Linux operating systems

An error might be displayed when you install DB2 9.5 on the following operating systems: Red Hat Enterprise Linux (RHEL) version 5 64-bit or SuSE Linux Enterprise Server (SLES) version 10 operating system for Intel Linux or Linux for System z. Follow the steps to resolve this issue.

Problem

When you install DB2 9.5 on RHEL 5 64-bit or SLES 10 operating system for Intel Linux or Linux for System z, an error message is displayed. For example,

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

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

Solution

When you install DB2 9.5 on RHEL 64-bit or SLES 10 operating system for Intel Linux or Linux for System z, you must consider the following requirements:

- DB2 9.5 requires RHEL 5 GA or later versions.
- Install all available compat-libstdc++ RPM-packages before you install DB2 9.5 and ensure that libstdc++.so.5 is available. This installation is required for DB2 servers and clients.

To know more about installation requirements for DB2 servers, see the IBM DB2 documentation at <https://publib.boulder.ibm.com/infocenter/db2luw/v9r5/topic/com.ibm.db2.luw.doc/welcome.html>.

DB2 diagnostic information in db2diag.log

The db2diag.log file contains DB2 diagnostic information. A user with appropriate privileges can set the fully qualified path for DB2 diagnostic information by using the **diagpath** parameter.

If the **diagpath** parameter is null, the diagnostic information is written to the files in the following directories.

For DB2 v9.5

On Windows platforms

- In Windows Vista environment, the DB2 diagnostic error logs are written to the ProgramData\IBM\DB2\ directory.
- In Windows 2003 and XP environment, the DB2 diagnostic error logs are written to the Documents and Settings\All Users\Application Data\IBM\DB2\Copy Name\instance, where Copy Name is the name of DB2 copy.

On AIX, Linux, and Solaris platforms

The DB2 diagnostic error logs are written to *INSTHOME*/sqllib/db2dump, where *INSTHOME* is the home directory of the instance.

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

SQL0964C error when large amount of data is loaded

An SQL0964C error, which indicates that the transaction log is full, might be displayed when you load large amounts of data from a file. Follow the steps to troubleshoot and resolve this problem.

When you load a file that contains many entries, you might receive the following error message, which indicates that the transaction log is full:

```
SQL0964C  SQLSTATE=57011
```

You can troubleshoot this error by increasing the transaction log size. Complete the following steps to increase the transaction log size:

1. At command-line issue the following command and the password for the user:

```
su - db2instownername
```

where, *db2instownername* is the name of the DB2 instance owner.

2. Determine the current log file size setting by issuing the command:

```
db2 get db config for db2instancename | grep -i logfilsiz
```

3. Increase the size of the log file size setting by issuing the command:

```
db2 UPDATE db cfg for db2instancename using LOGFILSIZ new_value
```

4. Stop the slapd process.

5. To apply the changes that are related to database, issue the following command:

```
db2 force applications all
```

6. Restart the slapd process.

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

Instance starts in config-only mode after DB2 fix pack

An IBM Security Directory Server instance might start in config-only mode after you apply a DB2 fix pack. Follow the steps to resolve this issue.

You might get an error or your directory server instance might start in the config-only mode after you apply a DB2 fix pack. You must follow the post-installation instructions that are provided in the current DB2 fix pack readme file to resolve this problem.

You must also check the `ibmslapd.log` and `db2cli.log` files for the error descriptions that might be logged.

The following example error messages might be displayed in the `ibmslapd.log` file after you apply the DB2 fix pack:

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

The following example error messages might be displayed in the `db2cli.log` file after you apply the DB2 fix pack:

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

To know more about the post-installation steps after you apply a DB2 fix pack, see Applying DB2 fix packs.

Chapter 8. Web Administration Tool and application server issues

The IBM Security Directory Server, Version 6.3.1 Web Administration Tool is installed on an application server, such as embedded WebSphere Application Server. Embedded WebSphere Application Server is included with the IBM Security Directory Server and administered through a console. WebSphere Application Server can also be used as the application server. Use the troubleshooting information to resolve issues that are related IBM Security Directory Server Web Administration Tool and application server.

Corruption of data entered in the Web Administration Tool

At times, you might see the data that you enter in the Web Administration Tool is corrupted. Follow the steps to resolve this issue.

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

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 that is 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 Application Server administrative console tree:

- Select **Servers**.
- Select **Application Server**.
- Select the server that you want; for example, `server1`.
- Click **Process Definition**.
- Click **Java Virtual Machine**.

- Click **Custom Properties**.
- Click to create a 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.

Migration of files before you patch or migrate Web Administration Tool

When you patch or migrate the Web Administration Tool, you must first back up the specified files before you uninstall.

You must back up the following four files before you uninstall the IDSSWebApp.war file (the Web Administration Tool) and restore them after you reinstall the WAR file:

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

Additional login panels fail

If you open more login panels in the same browser instance, it might result in a failure. Follow the steps to open more login panels.

When you use the Web Administration Tool, do not open more 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 cookies. Additional login panels must be opened from new instances of the browser.

For AIX, Linux, and Solaris systems:

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

```
mozilla &
```

For Windows systems:

- Internet Explorer - Open more Internet Explorer windows by 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.

Web Administration Tool in inconsistent state

When you use the **idsldapmodify** command, it might put the Web Administration Tool into inconsistent state. Follow the steps to resolve this issue.

If you are logged in to the Web Administration Tool and use the **idsldapmodify** command from the command line to change your password, the Web Administration Tool changes the server status to stopped. This status change occurs because the Web Administration Tool opens new connections to the server every time it starts a task. The Web Administration Tool tries to connect to the server with the old password because it is unaware that the password changed. Hence, 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 you work with the Web Administration Tool.

Incorrect language is displayed in Web Administration Tool

The tabs, table headers, and static list boxes in the Web Administration Tool are sometimes displayed in an incorrect language. This problem might be encountered only on the AIX operating systems. However, Solaris and Linux systems might encounter the same problem. Follow the steps to resolve this 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. It is the language that was used the first time that a user logs in to the Web Administration Tool application. If you change the locale on your system, 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, run **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

You can resolve problems that occur when you run the Web Administration Tool with Microsoft Internet Explorer by changing the cache setup.

Make 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 you use 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 that comes from the server are not displayed correctly in the HTML page. This display problem is because of the way that the HTML is rendered by the web browsers.

For example:

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

This display problem affects fields such as labels, list boxes, tables, and captions.

Web Administration Tool requires IBM JDK on Domino server

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

The following limitations apply to 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. To view entries, you must select the option with the plus sign (+) next to it and click **Expand**.

Templates with object class that has no attributes

The Web Administration Tool does not save templates that are created with an object class that has no attributes. You must use the object classes that contain specified attributes to create templates.

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

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

Non-editable fields are displayed as editable

When you use Ctrl+L to view links, non-editable fields might be displayed as editable. Data that are entered in these fields is not saved.

If you open the Web Administration Tool by using Home Page Reader Ctrl+L keystroke to view the links on a Web Administration Tool page, non-editable fields might be displayed as editable. A text box might be displayed next to the non-editable field. Although you can enter data in the non-editable fields, the data is not saved.

Back and Forward buttons not supported

The **Back** and **Forward** buttons on Internet browsers are not supported for the Web Administration Tool. You cannot use them to navigate the Web Administration Tool.

Log on issues in Internet Explorer

When you log on to the Web Administration Tool console in Internet Explorer, you might encounter errors. Follow the steps to resolve or avoid the problems.

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

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

If these conditions exist on your computer, use localhost when you log on to the Web Administration Tool. You can avoid errors by using localhost instead of an IP address or host name when you log on to the Web Administration Tool.

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

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

Web Administration Tool issues on Windows Server 2003

You might encounter some errors when you use the Web Administration Tool console on the Windows Server 2003 platform. Follow the steps to avoid or resolve these errors.

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

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

To avoid these errors:

1. If embedded WebSphere Application Server is running locally, add `http://localhost` to the list of trusted sites.
2. If embedded WebSphere Application Server is running on a remote system, add the IP address or host name of the remote computer to the list of trusted sites. For example: `http://IP address` or `http://hostname`.

To add a web address to the list of trusted sites:

1. Click **Tools > Internet Options > Security > Trusted Site > Sites**.
2. Type the web address in the website 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`

Web Administration Tool logon fails for new user

A new user might fail to log on to the Web Administration Tool for the first time. This problem occurs if the password policy is enabled and the **User must change password after reset** option (`pwdMustChange`) is set. Follow the steps to resolve this issue.

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

To resolve the problem, the user can use the `1dapchangepwd` command-line utility to reset the password and then use the new password to log on.

Web Administration Tool backup creates another folder

When you backup by using the Web Administration Tool to a backup location that is specified in an NLV string, another folder gets created. Follow the steps to resolve this problem.

A problem occurs if the locale of the browser with Web Administration Tool differs from the locale of the system with the directory server instance. An NLV string folder also gets created apart from the backup folder when the backup operation is initiated.

This problem occurs because the string entered as the backup location is used as a file path. The string must be representable in the local code page of the system. The Web Administration Tool attempts to translate the Unicode input to the local code page to create the file path. It encounters Unicode input characters that cannot be represented to the system locale, which causes this problem.

Embedded WebSphere Application Server - Express on AIX

An error might occur when you start the embedded version of WebSphere Application Server - Express on AIX. Follow the steps to resolve this issue.

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 that is 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. Replication issues

When you use IBM Security Directory Server, you might encounter errors during replication. Use the explanations and information to troubleshoot and resolve issues that are related to 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. Replication improves the response time for the request completion.
- Security and content filtering - Replicas can contain subsets of the data in a supplier server.

For more information about replication, see the *IBM Security Directory Server Administration Guide*.

Diagnosis of replication errors

To identify the source of replication errors, you must understand the replication topology, know how to monitor replication status, and view replication logs and messages.

Sample replication topology

Use the example of a basic replication topology to set up your replication topology correctly.

If you are not sure whether your topology is set up correctly, you can compare it against this example. This topology assumes that there is a suffix in the server configuration for o=sample.

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

```
version: 1

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

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

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

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

```

objectclass: ibm-replicaGroup

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

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

```

Add the example file to masterhost with following command:

```
ldif2db -r yes -i in
```

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

```
db2ldif -o out
```

The server configuration file for masterhost must contain:

```

dn: cn=Configuration

ibm-slapdServerId: masterhost-389

```

The configuration file for replicahost must contain:

```

dn: cn=Configuration

ibm-slapdServerId: replicahost-389

```

and the following entry

```

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

```

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

```

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

```

Replication status

You can use the **idsldapsearch** command to get replication status information. Use the operational attributes to search for various details that can help you monitor the replication status.

Note: The following **idsldapsearch** examples are based on the sample replication topology that is provided in the topic, “Sample replication topology” on page 59.

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** object class 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 whether the subtree was quiesced; for example:

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

The remainder of the status-related operational attributes is 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 that indicates the replication state for all the replication agreements:

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

The available attributes are:

ibm-replicationLastActivationTime

The time that the last replication session started between this supplier and consumer.

ibm-replicationLastFinishTime

The time that the last replication session finished between this supplier and consumer.

ibm-replicationLastChangeId

The change ID of the last update that is 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 is suspended or "held".

Error log full

The replication errors that occurred are more than can be logged. The number of errors that can be logged is based on the configured value for `ibm-slapdRep1MaxErrors`.

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 that is 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 you request 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 that are logged for the specified replication agreement.

ibm-replicationFailedChanges

Lists the IDs, DNs, update types, result codes, timestamps, numbers of attempts for failures that are logged for a specified replication agreement.

ibm-replicationperformance

Give the operation counts per connection for multi-threaded replication.

Viewing replication errors with the Web Administration Tool

You can use the Web Administration Tool to 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.

Procedure

1. Log in to 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 **Show topology** on the table.
4. Click **View errors**.

From the View errors panel, you can:

- View the details of a specific error in the replication agreement.
- Attempt the selected replication update again.
- Attempt 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:

- a. Select the replication error that you want to view from the **Replication error management** table.
- b. Click **View details** on the toolbar. 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 that is 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 that is associated with the failed update, including all the operational attributes.

Controls

The controls that are used during the update.

Viewing replication errors with the `idsldapsearch` command

You can use the `idsldapsearch` to display replication errors. Viewing this information can help you identify the source of your replication problem.

About this task

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

- `ibm-replicationFailedChanges`
- `ibm-replicationFailedChangeCount`

Procedure

1. Use the `idsldapsearch` command to display replication errors:

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

This command returns an output similar to the following output:

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

2. Use the `idsldapexop` command to show data for the update, try the update again, or remove the update from the replication error log. Use the following `idsldapexop` command to show data for the failed update:

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

This command returns an output similar to the following output:

```
dn: entry-85,o=sample
cn: entry-85
objectclass: person
objectclass: eperson
objectclass: organizationalperson
objectclass: inetorgperson
objectclass: top
```


Lost and found log

The lost and found log (`lostandfound.log`) archives entries that are replaced because of replication conflict resolution. You can use the log of these entries to recover the data in the replaced entries, if necessary.

The information that is 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 that is associated with the failed update, including all the operational attributes, such as **ibm-entryUUID**.

Write and replicated write messages

To identify replication conflicts and related issues, you must understand the difference between write and replicated write messages.

In an IBM Security Directory Server environment, you might see informational message, such as

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

This message in the `ibmslapd.log` file indicates that a directory server is participating as a peer server in a replication network of directory servers. A peer server can receive updates from other peer servers and from client applications. A stand-alone master server shows zero updates from other suppliers. However, it can have some updates from clients, which depend on the update activity in a specified hour. A directory server that is configured only as a replica shows some updates from suppliers and zero updates from clients. Updates that are sent to such a replica that were referred to a master server are not counted as updates from clients.

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

ibm-replicaSubentry object class in a replication topology

Understanding the behavior of the `ibm-replicaSubentry` object class (`ReplicaSubEntry`) in a replication topology can help you identify and troubleshoot replication errors.

When a directory server that is in a replication environment starts, it compares its server ID against the server IDs in the `replicaSubEntry` entry. If the server ID matches, then as per the attribute value of **ibm-replicationServerIsMaster**, the

server either plays the role of a supplier or consumer. If the server ID does not match, the server assumes that it is a consumer in a replication topology.

If `replicaSubEntry` is defined, then the respective server ID provided with the attribute **ibm-replicaServerId** becomes supplier or consumer, which depends on the attribute value of **ibm-replicationServerIsMaster**.

For example:

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

Here `replicaSubEntry` means that the server with the server ID `server1` is a supplier server in the replication topology.

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

Here `replicaSubEntry` means that the server with the server ID `server2` is a consumer server in the replication topology.

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

Command-line utilities to view replication status

To determine issues that are related to replication, it is important to view the status of a replication agreement. You can use command-line utilities to view the appropriate operational attributes that are associated with replication.

The two special attributes, `+ibmrepl` and `++ibmrepl`, are defined to request replication-related operational attributes in a search. The `+` and `++` are subsets of the operational attributes. The single `+` is less expensive. The `++` includes all operational attributes that are shown in the `+` attribute list and the attributes that are listed in the `++` column as shown in the table.

Table 8. Replication-related operational attributes

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

To search a specific replication agreement, issue the **ldapsearch** command in the following format:

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

For example,

```
idsldapsearch -h peer1 -p 1389 -D cn=root -w password -b cn=peer2:2389,\  
cn=peer1:1389,ibm-replicaGroup=default,0=sample objectclass=* ++ibmrepl
```

To search all agreements, issue the **ldapsearch** command in the following format:

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

For example,

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

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

IBMSLDAPD_REPL_UPDATE_EXTRA_SECS environment variable

You can use the *IBMSLDAPD_REPL_UPDATE_EXTRA_SECS* environment variable to extend the time duration for update operations in replication.

The default timeout for any change to be completed through replication is 60 seconds. The replication updates might involve many changes, such as adding a large group entry or adding or modifying entries that contains large objects such as credentials. In such cases, the update operation might require more than 60 seconds to finish. Update operations include add, delete, modify, or modifydn operations. If any such single update operation through replication takes more than 60 seconds, the supplier server times out that update operation. It tries again to send the same update through replication. To extend the timeout duration for update operations in replication, you can use the *IBMSLDAPD_REPL_UPDATE_EXTRA_SECS* environment variable.

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

- Minimum: 1
- Maximum: 2147483647

Note: For optimal result, a value of 180 is preferred for the variable. Setting the variable with a value greater than 600 is not preferred. Test the same update from a direct client against the consumer server from the supplier server. In this way, you can determine a value that is best suited for your environment.

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

Adding the variable to configuration file

Using the LDAP client utility:

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

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

2. Restart the directory server instance.

Using the Web Administration Tool:

1. Ensure that ibmslapd and ibmdiradm processes are running for the directory server instance.
2. Log on to the directory server instance by using the Web Administration Tool.
3. From the left navigation area, expand **Directory management** and then click **Manage entries**.
4. On the **Manage entries** panel, expand **cn=configuration**, and then select **cn=Front End** and click **Edit attributes**.
5. On the Edit and entry panel, click **Next** to open the Optional attributes panel.
6. Click **Multiple values** next to the **ibm-slappedSetenv** field.
7. In the resulting panel, enter **IBMSLAPD_REPL_UPDATE_EXTRA_SECS=180** in the **ibm-slappedSetenv** field and then click **Add**.
8. To save, click **OK**.
9. To effect the changes that are made, restart the directory server instance.

From command prompt

Set the environment variable *IBMSLAPD_REPL_UPDATE_EXTRA_SECS*.

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

```
export IBMSLAPD_REPL_UPDATE_EXTRA_SECS=180
```

On Windows systems:

```
set IBMSLAPD_REPL_UPDATE_EXTRA_SECS=180
```

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

Information for troubleshooting replication

Use the troubleshooting information to identify the cause of various replication issues and resolve them.

Replicated suffix

The replicated suffix must have the *ibm-replicationcontext* object class. Set the object class before you load your data in the database.

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

```
08/13/04 15:32:34 For the replica group entry
ibm-replicaGroup=default,o=sample, the parent entry
must be an ibm-replicationContext entry.
08/13/04 15:32:34 Parent entry does not exist for entry
cn=urchin,ibm-replicaGroup=default,o=sample.
```

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

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

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

where the `mod.ldif` file contains:

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

Verify that suffixes and replication agreements exist

If you experience errors with replication, use the `idsldapsearch` command to verify that your suffixes are configured to be replicated and that the replication agreements exist.

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

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

where `o=sample` is the replication context.

If this command does not return any results, the suffix is not configured to be replicated. You must configure the suffix to be replicated. See the *IBM Security Directory Server 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 repltxt
objectclass=ibm-replicationAgreement
```

where `repltxt` is the location where the replication agreements for a replication context are stored; for example, `o=sample`. If the command does not return results, the replication agreement might not exist. To replicate correctly, the correct replication agreements must exist. See the *IBM Security Directory Server Administration Guide* for instructions for adding replication agreements.

Peer-to-peer replication error

If you are running peer-to-peer replication, you might encounter the error "No such object occurred for replica." Follow the steps to resolve this problem.

You might see an error similar to the following error:

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

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

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

To correct this error, complete the following steps:

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

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

where:

- *out.ldif* is the name of the file to which you want to export the entry.
- *instance name* is the name of the instance.
- *subtree DN* is the DN of the entry you want to export.

3. Use the **idsldapadd** command to import the entry to the other peer server; for example:

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

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

Insufficient access error

When replication returns an insufficient access error, the reason might be related to the version of IBM Security Directory Server on the consumer. Follow the steps to troubleshoot and work around this problem.

Consider a scenario where a replication topology extended operation is issued to a server that is of IBM Security Directory Server, version 6.0 or later. However, if the consumer of the server has a release earlier than version 6.0, the operation fails. In the server trace, you can identify insufficient access as the cause of the failure.

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

Replication topology extended operation fails with result code 80

After you run a replication topology extended operation, you might see an error message that the operation failed with result code 80. There are several reasons why this error might occur. Complete the checks that are required to ensure that this error is resolved.

You might see following message after you run 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 error occurs, check for the following situations:

- If the replication context entry exists on the consumer server, be sure that the replication context entry has an object class of `ibm-replicationContext`.

Alternatively, 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.

- The supplier of the extended operation first sends all the replication topology-related entries under a replication context to the consumer. After that, the supplier sends the replication topology extended operation to the consumer to try to cascade the operation. There might be more than one tier of servers that are involved in a replication topology. In such cases, ensure that each supplier has the correct credential object to bind with its consumers.
- One of the consumer servers is down or not reachable at that instance.
- The replication context is a non-suffix entry and the consumer does not have the parent entry of the context.

For example, suppose that `cn=johndoe,cn=people,o=sample` is the context for the topology you want to replicate. If `o=sample` is the suffix on the consumer and `cn=people,o=sample` does not exist, the operation fails.

- The repl topology extended operation times out on a heavily loaded consumer. This problem results in the message GLPRPL098E.
- Suppose that a certain set of agreements exists on the consumer. The repl topology extended operation attempts to delete these agreements and before that attempts to purge the queue that is associated with that agreement. If the purge fails, the extended operation fails. This problem results in the message GLPRPL093E.

Replication command-line interface error

If you have a master server that is configured to do replication, you might see a command-line interface error. This error occurs only on Windows operating systems. Follow the steps to resolve this error.

You might see an error like the following error in the `ibmslapd` error log during updates:

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

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

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

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

Entries in LDIF file are not replicated

When you use the `idsldif2db` command with the `-r yes` option, you might find that entries are not being replicated. The `-r yes` option indicates that the entries in the file are to be replicated. Use the troubleshooting information to 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 that 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 situations:

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

Problem with cn=ibmpolicies subtree

A problem might occur with replicating or modifying the cn=ibmpolicies subtree, where this subtree becomes read-only or might not get replicated properly. Follow the steps to resolve this problem.

In IBM Security Directory Server, Version 6.0 and later, a partial replication configuration entry is automatically added to the cn=ibmpolicies subtree. The design of a directory server in version 6.0 and later versions have the ibm-replicationcontext, ibm-replicagroup, and ibm-replicasubentry setup automatically for the cn=ibmpolicies subtree the first time the LDAP server is started. Also, the partial entries are created with the default ibm-slappedServerId value (randomly generated when the instance is first created). As users often modify the **ibm-slappedServerId** value in the ibmslapd.conf file after the initial configuration, this subtree often become read-only or might not get replicated properly. To resolve this problem, you must consider removing these partial replication entries from all systems in the topology:

To remove the entries, complete the following steps:

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

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

```
dn: CN=IBMPOLICIES
cn: IBMpolicies
objectclass: container
objectclass: top
objectclass: ibm-replicationcontext
```

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

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


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

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

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

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

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

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

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

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

After you remove these entries, you can set up replication on the `cn=ibmpolicies` subtree just as you would on any other subtree.

Master server becomes unstable or stops

The master server might become unstable or stop when it serves a larger number of replica servers. The reason might be because the master server ran out of resources. Follow the steps to resolve this issue.

To resolve this problem, you can set the `Ulimits` DN entry in the configuration file as shown here:

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

Then, configure the system ulimit values to:

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

Restart the servers for the changes to take effect.

Stopping a multithreaded replication supplier

In a replication environment, abruptly stopping a supplier that uses multithreaded replication to accelerate replication between its consumers can cause problems. To avoid having replication-related errors, complete the specified steps before you stop a supplier server.

About this task

At any specified time, a supplier might send multiple updates to a consumer. The number of updates can be the number of consumer connections, which are multiplied by the depth of replication receive queue. A supplier can also have multiple consumers. In such cases, the number of replication updates at any specified time can be large. The replication status of the supplier is based on the most recent change replicated (successfully or otherwise) so far. If a supplier is restarted, it uses this replication status value to determine the changes that must be sent to the consumer. If a supplier is stopped abruptly before it receives a response from its consumers for the updates that it sent, the supplier sends the updates again. When these updates are sent again, it can cause errors to be reported. These errors are logged by the supplier. They can be cleared by using the Web Administration Tool or the command-line utility for managing the replication error log. If the logging of replication errors is enabled, these errors are counted against the limit for logged errors. Hence, they must be cleared.

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

In some cases, the administrator might not be aware of the problem and hence might not be able to resolve the problem. A replicated update from a supplier to its consumer might result in an error and the supplier might not be available to log the error. Depending on the last response that is received by the supplier, this update might not be replicated again.

Procedure

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

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

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

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

ibm-replicaServerId=wingspread-2389,ibm-replicaGroup=default,0=SAMPLE
```

3. Find all the replication agreements for the server.

```
#ldapsearch -D cn=admin -w password -p 2389 -b ibm-replicaServerId=wingspread-2389,\
  ibm-replicaGroup=default,0=SAMPLE objectclass=ibm-replicationAgreement 0.0

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

4. Verify the status for a specified replication agreement.

```
#ldapsearch -D cn=admin -w password -p 2389 -b cn=wingspread-1389,\
  ibm-replicaServerId=wingspread-2389,ibm-replicaGroup=default,0=SAMPLE \
  objectclass=* +ibmrep1
```

```

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

```

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

5. Suspend replication for the agreement.

```

# ldapexop -D cn=admin -w password -p 2389 -op controlrepl -action suspend \
  -ra cn=wingspread-1389,ibm-replicaServerId=wingspread-2389,\
  ibm-replicaGroup=default,0=SAMPLE

```

Operation completed successfully.

6. Verify the status of replication agreement.

```

# ldapsearch -D cn=admin -w password -p 2389 -b cn=wingspread-1389,\
  ibm-replicaServerId=wingspread-2389,ibm-replicaGroup=default,0=SAMPLE \
  objectclass=* ++ibmrepl

```

```

cn=wingspread-1389,ibm-replicaServerId=wingspread-2389,
  ibm-replicaGroup=default,0=SAMPLE
ibm-replicationChangeLDIF=N/A
ibm-replicationLastActivationTime=20080707152648Z
ibm-replicationLastChangeId=4855
ibm-replicationLastFinishTime=20080707152648Z
ibm-replicationLastResult=N/A
ibm-replicationLastResultAdditional=N/A
ibm-replicationNextTime=N/A
ibm-replicationPendingChangeCount=1
ibm-replicationState=on hold
ibm-replicationFailedChangeCount=0
ibm-replicationperformance=
  [c=0,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
ibm-replicationperformance=
  [c=1,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
ibm-replicationperformance=
  [c=2,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
ibm-replicationperformance=
  [c=3,l=10,op=0,q=0,d=0,ws=0,s=0,ds=0,wd=0,wr=0,r=0,e=0,ss=1,rs=1]
ibm-replicationPendingChanges=4856 modify CN=WINGSPREAD-1389,
  IBM-REPLICASERVERID=wingspread-2389,IBM-REPLICAGROUP=DEFAULT,0=SAMPLE

```

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

7. Determine whether there are any replicated updates that were sent to the consumer. In the output from step 6 related to the replication status, the attribute **ibm-replicationperformance** can be used to determine the number of

updates that were sent to the consumer. This attribute applies only to replication agreements by using multithreaded replications.

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

- c** Indicates the connection number. In the output of step 6, there are four connections. The first connection shows the most traffic. The workload determines how often the other connections are used.
- l** Indicates the size limit for each queue. In the example, the value is 10. For each connection, there are two queues of same length. One queue is for updates to be sent on the connection, which is called the send queue. The other queue is for updates that were sent but no response was received from the consumer, which is called the receive queue. When updates are sent, they are moved from send queue to the receive queue. When the receive queue reaches its size limit, no more updates are sent until some responses from the consumer are received. When the send queue reaches its size limit, no more updates are assigned to the connection. When the size limit for all the send queues of connections are reached, the supplier waits for the consumer to process the backlog.
- op** Specifies the replication ID of the last operation that is assigned to the send queue of the connection. Replication IDs are assigned to all updates that are to be replicated to a consumer. The process of assigning replication IDs must not stop even if replication is suspended.
- q** Specifies the current size of the send queue. This value must not change if replication is suspended.
- d** Specifies the count of dependent updates. An add request for an entry followed by a modify request of the same entry is counted as a dependency. All dependent updates must be assigned to the same connection so that they can be applied in correct order.
- ws** Indicates the number of times the send queue reached its size limit.
- ds** Specifies the number of dependent updates sent.
- wd** Specifies the number of times that the send queue waited for a dependent update before it sent more updates.
- wr** Indicates the number of times the receive queue reached its size limit.
- r** Indicates the number of replicated updates that is waiting for a response from the consumer.
- e** Specifies the number of replication errors reported by the consumer.
- ss** Indicates the session count of the sender thread. It is incremented when a connection to the consumer is established.
- rs** Indicates the session count of the receiver thread.

The value indicated by **r** might show that the number of replicated updates that are waiting for a response is 0. In this case, it is safe to stop the supplier for this consumer server. The value of **r** varies between 0 and the value of **l**, the size limit of the queue, which defaults to 10. If the value is not 0 for **r**, you must wait for it to be 0. The value of **r** depends on the size and type of the replication update and the workload on the consumer. After this value

becomes 0, the supplier sends the status of updates to the consumer. On restarting the supplier, it replicates only the updates that were not sent before.

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

```
#ldapexop -p 2389-D cn=admin -w password -op controlrepl -action resume \  
-ra cn=wingspread-1389,ibm-replicaServerId=wingspread-2389,\  
ibm-replicaGroup=default,0=SAMPLE  
Operation completed successfully.
```

Results

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

Synchronizing directory servers in a replicated environment

If directory servers in a replicated environment are out of synch, the replication queues might get blocked. To resolve this problem, you must resynchronize your replicated environment.

About this task

Consider a scenario where M1 is the master server with the most recent updated data. R1 and R2 are the two replica servers of the master server, M1. To resynchronize the directory servers, complete the following steps.

Procedure

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

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

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

```
idscfgdb -I instance_name -a dbadminDN -w dbadminPW -t databasename \  
-l dblocation -n
```
8. Synchronize the modified schema. Copy the V3.modifiedschema from M1 over to R1 and R2. The modified schema, V3.modifiedschema, is in the `instance_home/idsldapd-instance_name/etc` directory.

9. Synchronize the `ibmslapddir.ksf` file. To know more about Synchronizing two-way cryptography between server instances, see the *IBM Security Directory Server Administration Guide*.

Note: Only if the master and the replicas are on the same hardware and operating system, the `ibmslapddir.ksf` file can be copied over from master to replicas. The `ibmslapddir.ksf` file is in the `instance_home/idsslapd-instance_name/etc` directory.

10. Copy the `M1.ldif` file to replicas and load the data of M1 onto R1 and R2. Issue the following command:

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

Results

Note: On Windows platform, change the paths accordingly.

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

Multimaster configurations

The configuration must ensure that updates for the same entry or set of entries do not occur to several peer masters at the same time. The replication system can be configured in such a way that all writes go to one master except in the case of failover. Or, the system can be configured so that all writes for a specified subtree go to one master except in the case of failover.

If writes of the same entry occur on several masters, then before such write can be replicated, an update conflict might occur.

A directory server instance can be configured with conflict resolution. This configuration ensures that for almost all update conflicts, the latest change to a specified entry is preserved. It ensures that the content of all servers converges to the same value for the entry. However, update conflicts must be avoided. Conflict resolution might cause inherent loss of data because the later change to the entry is preserved but the earlier change is discarded. Conflict resolution can also affect replication performance, if the number of conflicts that are observed is large.

Sometimes, it is not possible to avoid configurations where update conflicts can occur. For example, there might be IBM Security Directory Server at two sites. Because of a temporary loss of network connectivity between the sites, all writes occurring at a specified site might occur on the server for that site. Update conflicts might occur as a result, and the IBM Security Directory Server conflict resolution procedures then converge the content of entries on the servers. However, in most configurations, nearly all update conflicts can be avoided.

If conflict resolution is used, the following condition must be ensured. The directory server must be loaded so that timestamps for the created entries are same on all servers in the topology at the outset. There are two ways to ensure this condition:

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

Options for replication filter and replication method are not available

When you create a master server in a replication topology, the options to specify replication filter and replication method are not available. This unavailability is a limitation of the Web Administration Tool. However, you can specify the filter and replication method options in a peer-to-peer replication.

The replication filter contains the existing filters under the selected subtree and the replication method specifies the type of replication, single-threaded or multi-threaded.

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

Consumer server that does not support SHA-2

Replication between an IBM Security Directory Server, Version 6.3 or later supplier and a consumer server with an earlier version that does not support SHA-2 cannot be started.

Consider the following scenario:

Replication is set up between an IBM Security Directory Server, Version 6.3 or later supplier server and a consumer server with an earlier version. The IBM Security Directory Server, Version 6.3 or later supplier server has SHA-2 family of encryption algorithm as the configured password encryption method. The consumer server with an earlier version does not support SHA-2 family of encryption algorithm.

In this case, the supplier logs a message that the replication operation cannot be started with the consumer and sets the replication state to `connecting`. Similarly, if attribute level encryption is set with SHA-2 family of encryption scheme on IBM Security Directory Server, Version 6.3 or later, replication to the consumer server fails to start and the replication state is set to `connecting`.

Chapter 10. Performance issues

If you are experiencing problems with the performance of your directory server, use this information to find possible fixes and workaround.

Identification of performance problem areas

Use the server audit log and **idsslapd** trace for identifying areas that might be affecting the performance of your directory server.

Server audit log

The server audit log shows the searches that are being done and the parameters that are used in each search. The server audit log also records when a client binds and unbinds from the directory. By observing these measurements, you can 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 turn on the trace, run the commands that you think might be giving you trouble.

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

```
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. Complete the steps to fully use the added memory.

About this task

To take advantage of added memory, you must complete the following steps.

Procedure

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 systems, if you have clients that are generating many connections to the server and the connections are being refused, set the `SLAPD_OCHANDLERS` environment variable to 15 before you start the server.

About this task

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

```
Feb 11 14:36:04 2004Communications error: Exceeding 64
connections/0CH - dropping socket.
```

If you see these errors, complete the steps in this procedure.

Procedure

1. Save a copy of your `ibmslapd.conf` file.
2. Insert the following in the section that starts with `dn:`
`cn=FrontEnd,cn=Configuration:`
`ibm-slapdSetenv: SLAPD_OCHANDLERS=15`
3. Stop and restart the server.

DB2 rollbacks and isolation levels

If you are experiencing rollback activities in DB2, check the isolation level. Rollbacks occur when one application process has a row that is locked while another application process tries to access that same row.

The default isolation level, repeatable read, can result in more rows that are locked than are required for the current read request. Hence, a more relaxed isolation level is normally required for LDAP applications.

For example, the read stability isolation level allows other applications to insert or update data in rows that were 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 to be returned on multiple reads, the isolation level must 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 that are using the current database instance are affected by this setting.

Default value of LOGFILSIZ must be increased

If you are adding a large group to your directory, you must modify the **LOGFILSIZ** parameter of your DB2 database.

Consider the following scenario: You might be adding a large group, for example, more than 50,000 members, to your directory. You migrated your database from a previous release. In this case, you must modify the **LOGFILSIZ** parameter of your DB2 database to be at least 2000. On migrated databases, this value might currently be set to 750 or 1000.

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

For AIX, Linux, and Solaris platforms:

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

To increase this value, issue the following command:

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

For Windows platforms:

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

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

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

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

Audits for performance profiling

To identify performance bottlenecks during operation execution, you can check the server audit log for the summary figures that indicate performance hotspots.

The following hotspots are identified for auditing:

- An operation waits in the worker thread queue for a long time before the worker thread actually starts the operation.
- The time that is spent for cache contention inside the back-end requires tracking.
- The time that is spent in handling client input and output, that is, the time that is spent in receiving the request and returning the result. This value can also be used for detecting bottlenecks because of slow clients or network issues.

Using the audited performance hotspot data, directory administrators can use the system audit facilities to log the LDAP audit record with the system-defined record format.

When you audit the performance profiling, you must consider the following points:

- The configuration options can be enabled to auditing for a combination of different types of operations. For example: auditing for add and modify operations only, along with the auditing for performance.
- At the end of operation execution, the audit information is stored in the server audit logs only. In a scenario where the server is having performance bottlenecks and is in a hung state, the `cn=workers, cn=monitor` search can be issued. This search gives information about where each worker is stuck. This information is obtained by accumulating information that is collected about the worker until that point in the audit records.

For each operation, performance data field in the audit records is controlled by using the configuration option **ibm-auditPerformance**. Currently, the following performance data fields are defined for each operation:

operationResponseTime

Represents the time difference in milliseconds between the time the operation was received and the time its response was sent. The operation received time and the response sent time of an operation are published in audit v3 header.

timeOnWorkQ

Represents time in milliseconds spent in the worker queue before execution is initiated on the operation. The value of this field is the difference between the time execution was initiated and the time the operation was received.

rdbmLockWaitTime

Represents time in milliseconds spent in acquiring locks over RDBM caches during operation execution. The value in this field helps administrators to determine the time that is spent for cache contention against real work.

The lock wait time over the following resources are also considered.

- Resource cache
- DN cache
- Entry cache
- Filter cache
- Attribute cache

Note: From IBM Security Directory Server, version 6.3 release, attribute cache is deprecated. Henceforth, users must avoid the use of attribute cache.

- Deadlock detector
- RDBM locks

This field is implemented by introducing of a field in the operation structure. This field is updated when the acquiring of lock is attempted during operation execution. In addition, wrapper functions are introduced for functions that attempt to acquire locks over RDBM caches. These wrapper functions take another operation pointer as parameter and update the lock wait time field of the operation if **ibm-auditPerformance** is enabled.

clientIOTime

Represents time in milliseconds that was spent in receiving the complete operation request and returning the complete operation response. This field

is implemented in the operation structure. It is updated on receiving the complete BER for operation request and on successfully returning the response BER message for the operation.

An example of the audit version 3 format for search operation with **ibm-auditPerformance** enabled is shown here:

```
AuditV3--2006-09-09-10:49:01.863-06:00DST--V3 Search--
bindDN: cn=root--client: 127.0.0.1:40722--connectionID: 2--
received: 2006-09-09-10:49:01.803-06:00DST--Success
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: o=sample
scope: wholeSubtree
derefAliases: neverDerefAliases
typesOnly: false
filter: (&(cn=C*)(sn=A*))
operationResponseTime: 591
timeOnWorkQ: 1
rdbmLockWaitTime: 0
clientIOTime: 180
```

To control server performance hits when information for performance data fields is collected, the **ibm-auditPerformance** field is introduced in the audit configuration section. The value of the **ibm-auditPerformance** field is false, by default and therefore no performance data is collected and published by default. When the value of the **ibm-auditPerformance** field is set to true, performance data is collected and published in the audit logs for each operation that is enabled to be audited. If the **ibm-auditPerformance** field is enabled, that is, set to true, in audit record section the four performance data fields are audited: **operationResponseTime**, **timeOnWorkQ**, **rdbmLockWaitTime**, and **clientIOTime**. The values of these performance data fields are times in milliseconds.

Chapter 11. Information for troubleshooting in various scenarios

When you use IBM Security Directory Server, there are several scenarios that you might encounter, which require troubleshooting. Use the solutions that are provided to resolve these problems.

Server is not responding

If the server seems to not respond, you must first verify whether the server is truly not responding, or its performance is too slow.

To determine whether the server is suffering from poor performance, see the *IBM Security Directory Server Performance Tuning and Capacity Planning Guide* for monitoring performance. Compare the operations that are initiated and operations completed values. Also, compare the values of add operations that were requested and completed to understand what is happening with the performance of your system.

If you determine that the server is not responding, run the **IBM Support Assistant Lite** tool. This tool gathers information that you can provide to IBM Software Support to help identify the problem. See “Tools for troubleshooting a directory server instance” on page 1 for information about the **IBM Support Assistant Lite** tool.

Memory leak is suspected

If a memory leak is suspected, run a script that gathers information about the memory sizes of the processes that are running on your system.

Note: The following script is an example for AIX systems. You must modify it for your operating system.

When the script finishes, send the monitor.out text file that is generated by the script to IBM Software Support for analysis.

The script is as follows:

```
#!/bin/sh
instance=ldapdb2
port=389
binpath=/opt/IBM/ldap/V6.1/bin

while [ true ]; do
echo | tee -a /tmp/monitor.out
echo 'Begin Monitoring.....' | tee -a /tmp/monitor.out
date | tee -a /tmp/monitor.out
echo 'Process info via ps auxw command: ' | tee -a /tmp/monitor.out
ps auxw | 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 return errors

If you are experiencing errors on SSL, use the **ldapsearch** command to verify that SSL is set up correctly.

Run the following command:

```

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 that is being searched; for example, -b o=sample

Record and send any errors to IBM Software Support.

Recovering data from a directory server instance where encryption seed value is lost

If an encryption seed value is lost for a directory server instance during an instance creation, then you cannot recover the lost encryption seed value. However, you can recover the data from the directory server instance for which the encryption seed value is lost.

The workaround is to create a directory server instance with a new encryption seed value and then use the **db2ldif** and **ldif2db** utilities to export and import data. You can supply the new encryption seed and salt value of the new instance to these utilities. The data would be preserved (along with the passwords) on this new instance. The steps to recover data on a Linux system are as follows:

1. Create a user for the instance. Issue the command of the following format:

```
idsadduser -u newinst -w newinst -l /home/newinst -g idsldap
```
2. Create and configure a new directory server instance. Issue the commands of the following format:

```

idsicrt -I newinst -e thisismynewencryptionseed -l /home/newinst -n
idscfgdb -I newinst -a newinst -w newinst -t newinst -l /home/newinst -n
idsdnpw -u cn=root -p root -I newinst
idscfgsuf -s "o=sample" -I newinst

```


Note: Save the encryption seed `thisismynewencryptionseed`.

3. After you set up the new instance, `newinst`, you must find and save the salt value that is generated by the directory server instance. To find the salt value, issue the command of the following format:

```
idsldapsearch -p port_number -D cn=root -w root -b "cn=crypto,cn=localhost" \
-s base objectclass=* ibm-slapdCryptoSalt
```

For example, consider the salt value of the new instance, `newinst`, as `newsaltvalue`.

4. To export data to an LDIF file from the directory server instance (for example, `oldinst`) for which the encryption seed is lost, use the **db2ldif** command of the following format:

```
db2ldif -o mydata.ldif -I oldinst -k thisismynewencryptionseed -t newsaltvalue
```

Note: After completion of this command successfully, the entire data from the directory server instance, `oldinst`, would be stored in the `mydata.ldif` file that is specified in the **db2ldif** command.

5. Finally, import the data from the LDIF file to the new directory server instance. Issue the **ldif2db** command of the following format:

```
ldif2db -i mydata.ldif -I newinst
```

Attribute encryption to be avoided for older versions

Attribute encryption must be avoided in an environment that includes versions of IBM Security Directory Server earlier than version 6.1.

Attribute encryption must not be used in an IBM Security Directory Server environment that includes server versions earlier than version 6.1. Storing encrypted attributes on some servers and not storing encrypted attributes on other servers defeats the purpose of encrypting attributes. In such an environment, for interoperability between servers you must not encrypt attributes. If attribute encryption is used in such an environment, the following situations might arise:

- Attempts to replicate schema definitions for encrypted attributes might fail because the target server does not recognize the new `IBMAttributeTypes` keywords.
- On a server earlier than version 6.1 and servers that do not have matching encryption keys, data is handled as described next. For an attribute that is defined but not encrypted, data is decoded for replication and stored in decoded format. If the servers have the matching keys, data is not decoded during replication. Hence, the data is rendered as useless on the servers that do not have matching keys.

In situations where RDBM startup encryption processing fails for a specified attribute, the processing can be skipped for the attribute by deleting or commenting the **ibm-slapdMigrationInfo** line from the configuration file for that entry from the RDBM. For example:

```
dn: cn=Directory, cn=RDBM Backends, cn=IBM Directory, cn=Schemas, cn=Configuration
...
#ibm-slapdMigrationInfo: encrypt secretAnswer
```

Character sets larger than 7-bit ASCII in passwords

There are certain limitations with use of character sets that are larger than 7-bit ASCII in a user password. Understand the limitations and how to work around them.

Portable characters (common character set) or 7-bit ASCII characters use the first 7 bits to form characters (128 characters, 0 through 127). These characters are used on most of the code pages. In IBM Security Directory Server, **userpassword** is a binary attribute and it is not converted from the client code page (for example, IBM-437, IBM-850, or Windows 1252) to UTF-8 and then back to the server code page like text attributes. Code pages differ from the portable character limitations. You might use non-portable ASCII characters (beyond the first 127 or 7-bit ASCII) in a user password. Then, the password matches only if it is provided from the same code page in which it was originally created.

For example, if you use the Web Administration Tool to create the password, `as12÷÷qw`, for the entry, `cn=Bob Garcia,ou=austin,o=sample`, and then if you do a search by using the **ldapsearch** command from the command line as Bob Garcia, the following results are displayed:

```
ldapsearch -D "cn=Bob Garcia,ou=austin,o=sample"\  
-w as12÷÷qw -b "o=sample" "objectclass=*"\  
ldap_simple_bind: Invalid credentials
```

This occurs because the Web Administration Tool and the command line use different code pages and the password `as12÷÷qw` contains non-portable characters. Therefore, ensure that the client always authenticates by using the same code page that was used when the password was created. Or else, you must limit passwords to portable characters (7-bit ASCII). Using non-portable characters is a permanent limitation.

Premature expiry of user password

If the password of a user expires prematurely, the reason might be because of timezone and daylight saving.

Comparisons pertaining to password policy, such as validation of maximum age of password (`pwdMaxAge`), are done in Coordinated Universal Time. However, in geographical areas that follow daylight saving, the password might expire prematurely or later than due time. Premature expiry happens because the timestamps are not monitored in Coordinated Universal Time. Users must convert the time in their timezone to Coordinated Universal Time. Then, if they do password expiry calculations, the password would expire at the expected time that was set.

Troubleshooting the limitation in the `idssethost` command

You might configure a directory server instance to listen on a specific interface by using the **idssethost** command or the Instance Administration Tool (**idsxinst**). The **idssethost** command configures the directory server instance to listen only on the specific IP address. Follow the steps to work around this limitation.

About this task

An entry is added to the `ibmslapd.conf` file:

```
ibm-slapdIpAddress: IP_address
```

However, this leads to unexpected behavior from the perspective of the user because the `ibmslapd` no longer listens on the loopback address (127.0.0.1). All LDAP client utilities that run locally attempt to connect to `ibmslapd` over the loopback interface. As a result, when the commands on the local system are run, they fail to contact the directory server, unless the `-h` option is used to point specifically at the interface that `ibmslapd` is listening on.

Additionally, the `idssethost` command does not allow configuring the directory server to listen on the loopback interface. Any attempt to do this configuration returns the following error:

```
idssethost -I ldapdb2 -i 127.0.0.1 -n
GLPCTL062E The specified IP Address '127.0.0.1' is not a valid IP address
for this machine.
```

For the current versions of IBM Security Directory Server, if `ibmslapd` is required to listen on a specific interface and the loopback interface, the directive to listen on loopback must be added manually. Perform the following steps:

Procedure

1. Add the IP addresses that you want the server to listen to. Issue the `ldapmodify` command of the following format:

```
idsldapmodify -p port -D cn=adminDN -w adminPW -f filename
```

where, *filename* contains:

```
dn: cn=Configuration
changetype: modify
add: ibm-ibm-slapdIpAddress
ibm-slapdIpAddress: 10.10.10.10
-
add: ibm-ibm-slapdIpAddress
ibm-slapdIpAddress: 127.0.0.1
```

2. Query the DN entry `cn=Configuration` in the `ibmslapd.conf` file to see the existing IP addresses to which `ibmslapd` listens to. Issue the `ldapsearch` command of the following format:

```
idsldapsearch -p port -D cn=adminDN -w adminPW -s sub \
-b "cn=Configuration" -L objectclass=*
```

An example excerpt of the output of the command is as follows:

```
n: cn=Configuration
cn: Configuration
ibm-slapdAdminDN: cn=root
ibm-slapdAdminGroupEnabled: true
ibm-slapdAdminPW: {AES256}ohtCABBYFbFo7jREOPz/zQ==
ibm-slapdCryptoSync: vDydxlFDW0xKtWBL
ibm-slapdDerefAliases: always
ibm-slapdIpAddress: 10.10.10.10
ibm-slapdIpAddress: 127.0.0.1
ibm-slapdPort: 389
#ibm-slapdPwEncryption must be one of:
# none/aes128/aes192/aes256/crypt/sha/ssh/md5/
# sha224/sha256/sha384/sha512/ssh224/ssh256/ssh384/ssh512
ibm-slapdPwEncryption: aes256
ibm-slapdServerBackend: RDBM
...
```

3. Restart the directory server instance.

Results

If there are no `ibm-slapdIpAddress` directives, the default behavior for `ibmslapd` is to listen on all available interfaces. After a specific (or multiple) `ibm-slapdIpAddress` entries are added to the `ibmslapd.conf` file, `ibmslapd` no longer listens to any interfaces that are not explicitly listed in the configuration file. To reset a directory server so that it listens on all available interfaces, you can remove all the `ibm-slapdIpAddress` entries from the `ibmslapd.conf` and restart the server.

Environment with SNMP agent configured

Sometimes, the environment in which an SNMP agent is configured might require tuning. IBM Tivoli Directory Integrator is set up to get the required result when you use an SNMP agent for monitoring directory server instances for performance and availability.

Some of the likely scenarios and their workaround are listed:

- When the **idssnmp** tool is running for a long time, it is observed that the `LDAP_HOME/idstools/snmp/logs/ibmdi.log` file size grows large.
If a user wants **idssnmp** to generate or keep less amount of log, the user can modify the `TDI_HOME/etc/log4j.properties` file and configure an appropriate log appender. For more information about the list of appenders, see *IBM Tivoli Directory Integrator version 7.1 Installation and Administrator Guide*.
- To run the **idssnmp** tool over SSL, user must edit the `LDAP_HOME/idstools/snmp/solution.properties` file and specify the certificate information.
- If a user wants to run **idssnmp** over SSL and the `solution.properties` file is not present, the user can create the solution files that are required by `idssnmp` by running the following command:

```
TDI_InstallDirectory/ibmdisrv -s LDAP_HOME/idstools/snmp -v
```


This command creates the `solution.properties` file in the `LDAP_HOME/idstools/snmp` directory.
- The **idssnmp** tool parses log files sequentially. For example, the **idssnmp** tool parses `slapd.log` for the newly generated logs and then proceeds to parsing the next log file, `audit.log`. As a result, the traps are sent in an order, which is not the same as the messages were generated in the log files. The trap messages contain information about the time with the log line when it was generated and the log file identifier. Based on the timestamp, the user is required to identify the occurrence order of traps.

Configuration update to directory server instance after reinstallation

Consider a scenario, where you might create a directory server instance, and then uninstall the IBM Security Directory Server product. Later you reinstall the product. In this scenario, you must take certain actions to restore services for the existing directory server instance.

On Windows systems

The directory server and administration server services are created for a directory server instance when a directory server instance is created. The services are removed when the instance is dropped or when the IBM Security Directory Server product is uninstalled.

If a user reinstalls the IBM Security Directory Server product, for the existing directory server instance the services do not get automatically added. To restore the services for the existing directory server instance, enter the following command at the command prompt:

```
ibmslapd -I instance_name -i  
ibmdiradm -I instance_name -i
```

Note: These commands create the services entry in a Windows operating system and do not create the directory server instance.

Start the services automatically for a directory server instance when the computer is started. Open the Services window from Administrative Tools and set the **Startup Type** to Automatic for the services that are associated with the directory server instance.

On nonWindows systems (such as AIX and Linux)

The entry that causes the ibmdiradm daemon to start automatically for a directory server instance gets removed. This entry is not restored for an existing directory server instance when you reinstall the IBM Security Directory Server product.

To restore an entry, compare the inittab file that you saved with the existing /etc/inittab file. Copy any lines in the saved file that are not in the existing file, and add them to the inittab file.

An entry in the /etc/inittab file is of the following format:

```
unique_id:run_states:action:path_to_script_to_be_run
```

An example entry:

```
ids0:2345:once:/opt/ibm/ldap/V6.3/sbin/ibmdiradm -I myinst1 > /dev/null 2>&1
```

Tombstone entries in a directory server

Before entries are permanently deleted from the RDBM database, a subtree is created to hold the entries to be deleted with operation attributes. This feature is available in IBM Security Directory Server, Version 6.2 and later versions.

The to-be-deleted entries are moved to the tombstone subtree, cn=Deleted Objects, and the attribute table is updated for the entry to mark the entry as deleted by adding attribute such as **isDeleted**. This feature is supported only on the primary RDBM back-end of the directory server. Tombstones are not supported in configuration, schema, or change log back-end.

There might be situations where data inconsistency gets introduced by entry deletions when this feature is enabled, which requires the intervention of the directory administrator. For performance reasons, there is no check that is provided, which can possibly prevent tombstones entries with the attribute **isDeleted** set to TRUE from being accidentally created or modified under any subtrees.

You can identify these entries in an RDBM back-end database by comparing the searches. Compare the search results that are returned by a normal search with a search base to that returned by a null base search.

For example, consider an RDBM back-end database with two subtrees: o=sample and cn=Deleted Objects. where, o=sample, contains three entries: cn=A, cn=B, and cn=C (with isDeleted=TRUE). The subtree cn=Deleted Objects containing entries, cn=X, cn=Y, and cn=Z (without isDeleted=TRUE).

When searches that use a search base and null base are requested without including the return deleted object control, the following results are displayed.

- In a search with a search base o=sample and search filter, objectclass=*, all entries under the search base, including entries with isDeleted=TRUE, are displayed.
- In a null base search with search filter, objectclass=*, all entries except for those entries with isDeleted=TRUE are displayed.

*Table 9. The results of different search base with search filter, objectclass=**

Subtree search base	Search filter	With control	Search results	Remarks
o=sample	objectclass=*	No	cn=A cn=B cn=C	cn=C is a normal entry with isDeleted=TRUE
null	objectclass=*	No	cn=A cn=B cn=Z	List LDAP_ENTRY table with isDeleted!=TRUE. cn=C is not qualified.

It is possible that the isDeleted attribute is accidentally deleted or is set to FALSE for entries under the tombstone subtree. When searches that use a search base and null base are requested with the return deleted object control, the following results are displayed.

- In a search with a search base, cn=Deleted Objects, and search filter, objectclass=*, all entries under search base are returned. However, when a search with a search base, cn=Deleted Objects, and search filter, isDeleted=TRUE, is requested, entries with isDeleted=FALSE are not returned.
- In a null base search with search filter, objectclass=*, all entries in the database are displayed. However, when a null base search with search filter, isDeleted=TRUE, is requested, only the entries with attribute isDeleted=TRUE in the database are displayed.

Table 10. The results of different search base with different search filters

Subtree search base	Search filter	With control	Search results	Remarks
cn=Deleted Objects	objectclass=*	Yes	cn=X cn=Y cn=Z	cn=Z is a tombstone without isDeleted=TRUE
cn=Deleted Objects	isDeleted=TRUE	Yes	cn=X cn=Y	

Table 10. The results of different search base with different search filters (continued)

Subtree search base	Search filter	With control	Search results	Remarks
null	objectclass=*	Yes	cn=A cn=B cn=C cn=X cn=Y cn=Z	List the LDAP_ENTRY table, including entries with isDeleted=TRUE.
null	isDeleted=TRUE	Yes	cn=C cn=X cn=Y	

Note: Deletion of schema attributes might fail because some of tombstone entries still reference them. A delete, rename, or restore of a tombstone entry is not replicated. It might result in data inconsistencies on replicas such as rename and restore cases.

Directory server instance backup

You can take multiple backups of the directory server instance, both offline and online, at multiple locations at different points in time. Understand how to work with the directory server instance backup in different scenarios.

For example, if myinst1 is the directory server instance and *instance-location/idsslapd-myinst1/backup1*, *instance-location/idsslapd-myinst1/backup2*, and *instance-location/idsslapd-myinst1/backup3* are the locations where backups are stored at T1, T2, and T3 time (where, T1<T2<T3).

When the instance, myinst1, is dropped or the database instance for the instance is unconfigured, the database configuration details (**dbbackuponline** and **clbackuponline**) are set to FALSE in the *instance-location/idsslapd-myinst1/backup3/dbback.dat* file.

Scenario 1

If the instance, myinst1, is re-created and configured with the backup location set to *instance-location/idsslapd-myinst1/backup2*, where the offline backup of the previous instance was stored before it was dropped. In this case, the results of the monitor search, "cn=backup,cn=monitor", will be consistent for searches that are done before the directory server instance is started (server state as stopped) and after the directory server instance is started (server state as running).

Scenario 2

If the instance, myinst1, is re-created and configured with the backup location set to *instance-location/idsslapd-myinst1/backup1*, where the online backup of the previous instance was stored before it was dropped. In this case, the results of the monitor search, "cn=backup,cn=monitor", are not consistent for searches that are done the directory server instance is started (server state as stopped) and after the directory server instance is started (server state as running).

The reason is when the server is in stopped state monitor search refers *instance-location/idsslapd-myinst1/backup1/dbback.dat* file for the

online status of database and change log (which is true since the previous online backup was stored at *instance-location/ldsslapped-myinst1/backup1*), and when the server is in running state monitor search refers directory server for the online status of database and change log (which is false as database is not configured for online backup). Suppose that a user starts the directory server. Then, the user does online backup that is based on the monitor search results that the user received when the directory server was in stopped state. In this scenario, the user would get unexpected behavior because the database is not configured for online backup.

If a user intends to restore from an existing backup for a re-created directory server instance, it is alright to configure to a previous backup location. However, suppose that the user intends to back up the re-created directory server instance. Then, to avoid the situation as mentioned in scenario 2, it is advisable to remove previous backup files from the location. Or else the user must specify a location that does not have any backup image.

Configuration of preaudit records for serviceability

Sometimes, when a directory server locks up or stops, the audit log might not have a record of the operation that causes the problem. You can configure the auditing of preaudit records to record operations that were not completed.

The audit log does not record the operation that causes the problem because the audit logs are updated after the directory server back-end completes the operation. So, any problems that occur before the audit records get updated are not logged and the result of the operation is unknown.

You can configure auditing of preaudit records to record operations that were not completed. When preaudit records are enabled, the audit plug-in is called to update audit records before the operation completes. When preaudit is enabled, the thread ID is also logged in the audit header. To enable pre-auditing, you must set the value of the *IBMSLDAPD_PREOP_AUDIT* environment variable to YES. You can set this value by accessing the environment variable or by using the **ldapmodify** command with the following format:

```
ldapmodify -D adminDN -w adminPW
dn: cn=Front End, cn=configuration
changetype: modify
add: ibm-slapdSetEnv
ibm-slapdSetEnv: IBMSLDAPD_PREOP_AUDIT=YES
```

An example of a pair of diagnostic audit records when preaudit is enabled, where the sequence identifier is 3: *PREOP: 3* and *POSTOP: 3*, is as follows:

```
AuditV3--2007-08-29-11:44:32.912-06:00DST--V3 PREOP: 3 threadId: 1161116592
Search--bindDN: cn=root--client: 127.0.0.1:1044--connectionID:
3--received: 2007-08-29-11:44:32.912-06:00DST--Success
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: o=sample
scope: baseObject
derefAliases: neverDerefAliases
typesOnly: false
filter: (objectclass=*)
```

```
AuditV3--2007-08-29-11:44:33.092-06:00DST--V3 POSTOP: 3 threadId: 1161116592
Search--bindDN: cn=root--client: 127.0.0.1:1044--connectionID:
3--received: 2007-08-29-11:44:32.912-06:00DST--Success
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
```



```
criticality: false
base: o=sample
scope: baseObject
derefAliases: neverDerefAliases
typesOnly: false
filter: (objectclass=*)
```

Entries that are displayed to root and anonymous users

No entries are displayed to root and anonymous users when logged on to a proxy server instance with the Web Administration Tool. You must change the page control restrictions for such users.

Root and anonymous users are not able to view entries with the Web Administration Tool when logged on to a proxy server, even with the **ibm-slapdAllowAnon** attribute under the DN entry `cn=Connection Management, cn=Front End, cn=Configuration` set to true. The reason is because the Web Administration Tool uses page control to browse through entries. If paging is enabled only for administrators by setting the **ibm-slapdPagedResAllowNonAdmin** attribute to false under the DN entry `cn=ProxyDB, cn=Proxy Backends, cn=IBM Directory, cn=Schemas, cn=Configuration`, then non-administrators are not allowed to browse through entries with the Web Administration Tool. This restriction is also applicable in the case of RDBM servers.

For the root and anonymous users to view the entries with the Web Administration Tool, the following must be considered:

- Set the **ibm-slapdAllowAnon** attribute under the DN entry `cn=Connection Management, cn=Front End, cn=Configuration` to true.
- Set the **ibm-slapdPagedResAllowNonAdmin** attribute under the DN entry `cn=ProxyDB, cn=Proxy Backends, cn=IBM Directory, cn=Schemas, cn=Configuration` to true.

After you set the attributes, restart the directory server instance.

Note: The root user is an anonymous user in the case of proxy server for DIT-related operations.

Directory server instance is restored to latest consistent state

When you run the restore operation on a directory server instance, the directory server instance might get restored to latest consistent state. It is restored to this state instead of getting restored to the point when the backup was done. Follow the steps to resolve this issue.

Observed

1. Create and configure a directory server instance for online backup.
2. Stop the directory server instance and run the initial offline backup either with the Web Administration Tool or the **idsdbback -u -k** command.
3. Add the suffix, `o=sample`, and start the directory server instance.
4. Add the entry `o=sample`.
5. Verify that the database parameter `LOGARCHMETH1` is correctly set.
6. Perform a restore operation either with the Web Administration Tool or the **idsdbrestore -k** command.

7. Verify that the suffix, `o=sample`, is not present (since backup was done before the suffix was added).
8. Add the suffix, `o=sample`, and start the directory server.
9. Run **ldapsearch** for the entry `o=sample` and you can observe that the entry is present.

Expected result

The entry, `o=sample`, must not be present because only the suffix, `o=sample`, is added after restore on a clean database (no data).

Reason

During roll-forward, DB2 scans the current logs in the `newlogpathlocation`. Because of the options that are specified in the roll-forward, DB2 scans the logs until the end, and restores a database to the latest consistent state.

For example, at the time of backup, suppose that you have 100 entries in the directory server. After the backup operation, if you delete five entries and then run the restore operation. You might still find the 95 entries in the directory rather than the 100 entries that you backed up. This reason is because, the latest consistent state of the database was after the deletion on five entries.

However, you can modify the options in the rollforward recovery operation such that the database is restored to the point where it had 100 entries. You must specify the timestamp of the last committed change. This timestamp is the one at which the 100th entry was added and to obtain this value of timestamp is difficult.

Online backup and restore limitation

When you change the original backup location, you might encounter an error. Follow the steps to work around this limitation with the online backup and restore feature.

During online backup and restore, suppose that the folder name (backup location) to which online backup was initially configured is changed for the subsequent backups. It is observed that no error is thrown during the backup operation (**idsdbback**), but during the restore operation (**idsdbrestore**) the following error messages might be displayed:

```
...
GLPCTL101I Restoring backup database rdsdb to configured database rdsdb.
GLPCTL103E Failed to restore backup database rdsdb to configured database rdsdb.
GLPDBR004E Failed to restore directory server instance 'tdsadmin'.
GLPDBR028W The program did not complete successfully. View earlier error messages
for information about the exact error.
...
```

Reason

As per the **idsdbback** and **idsdbrestore** (also available as **dbback** and **dbrestore**) design for online backup, the first-time backup must be a complete offline backup while the `ibmslapd` process is in stopped state. After the first offline backup, the online backup feature can be used while the `ibmslapd` process is running.

During the first offline backup, the **idsdbback** command takes the following options:

```
idsdbback -I instance_name -k /path/backupfolder1 -u [-a /path/logarchivefolder]
```

If the optional path for logarchive folder is not provided, the command uses a folder inside the backupfolder1 folder (as per the example) to configure the logarchive folder and sets this value in the corresponding DB2 database configuration parameter, **LOGARCHMETH1**.

If the backup folder is changed for a subsequent online backup, **idsdbrestore** fails if the previous backup folder does not exist, since the **LOGARCHMETH1** still points to the previously configured value.

Confirming the problem

To confirm, verify the **LOGARCHMETH1** variable for the corresponding database configuration.

```
su - instance_name
db2 list db directory
db2 get db configuration for databasename | grep -i LOGARCHMETH1
```

Note: Replace the *instance_name* and *databasename* with the appropriate names.

Resolving the problem

If you want to change the backup location after the first offline backup, or even after subsequent online backups, complete the following procedure to update the backup folder and logarchive folder values:

1. Stop the **ibmslapd** process.

```
ibmslapd -I instance_name -k
```

2. Use the **idsdbback** command to update both the backup folder and logarchive folder.

```
idsdbback -I instance_name -k /path/backupfolder2 \
-a /path/backupfolder2/INACTIVE_LOGS -u -n
```

3. Start the **ibmslapd** process.

```
ibmslapd -I instance_name -n -t
```

Log management servers fails to stop

When you attempt to stop the log management service after you start the service with the Web Administration Tool, it fails to stop. Follow the steps to work around this problem.

Observed

The log management service gets started when you use the Start/Stop log management panel of the Web Administration Tool. However, when you attempt to stop the log management service by using the Start/Stop log management panel, the panel displays that the service is stopped. However, the log management service is running in the background.

Expected behavior

The expected result is that the log management service must stop.

Reason

For the log management to work, IBM Tivoli Directory Integrator is required. The IBM Tivoli Directory Integrator, version 7.1 installation wizard prompts for a installation location for the Directory Integrator Solutions Directory. If you opt for the default preselected option, "Use a subdirectory named Directory Integrator under my home directory," this problem with log management service occurs.

Workaround

If you opt for the **Use Install Directory** option from the IBM Tivoli Directory Integrator, version 7.1 installation wizard for the Solutions Directory, the mentioned problem with log management service does not occur. For example, if the IBM Tivoli Directory Integrator is installed in the location `/opt/IBM/ldap/V7.1/TDI`, and you opt to provide the Directory Integrator Solutions Directory within the `/opt/IBM/ldap/V7.1/TDI` directory, then the problem with log management service is not observed.

Instance does not start and returns error GLPCRY007E

In certain scenarios, the directory server instance might not start and might return error GLPCRY007E. Follow the steps to resolve this issue.

Scenario

1. Create a directory server instance, `inst1`, configure the instance, and start the instance. The encryption seed that is used to create the instance, `inst1`, is `thisismyseed`.
2. Drop the instance, `inst1`, without dropping the database that is associated with it.
3. Re-create the instance with the encryption seed, `thisismyseed`, and configure the instance with the existing database.
4. Start the instance.

Observed

The instance does not start and returns error:

GLPCRY007E The directory key stash file is inconsistent with the associated encrypted data.

Reason

When a directory server instance is created and is started, some information from key stash file (`.ksf`), is stored in the database. Therefore, an existing database cannot be used with a key stash file that gets created when an instance is re-created.

Workaround

In such case, if you intend to use an existing database with a new instance, then at the time of instance creation you must use **-e** and **-g** options to specify the encryption seed and encryption salt values for the new instance. This encryption seed and salt value must be same as the encryption seed and salt value of the dropped instance.

If you did not provide the salt value with the **-g** option for the instance that you are intending to drop, then the salt value must be determined before an instance is dropped. Issue the **idsldapsearch** command of the following format to retrieve the salt value.

```
idsldapsearch -h IP address -p port -s base -b "cn=crypto,cn=localhost" \
objectclass=* ibm-slapdCryptoSalt
```

Chapter 12. Interoperability

Use the information on interoperability between IBM Security Directory Server and other directory servers to troubleshoot related issues.

Interoperability with Novell eDirectory Server

When you do a simple bind by using IBM Security Directory Server client utilities against Novell eDirectory Server, you might encounter an error message. Run the configuration command to resolve this problem.

You might encounter an error message such as `ldap_bind: Confidentiality required`.

Run the following command:

```
#ldapconfig set "Require TLS for Simple Binds with Password=no"
```

Interoperability with Microsoft Active Directory

If IBM Security Directory Server is configured over SSL by using `serverClientAuth` authentication, follow the steps to make it work with Microsoft Active Directory client **LDP.exe**.

To make IBM Security Directory Server configured over SSL by using `serverClientAuth` authentication to work with Microsoft Active Directory client **LDP.exe**, complete the following steps.

1. Select **Internet Information Services (IIS) Manager** from Administrative Tools in Control Panel.
2. On the left navigation panel, select the **Web Site** node.
3. Under the website node, right-click **Default Web Site**, and then select **Properties**.
4. On the Default website Properties dialog box, select the **Directory Security** tab.
5. To request for a new certificate, click **Server Certificate** under the Secure communications area. The Web Server Certificate Wizard is opened.
 - a. On the Server Certificate page in the IIS Certificate Wizard dialog box, select the **Create a new certificate** option and click **Next**.
 - b. On the Delayed or Immediate Request page, enter the required options and click **Next**.
 - c. On the Name and Security Settings page, in the **Name** field enter the host name of the system and click **Next**.
 - d. On the Organization Information page, specify appropriate names and click **Next**.
 - e. On the Your Site's Common Name page, in the **Common name** field, enter the host name of the system and click **Next**.
 - f. On the Geographical Information page, specify appropriate values and click **Next**.
 - g. On the Certificate Request File Name page, in the **File name** field specify the path name and file name for the certificate request and click **Next**.
 - h. The summary of the values is displayed. Click **Next**.

- i. Click **Finish**.
6. Send the certificate request by using these steps to any certificate authority (CA) to issue a certificate.
7. After you receive the server certificate, add the certificate by using IIS Certificate Wizard.
 - a. On the Pending Certificate Request page, select the **Process the pending request and install the certificate** option and click **Next**.
 - b. On the Process a Pending Request page, in the **Path and file name** field specify the path name and file name of the certificate. You can also use Browse to select the certificate. Click **Next**.
8. Export the personal certificate to pfx or p12 format by using IIS Certificate Wizard.
 - a. On the Modify the Current Certificate Assignment page, select the **Export the current certificate** to a .pfx file option and click **Next**.
 - b. On the Export Certificate page, in the **Path and file name** field enter the path name and file name where pfx certificate to be stored. Click **Next**.
 - c. On the Certificate Password page, in the **Password** and **Confirm password** fields enter the password and click **Next**.
 - d. On the Export Certificate Summary page, the summary of the provided values is displayed. Click **Next**.
 - e. Click **Finish**.
9. To import the certificate, double-click the stored pfx certificate. The Certificate Import Wizard is opened.
 - a. On the File to Import page, in the **File name** field enter the path and file name of the pfx certificate and click **Next**.
 - b. On the Password page, enter the password and click **Next**.
 - c. On the Certificate Store page, select the **Place all certificate in the following store** option and click **Browse** and select **Personal** from the **Select the certificate store you want to use** list in the Select Certificate Store dialog box. Click **Next**.
 - d. Click **Finish**.
10. To export the personal certificate in BER format, complete the following steps.
 - a. Open Internet Explorer, select **Tools > Internet Options**.
 - b. Select the **Content** tab in the Internet Options dialog box, and select **Certificates** under the Certificates area.
 - c. On the **Personal** tab in the Certificates dialog box, select the certificate and click **Export**. The Certificate Export Wizard.
 - d. On the Export File Format page, select the **Base-64 encoded X.509 (.CER)** option and click **Next**.
 - e. On the File to Export page, in the **File name** field enter the file name that you want to export and click **Next**.
 - f. Click **Finish**.
11. On a system on which an directory server instance is running, open the directory server key database file by using the GSKit key management application, **ikeyman**.
12. Add the exported certificate as a signer in the server key database.

Chapter 13. Known limitations and general troubleshooting

Use the list of known limitations and guidelines for general troubleshooting to identify and resolve problems that are related IBM Security Directory Server.

Known limitations

Use the descriptions of the known limitations in IBM Security Directory Server, version 6.1 and later of directory servers to identify and work around the problems.

IBM Installation Manager preinstallation summary page shows existing software for installation

You can verify the features selected for installation on the IBM Installation Manager preinstallation summary page.

You can use IBM Installation Manager for installation of IBM Security Directory Server and its corequisite software, such as IBM DB2 or GSKit. If your computer contains a corequisite software that is supported by IBM Security Directory Server, you can choose to select that software instead of installing the software again.

If you select the existing corequisite software in IBM Installation Manager, it shows the software in preinstallation summary page even though the program does not install them.

If you choose to continue with the existing DB2 or GSKit during the installation, IBM Installation Manager updates its registry with the feature entries that you selected.

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
```

Invalid data entered on command-line utilities

A known limitation is that some types of invalid data that are entered on command-line utilities do not produce an error.

If you enter a command that contains invalid data after all required options are specified, you will not receive an error message. For example, the following command contains the required options for the **idsdnpw** command, but the `--` characters that follow 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

No locking mechanism exists to prevent conflicting commands from running at the same time for the same directory instance.

For example, you can run a command to configure a database and drop the database at the same time.

Unable to drop database

On Windows systems, you might not be able to drop the database immediately after you stop a directory server instance.

This problem occurs in a scenario where all of the following conditions are true:

- 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 you stop 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 you run the **idsidrop** or **idsucfgdb** commands. Alternatively, wait at least 2 minutes after you stop the server, and then drop the database.

Partial replication

Partial replication is a replication feature that replicates only the specified entries and a subset of attributes for the specified entries within a subtree. The entries and attributes that are to be replicated are specified by the LDAP administrator. Using partial replication, an administrator can enhance the replication bandwidth according to deployment requirements.

For instance, an administrator might choose the entries of the object class person with **cn**, **sn**, and **userPassword** attributes to be replicated and description attribute not to be replicated.

There are situations when administrator's intervention is required for the smooth running of partial replication. These scenarios are listed.

Creating missing parent entries on the consumer

In filtered replication, an entry addition might fail displaying "No such object" error because the parent entry does not exist on the consumer. It happens because the parent entry did not match the filter and was not replicated. In such cases, if the **ibm-replicationCreateMissingEntries** attribute is set to TRUE, the supplier must detect this error case and then generate and submit an add request for the missing entry before the supplier tries the add operation again instead of processing this case as an error. The missing entry must have the same DN as the immediate parent of the entry whose add failed. The missing entry belongs to the objectclass **extensibleObject** and contains operational attributes for create and modify timestamps as present on master server, that is, the timestamps are not modified when the entry is created on consumer. The missing entry must have ACLs as on the supplier server and must also have the description attribute value that is set to Missing entry created by *master server*.

Scenario

Sometimes the method to generate and submit a request to add a missing entry is recursive. The end condition is either a successful addition of all missing ancestors in the chain or a failure. The failure might occur during addition of any of the missing ancestors (for any reason other than `NO_SUCH_OBJECT`). If there is a failure, the change cannot be replicated and administrator intervention is required.

Workaround

The administrator must manually take care of handling errors when the `ibm-replicationCreateMissingEntries` attribute is set to `FALSE`. Administrators can also use error logs to identify the replication failure error messages that are logged in to error logs.

Modification in replication filter

Scenario

In partial replication, changes to replication filter can be dynamic. When a replication filter is changed, the data on the consumer would be in sync with the supplier cannot be assured.

Workaround

In cases where replication filter is changed, the administrator must take of such changes and reinitialize the consumer as per the new replication filter.

Note: The replication filter entry cannot be deleted if it is in use.

Replication is not initiated

In a replication environment, if a supplier uses a password encryption setting that is not supported by the consumer, then replication is not initiated.

Also, the supplier logs a message and sets the replication state to “error xxxx” where xxxx is the ID of the message that describes the problem.

Migration of an instance from version 6.1 or later to version 6.3.1

To migrate an existing instance from IBM Security Directory Server, Version 6.1 or later to version 6.3.1, you can use the `idsimigr` command-line utility.

The `idsimigr` command retrieves the schema and configuration files of the instances from the standard location specific to the instances. When you migrate from IBM Security Directory Server, Version 6.1 or later to version 6.3.1, certain checks must be done. Otherwise, the tool might exit and display error messages.

- If an instance exists, the backup directory must not be specified. If the backup directory is specified, the tool exits and displays error messages.
- If the IBM Security Directory Server instance is earlier than version 6.3.1 that is to be migrated is dropped before you run `idsimigr`, the backup directory must be specified. In such a scenario, the encryption key is not required but if the encryption key is specified, the tool exits, and displays error messages.
- During migration, the Windows service entry for each directory server and the Directory Administration server are migrated to IBM Security Directory Server, version 6.3.1. To avoid unforeseen errors, you must back up the schema, configuration, and key stash files before migration, even if you did not drop the instances.

Alias dereferencing does not work

In IBM Security Directory Server, Version 6.1 and later, alias dereferencing might not work when persistent search is run on a server with no alias entries.

If persistent searches are run before any alias entries are added to the server, then persistent searches do not dereference aliases. That means, only if alias entries exist on the server before you run persistent searches, the dereferenced aliases are displayed.

Operation times out

Suppose that both proxy and back-end servers are configured to use PKCS#11 mode. They are required to communicate with a remote nCipher cryptographic hardware for SSL operation. In this scenario, the operation times out. To increase the operation timeout duration, you must increase the number of times that a proxy server must try to attempt to establish a connection.

You must increase the number of times that a proxy server tries to establish a connection because:

the total time for which a proxy server waits to establish a connection =
maximum time for which proxy waits to establish connection *
number of retries by a proxy server to establish a connection

To increase the number of times that the proxy server must try again, export the environment variable, *SERVER_ATTEMPT_TIME*, with the required count. Set the count for trying again to greater than 12, if the cryptographic hardware used for SSL operation is at a remote location.

Instance stops when nCipher cryptographic hardware client is restarted

IBM Security Directory Server, Version 6.2 or later instance stops when nCipher cryptographic hardware client is restarted.

Scenario

The following steps describe the situation in which a directory server instance might stop.

1. Start a directory server instance. The instance is configured over SSL with server client authorization to use PKCS#11 in key storage and accelerator mode.
2. Perform search operation by using an LDAP client in SSL mode.
3. Restart the cryptographic hardware used.
4. Perform search operation by using an LDAP client in SSL mode.

Reason

You must not restart the cryptographic hardware if IBM Security Directory Server, Version 6.2 or later instance uses PKCS#11 in key storage or accelerator mode. If the cryptographic hardware that is used by a server instance for cryptographic operations is reset, then the instance stops logging appropriate messages in trace file.

Error with `idsldapdiff` tool query

When you query an entry of large size by using the `idsldapdiff` tool, an error might occur.

The Java implementation of the **idsldapdiff** tool has a limitation. Because of this limitation, it is unable to handle entries on IBM Security Directory Server that are more than 50 MB in size. As a result, the tool might throw an Out of Memory exception when it deals with entries with more than 50 MB in size.

Synchronization of entries with **idsadsrun** utility fails

The **idsadsrun** utility might fail when it is synchronizing many entries. An exception such as size-limit or time-limit occurs.

To avoid exceptions like size-limit or time-limit, you must consider the following settings:

1. Before you start synchronization, configure Microsoft Active Directory parameters to set the following values:

MaxPoolThreads	4
MaxDatagramRecv	4096
MaxReceiveBuffer	10485760
InitRecvTimeout	120
MaxConnections	5000
MaxConnIdleTime	900
MaxPageSize	1000000
MaxQueryDuration	1000
MaxTempTableSize	10000
MaxResultSetSize	262144
MaxNotificationPerConn	5
MaxValRange	1500
2. In the **ibmdisrv** file, tune the JVM parameters, for example, for a system with 1 GB of RAM, the parameter values can be **Xms254m-Xmx1024m**. You can tune the parameters according to your system configurations. For best results, use a system with high-end configurations to run the Active Directory synchronization tool.
3. Also, synchronizing approximately 100000 entries by using the “Run full synchronization of the entries from Active Directory Server to IBM Security Directory Server followed by real-time synchronization” mode while you run the **idsassrun** tool gives best results. With the “Run real-time synchronization” mode, up to 400000 entries can be synchronized.

idsadsrun utility fails when instance is run on a different port

The **idsadsrun** utility fails if a directory server instance is run on a different port with the **-p** option.

Presently, the Active Directory synchronization tool detects the administrator DN, password, LDAP URL, and port number from the instance name. Therefore, when a directory server instance is run on a different port with the **-p** option, the tool is unable to detect the port number that is specified by using the **-p** option.

Operations error during null base search

Operations error is displayed when null base search is run against a proxy server.

Proxy server does not support null base search and gives an operations error if null base search is fired against it.

User account gets locked

When the **pwdLockout** attribute is set to true, the user account might get locked even if the number of invalid bind attempts is less than the **pwdMaxFailure** value.

A user account might get locked when all the invalid bind attempts are made within a specified time interval that is set in the **pwdFailureCountInterval** attribute. For example, consider the following attributes are set to:

```
ibm-pwdPolicyStartTime=20070217044605Z
pwdInHistory=0
pwdCheckSyntax=1
pwdGraceLoginLimit=0
pwdLockoutDuration=0
pwdMaxFailure=3
pwdFailureCountInterval=0
passwordMaxRepeatedChars=0
pwdMaxAge=99
pwdMinAge=0
pwdExpireWarning=0
pwdMinLength=5
passwordMinAlphaChars=0
passwordMinOtherChars=0
passwordMinDiffChars=0
ibm-pwdPolicy=true
pwdLockout=false
pwdAllowUserChange=true
pwdMustChange=false
pwdSafeModify=false
ibm-pwdGroupAndIndividualEnabled=true
```

With this setting, if a user makes three invalid bind attempts, the user can still continue with bind attempts because the **pwdLockout** attribute is set to false. However, **pwdFailureTime** is registered even when **pwdLockout** is false, therefore if user does three invalid bind attempts with **pwdLockout=false**, **pwdFailureTime** has timestamps of the consecutive authentication failures.

Set the **pwdLockout** attribute to true:

```
# idslapmodify -D cn=RDN_value -w password
-p port_number -h host_name
dn:cn=pwdpolicy,cn=ibmpolicies
pwdLockout:true
```

Now, when the **pwdLockout** attribute is set to true another invalid or valid bind attempt causes lockout of user account. The lockout is because the invalid bind attempts made when **pwdLockout=false** is also taken into account according to the number of values in the **pwdFailureTime** attribute that are younger than **pwdFailureCountInterval**.

Description attribute for groups does not sync from Active Directory

The description attribute for groups does not sync from Active Directory to IBM Security Directory Server.

Active Directory synchronization solution synchronizes only the user entry attributes provided with **TDSOptionalAttributes** in the **adsync_public.prop** file.

Possible memory leak with PKCS#11 support configured

When you configure a directory server over SSL to use PKCS#11 SYMMETRIC acceleration support, there are chances for memory leak.

Consider a scenario where a directory server over SSL is configured to use PKCS#11 SYMMETRIC acceleration support. This configuration is for

cryptographic operations by using nFast cryptographic library. In this scenario, memory leak is observed during operations.

Note: nFast cryptographic library is a third-party library. It is used for PKCS#11 support that is provided by IBM Security Directory Server.

LDIF files with SHA-2 encrypted password or attributes

When you import LDIF files that contain SHA-2 encrypted password or encrypted attributes to versions earlier than 6.3, the data is encrypted based on the value of **ibm-slapdPwdEncryption** attribute.

The **db2ldif** utility exports user password data and other encrypted data as it is stored in the directory. If the data is encrypted, then it is exported in the same format to the LDIF file.

The LDIF import utilities, **ldif2db** and **bulkload**, check the format of the data in the LDIF file to verify whether it is in a recognizable encryption format. If any data that is not recognized from the supported encryption method is assumed to be in clear text format, and is encrypted or not encrypted based on the configured value of the **ibm-slapdPwdEncryption** configuration attribute. Therefore, when the earlier versions of IBM Security Directory Server import IBM Security Directory Server, Version 6.3 or later LDIF file that contains data encrypted using the SHA-2 family of encryption scheme, the earlier version of servers assume that the data is in clear text since SHA-2 family of encryption scheme is unknown encryption format, and encrypts the data according to the configured value of the **ibm-slapdPwdEncryption** configuration attribute.

Multivalued attributes in a virtual list view search

Duplicate entries might be returned in a virtual list view search if the sort key is a multivalued attribute.

Explanation

In virtual list view searches, the search filter resolutions are done in the database. The entire result set is not read from the database at one time. However, in a normal search operation a list of EIDs is maintained in the memory. It ensures that duplicate entries are not returned to clients, even if DB2 returns duplicate EIDs.

Because the entire result set (list of EIDs) is not read into the memory, the constraint of identifying and preventing duplicates exists. Suppose that a virtual list view search is done with a primary sort key attribute that has multiple values. Then, the entries that are returned might not be in sorted order. Additionally, duplicate entries might also be returned.

Example

Consider a directory server with the following data set:

Table 11. Entries and multivalued attribute of the entries

EID	Values of the cn attribute
1	A, Y
2	C, J
3	E

In a normal search with cn as the sort key, the entries are returned in the following order: 1, 2, 3. However, the search filter resolution for the DB2

query returns EIDs in the following order: 1, 2, 3, 2, 1, based on the values of cn. In this case, the duplication is prevented by maintaining the list of EIDs at the server end.

In a virtual list view search, the entire result set is not maintained in memory and therefore preventing duplication is not possible.

Consider a virtual list view search that is sent with the following values: before count = 1, after count = 1, offset = 3, and content count = 0. If the virtual list view control is applied over the DB2 result set, the entries 2, 3, 2 are returned. Here, the entry with EID=2 is returned twice. The result shows that there is a possibility of returning duplicate entries in a virtual list view search if the sort key is a multivalued attribute.

Distributed directory environment search scope

In a distributed directory environment, only base scope search with `ibm-allMembers` is supported.

If distributed group and dynamic distributed group are enabled in the configuration file, then only base scope search with `ibm-allMembers` is supported. If `onelevel` or `subtree` scope search is attempted with `ibm-allMembers`, then an appropriate error message is logged in the `ibmslapd.log` file and `LDAP_UNWILLING_TO_PERFORM` is returned.

However, if distributed group and dynamic distributed group are disabled in the configuration file, then a search for `ibm-allMembers` is forwarded to a single back-end server. In this case, the search returns group members for all search scopes.

Instance fails to start if system date is modified

A directory server instance might fail to start if the system date is modified.

Suppose that the system date is set to a previous date. For example, it might be set to one month before the date when the directory server instance was configured on the system. If such a significant change is made to the system date, then the directory server instance might fail to start. The following error messages can be seen:

```
GLPRDB001E Error code -1 from function:" SQLTables ".
GLPRDB001E Error code -1 from function:" SQLFetch ".
GLPRDB001E Error code -1 from function:" SQLFetch ".
GLPRDB001E Error code -1 from function:" SQLFetch ".
GLPSRV064E Failed to initialize be_config.
```

In this scenario, the directory server instance is run with server trace set to ON and the debug level is set. The following error messages can be seen in the server trace:

```
188:22:35:24 T1 retrieving SQLGetDiagRec info
188:22:35:24 T1 Error - map_rc fnc: henv=0,hdbc=0,hstmt=10001,native
retcode = -443; state = "38553"; message = "[IBM][CLI Driver][DB2/SUN64]
SQL0443N Routine "SYSIBM.SQLTABLES" (specific name "TABLES") has returned
an error SQLSTATE with diagnostic text "SYSIBM:CLI:-727". SQLSTATE=38553"
```

These error messages can also be seen in the `db2diag.log` file.

Format of the DN gets changed

In the configuration file, the format of the DN gets changed when a composite DN is added as suffix.

Suppose that a composite DN is added as suffix. Then, the format of the DN that gets added to the configuration file is different from the DN value that was provided. For example, a composite DN, o=sample+c=in gets updated in configuration files as c=i\20 + o=sample.

idsdbmaint tool error message

The **idsdbmaint** tool might give an error message, which states that it is unable to estimate the database size. This error is related to the privileges of the instance owner.

When the **idsdbmaint** tool is run with root or administrator privileges the tool inherits those privileges and therefore, the tool is able to access a directory even if it does not have write permissions or sufficient privileges for the directory instance owner. The **idsdbmaint** tool attempts to estimate the directory size with the privileges of the instance owner. In this case, if the instance owner does not have sufficient privileges to run the operation, the tool gives the following error.

GLPDBA054E Unable to estimate the database size.

Error opening filename.cat

An error message which states that there is an error opening filename.cat gets displayed when you run a directory server. This error is related to the language pack or locale.

If a directory server is set to a locale that does not have corresponding message files for that locale, then an error message Error opening filename.cat is displayed along with an appropriate message in English locale.

The reason for this error can be the following conditions:

- An incorrect language pack is installed on the system.
- IBM Security Directory Server does not support that particular locale.

The values TRUE and FALSE are not translated

The directory server messages do not translate the values TRUE and FALSE to the corresponding locales of the translated version. You can see the issue in the translated versions of IBM Security Directory Server, the graphical user interface (GUI) tools, such as the Web Administration Tool.

Some schema-related keywords are not translated

The values of some schema-related keywords such as syntax and matching rules are not translated. You might see the issue in the Web Administration Tool for the translated versions of IBM Security Directory Server.

Date is not displayed properly for the Russian locale

You might see the following issue in the translated version of the Web Administration Tool in the Russian locale: Sometimes, the date format either gets displayed in wrong format or the last character of the month name gets truncated. This issue is a limitation with the tool.

Date and time are displayed in English in translated versions

You can see the following issue on certain panels, such as Manage backup and restore, in the translated versions of the Web Administration Tool: The date and time values that are displayed on the panels are in the English locale instead of the locale of the translated version.

Error logo is not displayed with error messages

If you access panels on the Web Administration Tool when the directory server is in the stopped state, an error panel is displayed with error messages. However, on this error panel, the error logo is not displayed. This issue is a limitation with the Web Administration Tool.

Mnemonics missing from tool panels

Mnemonics are missing from the panels of the Instance Administration Tool and Configuration Tool. This limitation in the tools is specific to the French and Korean translated versions.

On the panels of Instance Administration Tool and Configuration Tool, the mnemonics for the buttons, such as Help, Finish, and Cancel, might not be available. This limitation in the tools is specific to the French and Korean translated versions.

In the French and Korean translated versions of the tools, you can use the following keys for the buttons:

- Help - The mnemonic key is "H".
- Finish - The mnemonic key is "F".
- Cancel - The mnemonic key is "C".

Alternatively, you can use hot keys (a combination of keys) to access the following buttons:

- Alt + H - Help
- Alt + F - Finish
- Alt + C - Cancel

Attribute encryption in RDN of an entry

The encrypted attribute of the RDN is displayed in clear text instead of being displayed in the encrypted format. This issue is a known limitation.

When you add an entry to a directory server instance with an RDN that has encrypted attribute in it, the following error message is displayed:

```
GLPWDM003E An error occurred while adding entry uid=5user,uid=5user,o=ibm,c=us :  
uid=5user,uid=5user,o=ibm,c=us: [LDAP: error code 34 - GLPSRV156I  
Encrypted attributes are not allowed in entry distinguished names.]
```

An attribute that is already present in the RDN of an entry can be encrypted without getting any error message. When the directory server instance is started, the encrypted attribute of the RDN must be displayed in the encrypted format. Instead, the entry displays the RDN in clear text.

This inconsistency is a limitation in the existing design.

LDAP search filters that exceed 4K are not supported

If an LDAP search filter exceeds the 4K limit, then the server might throw an `ldap_search:bad search filter` error. To avoid this error, you must use search filters that do not exceed the 4K limit.

An error message might also be logged in the `db2cli.log` file, which indicates a syntax error in the query sent to DB2. For example, the error message that is logged in the `db2cli.log` can be of the following format:

```
12/04/07 10:38:24 native retcode = -104; state = "42601"; message =  
"[IBM][CLI Driver][DB2/6000] SQL0104N An unexpected token  
"END-OF-STATEMENT" was found following ".ORGANIZATIONDN WHERE". Expected  
tokens may include:"). SQLSTATE=42601
```

Also, the `ibmslapd.log` file might contain the following error:

```
12/04/07 10:37:44 AM GLPRDB001E Error code -1 from function:" SQLExecute " .
```

To avoid these errors, you must use search filters that do not exceed the 4K limit.

Error during creation of a directory server instance from an existing instance

If the version of DB2 on the source and target server are different, the **idsideploy** tool displays an error when you create a directory server instance from an existing directory server instance.

During the creation of a directory server instance from an existing directory server instance, the **idsideploy** tool takes online backup of the source database, including the logs. At the target server, the database is restored with rolling forward of logs to bring the database to a consistent state. However, there is a limitation when the target database is DB2 version 9.5 and the source database is a previous version, DB2 version 9.1. The rolling forward of logs from a previous level to DB2 v9.5 is not supported.

Therefore, when you use the **idsideploy** tool, you must use the same DB2 versions on the source and target server.

To know more about supported platforms for DB2 backup and restore operations, see Backup and restore operations between different operating systems and hardware platforms.

The idsideploy tool fails to restore a database

The **idsideploy** tool might fail to restore a database if the backup location has backup images of the database.

When the **idsideploy** tool is run to create a copy of a directory server instance with data of an existing directory server instance, you must ensure that the directory path specified with the **-L** option does not already have backup image of the database, which the tool is attempting to restore. If a backup image is already present, then the restore operation of **idsideploy** fails.

Creation of online backup image fails

The **idsdbback** command might fail to create an online backup image of a directory server instance that is created by the **idsideploy** tool.

When the **idsideploy** tool is used to create a copy of a directory server instance (along with database), the tool backs up the source database and restores it on the target server. During this process, all the internal database settings are also copied on to the target server as it is. The error messages that might get displayed are:

```
GLPDBB051E Failed to create path '/export/home/mybkup/back/INACTIVE_LOGS'
for logging inactive log files.
GLPDBB010E Failed to back up directory server instance 'inst2'.
```

One of the reasons for this error is that the target database uses the same settings as the source database. You must set the archive path for the target server instance before you do the online backup operation for the target server instance. Otherwise, the online backup might fail. To know more about the **idsdbback** and **idsideploy** commands, see *IBM Security Directory Server Command Reference*.

Unable to connect from OpenLDAP client over DIGEST-MD5

A directory server instance of version 6.2 fails to authenticate an OpenLDAP client over the DIGEST-MD5 SASL mechanism, if the version of OpenLDAP client is 2.4.11. However, with the directory server instance of version 6.2, you can use OpenLDAP clients version 2.3.33.

Inconsistent data when transaction updates are replicated

There is a possibility of inconsistent data on a directory server when transaction updates are replicated in an environment with failover setup.

When transactional updates are replicated by a supplier, the updates are not replicated in a transactional manner by the supplier to its consumers. In a replicated environment, the supplier replicates the transactional updates to its consumers only when the transaction is complete (committed or rolled back state). If a supplier goes offline during replication of the transactional updates, it is possible that only a part of the update is replicated to its consumers. In this case, when the supplier is brought online the remaining updates that are in its replication queue is replicated automatically.

However, in a replication environment with failover, if the primary master fails during replication of updates, the proxy server fails over to the peer server. The data might not be entirely consistent because it is possible for the peer server to not get all the updates made to master.

Directory server instance creation fails

IBM Security Directory Server might fail to create a directory server instance.

On AIX, Solaris, and Linux systems, IBM Security Directory Server might fail to create a directory server instance because of one of the following reasons:

- Not enough disk space in the /home or /export/home directory
- The root user might not have write permission on the /home or /export/home directory

Unable to log on to a system

When migrated users use the LDAP operating system authentication mechanism, they might not be able to log on to the system. Follow the steps to work around this limitation.

IBM Security Directory Server does not support the password encryption mechanism that UNIX supports. Hence, the migrated users might not be able to log on to the system by using LDAP operating system authentication mechanism.

IBM Security Directory Server supports CRYPT and MD5 encryption schemes. However, the UNIX system uses a mix of MD5 and CRYPT password encryption scheme, which IBM Security Directory Server does not support.

You can use one of the following workaround for this problem:

- LDAP administrator can reset the user password for the migrated users on the LDAP system.
- Create new users on the LDAP system for LDAP - operating system authentication.

Propagated schema updates rejected

In some scenarios, a back-end server that is configured as the primary write server might be from earlier versions of IBM Security Directory Server. In this case, the back-end server rejects the propagated schema updates with an error.

When schema updates are requested by a global administrator group member, the schema updates are first applied to the IBM Security Directory Server Proxy Server, Version 6.3 or later. Then, the updates are propagated to the back-end servers. In some cases, the back-end server that was created by using the earlier versions of IBM Security Directory Server might be configured as the primary write server. Then, the back-end server rejects the schema updates with an appropriate error.

You might have an environment where back-end servers are configured by using multiple versions of IBM Security Directory Server. In this scenario, you must configure the back-end server that must be set as primary write server by using the IBM Security Directory Server, Version 6.3 or later. This configuration ensures that schema updates that are propagated from the IBM Security Directory Server Proxy Server, Version 6.3 or later are applied to the back-end servers.

Accessibility tool is unable to read messages in the Configuration Tool

The Accessibility tool, JAWS, is not able to read the message that is displayed on two dialog boxes of the Configuration Tool, which is a limitation.

The JAWS tool is unable to read the message that is displayed on two dialog boxes because of the limitation in the design that implements messages. The following messages that are displayed on the dialog boxes are not read by the JAWS tool:

- Configuration of the instance was changed, causing this task to become invalid. Would you like to dispose this task?
- Are you sure you want to close this window?

General troubleshooting

Use the workaround and solutions to resolve general issues in IBM Security Directory Server, version 6.1 and later directory servers.

IBM Installation Manager generates an error when the GSKit repository contains multiple installable

To provide secure communication mechanism, IBM Security Directory Server uses GSKit. IBM Security Directory Server requires both GSKit SSL package and GSKit crypt package to be installed on the computer.

For the installation of GSKit with IBM Installation Manager, you must provide a path that contains GSKit installable. If the path contains installable for multiple GSKit, version 8.0 or later, then IBM Installation Manager generates an error when installing GSKit.

The reason for the error is that GSKit, version 8.0 or later, are installed at the same location by default. IBM Installation Manager might not be able to sequence the order of versions and install an appropriate version.

To avoid such problems, you must store SSL and crypt packages for a GSKit version in a directory. You can use the same directory to store both 32-bit and 64-bit GSKit packages of a version if they contain unique file names.

Instance owner unable to access core file

The instance owner is sometimes unable to access the core file for a core file that is produced during server initialization. Follow the steps to work around this issue.

If the root user starts the server, a core file might be 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 required.

This problem occurs only on AIX, Linux, and Solaris operating systems.

Key labels do not match

If the key labels in the .kdb file and ibmslapd.conf file do not match, follow the steps to resolve this error.

If the key label in the SSL key database certificate does not match the key label in the directory server configuration file (ibmslapd.conf), the following error occurs:

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 information about how to create a self-signed key database certificate, see the *IBM Security Directory Server Administration Guide*.

GSKit certificate error

If the GSKIT fails with an error when you try to import a signer or personal certificate, follow the steps to troubleshoot and resolve this error.

When you import a signer or personal certificate from an external certificate authority (CA) such as Entrust, the GSKIT might fail with the following error:

An error occurred while receiving the certificate from the given file.

The problem might occur because 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 method is one way to obtain the root certificate.

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 do the following steps:

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: 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

The server instance might fail to start because of incorrect file permissions. You must ensure that the file permissions are readable by the user ID, **idsldap**.

On AIX, Linux, and Solaris systems, file permissions are frequently altered inadvertently by the actions of copying or editing a key database file. Because these actions are generally done as the user ID root, file permissions are set for the user root. For the directory server instance to use 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 host name is incorrect

The server instance fails to start because the localhost host name is not set correctly. You must ensure that the host name meets the specified requirements.

The localhost host name must correspond to the local loopback address of 127.0.0.1. If the localhost is renamed or the TCP/IP address changes, the directory server instance does not start.

Server instance cannot be started except by instance owner

If a user other than the instance owner cannot start a server instance, you must verify the group and access rights of the user.

On AIX, Linux, and Solaris systems, a user other than the directory server instance owner might not be able to start the directory server instance. The user must meet the following requirements to start the directory server instance.

- 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.

For more information about users and groups, see the *IBM Security Directory Server Installation and Configuration Guide*.

Error opening slapd.cat file on Windows systems

If an error occurs when you open slapd.cat file on a Windows system, you must check the *NLSPATH* variable.

On Windows systems, you might receive an error that includes the following message:

```
Error opening slapd.cat
Plugin of type DATABASE is successfully loaded from
C:/Program Files/IBM/LDAP/V6.2/bin/libback-config.dll.
Error opening rdbm.cat
```

If this error occurs, check the *NLSPATH* environment variable. The installation program sets the *NLSPATH* environment variable as a system environment variable. However, if the *NLSPATH* variable is also set as a user environment variable, the user *NLSPATH* environment variable overrides the system setting.

Append the *NLSPATH* information from the system environment variable to the information in the user environment variable.

DSML file client produces error

When a user tries to connect to an LDAP server that does not use SSL, the DSML file client might give an error. This error is not serious.

The DSML file client produces the following error when the DSML file client is set up to communicate by using SSL. The error occurs when 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 serious and the output XML file is generated.

Non-default log files need valid path

If you want to store your log files in a nondefault path, you must ensure that the path exists and is valid.

You must create the directory before you can configure the log files.

Null searches retrieve entries of deleted suffixes

If you deleted a suffix without first removing the entries of the suffix from the database, those entries are returned by the null search. The entries are returned even though the suffix no longer exists.

A null search (`ldapsearch -s sub -b "" objectclass=*`) returns all the entries that are found in the database.

Error occurs with the `idsldapsearch` command

The `idsldapsearch` command with `-h` option gives an error with the DIGEST-MD5 mechanism. Follow the steps to resolve this error.

The DIGEST-MD5 SASL bind mechanism requires the client to be able to resolve the fully qualified host name of the server. If the client cannot resolve the server's fully qualified host name, the bind fails with an `LDAP_PROTOCOL_ERROR`. To correctly resolve the host name, you might be required to make system changes or make DNS configuration changes, such as enabling reverse DNS mapping.

For example, AIX, Linux, Solaris, and HP-UX (Itanium) systems have lines in the `/etc/hosts` file with the syntax:

IP address fully qualified distinguished name alias

This syntax is used to define the local host name 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, which causes DIGEST-MD5 to fail.

For the DIGEST-MD5 mechanism to work correctly, the syntax must be similar to the following syntax:

`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.

Server behavior when language tags are disabled

To avoid potential problems and unexpected behavior after you enable language tags, you must not disable the language tags.

After you enable the language tag feature, if you associate language tags with the attributes of an entry, the server returns the entry with language tags. This behavior occurs even if you later disable the language tag feature. The behavior of the server might not be what the application is expecting. Hence, to avoid potential problems, do not disable the language tag feature after it is enabled.

Key database certificate

You must create the key database certificate before you set up SSL.

Before you set up SSL communications on your server, you must use the GSKit utility, **keyman**, to create the necessary certificates. For more information about **keyman**, see the *IBM Security Directory Server Administration Guide*.

idsbulkload hangs during parsing phase

If the **idsbulkload** seems like it is hanging during the parsing phase, you can resolve this issue by changing some variable values.

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 member groups 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** seems like it is hanging before it shows 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, and Solaris systems: `export BULKLOAD_REPORT_CHUNK=100`
- On Windows systems: `set BULKLOAD_REPORT_CHUNK=100`

idsbulkload then reports 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.
..
```

Size of log file exceeds the system file size limit

A directory server might fail if the size of any log file exceeds the system file size limit. This failure typically occurs when tracing is enabled on the server.

Unable to connect to directory server over SSL

When you use the **idsxinst** tool to copy an instance you might not be able to connect to a directory server over SSL. Follow the steps to resolve this problem.

The reason for this problem might be incorrect configuration. To resolve this problem, complete the following steps:

1. Verify that GSKit is installed on the server.
2. Verify that the `gskkm.jar` file is present in the `DS_ldap_home/java/jre/lib/ext` directory.
3. In the `java.security` file under the `DS_ldap_home/java/jre/lib/security` directory, check if the CMS provider entry exists. If the entry does not exist, add this entry in the `java.security` file by entering the following statement:
`security.provider.X=com.ibm.spi.IBMCMSProvider`

where, X is the next number in the order.

4. Ensure that `/lib` exists in the system path.

5. When you connect to source server over SSL, providing the 'Key name' is not mandatory and can be left blank.

Directory server fails to start after running **bulkload**

When you run ldap operations after you run **bulkload**, the directory server fails to start or shows an error. Follow the steps to troubleshoot and resolve this issue.

After performing **bulkload**, if the directory server fails to start or displays error when performing LDAP operations, it could be because of one of the following reasons:

- Check the log file, `db2diag.log`, if there is an error that states ACCESS TABLE WHEN IN RESTRICTED STATE. This means that loading data or **bulkload** was not complete or was unsuccessful.
- The table is in the "Load Pending" or "Locked" state. A previous LOAD attempt on the table might have resulted in failure. Accessing the table is not allowed until the LOAD operation is restarted or terminated.

Consider the following options to rectify the problem:

- Stop or restart the failed LOAD operation on the table by issuing LOAD with the TERMINATE or RESTART option.
- Check if the `bulkload_status` file is present. This file is created in the home directory of the instance. If this file is present, it means that **bulkload** was unsuccessful. Check the file for errors and rectify it, and try running the bulk load utility again.

The **idsadsrun** tool fails

The **idsadsrun** tool might fail for some instances when run simultaneously for multiple instances on the same system. You must install the latest IBM Tivoli Directory Integrator fix pack to resolve this issue.

When you run the **idsadsrun** tool simultaneously for multiple instances on the same system, you might the following exception:

```
org.apache.derby.client.am.DisconnectException: java.net.ConnectException :  
Error opening socket to server host_name on  
port port_number with message : Connection refused
```

You must apply the latest available fix pack to resolve this error.

To get this fix, go to IBM Tivoli Directory Integrator support site:
<http://www-306.ibm.com/software/sysmgmt/products/support/IBMDirectoryIntegrator.html>

Startup messages are displayed in two different locales

On Windows operating system, directory server startup messages might get displayed in two different locales. This problem usually occurs when a language other than English is specified for the directory server instance.

The cause of this problem might be because different locales are configured on the respective systems of the currently logged in user and the instance owner.

You can consider one of the following ways to rectify this problem:

- Set the *LANG* environment variable explicitly to the language you want use. For example, set *LANG=de_DE* (or any other supported language). You must then start the server from the same window.
- Modify the regional and language settings on the Regional and Language Options dialog box. Ensure that both the currently logged in user and the instance owner have the same set of regional and language settings. Verify that the server messages are now displayed in the same language.

Unable to open a new connection for an LDAP client

You might be unable to open a new connection for an LDAP client to connect to a directory server instance. This issue is specific to the directory server instance that is running on a Linux or Solaris operating system. You can work around this restriction by increasing the limit on open file descriptors.

On Linux and Solaris operating system, there is a limit on the maximum number of file descriptors that can be opened by a process. The default value of the maximum number for open file descriptors is 1024 for Linux operating systems and 256 for Solaris operating systems.

An IBM Security Directory Server, Version 6.1 instance uses 15 file descriptors for logging messages. So on Linux, an IBM Security Directory Server, Version 6.1 instance stops accepting new connections after 1009, that is, 1024 – 15 concurrent client connects. Whereas on Solaris, an IBM Security Directory Server, Version 6.1 instance stops accepting new connections after 241, that is, 256 – 15 concurrent client connects. If an error is encountered when new connections are opened, an appropriate message is logged. This error does not affect any existing connections; only new LDAP clients fail to connect to the directory server.

To increase the maximum open file descriptors, user must issue the following command and restart the server from the same command prompt.

```
#ulimit -Hn number of connections
```

Note: The performance with a high number of concurrent client connections depends on the hardware and the operations that are run. With thousands of concurrent client connections that are sending operations simultaneously, the performance of the directory server might decrease.

Error occurs when you deploy with idsideploy tool

When you deploy a replica or a peer in a replication environment by using the **idsideploy** tool, an error might occur. You can resolve this issue by ensuring that there is only one replication subentry.

When you deploy a replica or a peer in a replication environment by using the **idsideploy** tool, if the tool detects more than one replication subentry that contains the same **serverID** value for the attribute **ibm-replicaServerId** with the attribute **ibm-replicationServerIsMaster** set to true, the tool gives an error.

For any replication context, multiple replication subentries are not required, only one replication subentry is required. For example, if the entries are made as shown in the following example, **idsideploy** fails.

```
dn: ibm-replicaServerId=Peer1,ibm-replicaGroup=default, ou=ouunit1, o=sample
objectclass: top
objectclass: ibm-replicaSubentry
ibm-replicaServerId: Peer1
ibm-replicationServerIsMaster: true
```

```
cn: Peer1
description: Peer1

dn: cn=Peer1_entry,ibm-replicaGroup=default, ou=ounit1, o=sample
objectclass: top
objectclass: ibm-replicaSubentry
ibm-replicaServerId: Peer1
ibm-replicationServerIsMaster: true
cn: Peer1_entry
description: Peer1
```

To rectify the problem in this example, you must create only one entry.

Error occurs because environment variable values contain spaces

The **idsadscfg**, **idssnmp**, and **idslogmgmt** tools might give errors if the environment variable values contain spaces. Ensure that you do not use spaces in the value of the environment variables.

If you installed IBM Tivoli Directory Integrator in a different location other than the default location, set the following environment variable:

- For the Log management (**idslogmgmt**) tool, Active Directory synchronization (**idsadscfg**), and SNMP (**idssnmp**) tools function correctly, you must explicitly set the *IDS_LDAP_TDI_HOME* environment variable to point to the directory where you installed IBM Tivoli Directory Integrator.

The value that you set for the environment variable *IDS_LDAP_TDI_HOME* must not have space or double quotation marks, otherwise the tools do not work properly. On Windows, the tools work properly when tilde, “~” (that is, short path or path with no spaces) is used.

The idsadsrun tool stops during synchronization

The **idsadsrun** tool stops and exits when it attempts to synchronize Active Directory and IBM Security Directory Server after the directory server is restarted. You can restart the Active Directory synchronization to run in real-time synchronization mode.

When you run **idsadsrun** tool in full synchronization mode after you configure Active Directory synchronization solution, it is observed that even if the directory server instance that is associated with Active Directory synchronization stops, the **idsadsrun** tool remains in active state. However, when Active Directory is updated for entries, it is observed that instead of synchronizing with the directory server instance, the tool stops and exists.

After the **idsadsrun** tool runs a pass of full synchronization of directory server instance with Active Directory, it does real-time synchronization. If a pass of full synchronization is done and real-time synchronization is running, restart the Active Directory synchronization solution to run in real-time synchronization mode.

The idsadsrun utility fails to synchronize

The **idsadsrun** utility might fail to synchronize as is from Active Directory Server to a directory server instance. You must check the system configuration and Active Directory settings to troubleshoot this issue.

When you run the **idsadsrun** utility to synchronize many users and groups entries (for example, more than 100000 users and 10000 groups) from Active Directory

server to the directory server instance, you might observe that all entries in Active Directory Server are not synchronized in the directory server instance. It is observed that some user entries and their corresponding group entries might fail to sync.

This issue might occur depending on the system configuration and the JVM parameters that are tuned on your system. For example, for a system with 1-GB RAM, the parameter values can be `Xms254m-Xmx1024m`. You can tune the parameters that are based on your system configurations.

For best results, users can also configure Microsoft Active Directory and set the parameters to the following values:

<code>MaxPoolThreads</code>	<code>4</code>
<code>MaxDatagramRecv</code>	<code>4096</code>
<code>MaxReceiveBuffer</code>	<code>10485760</code>
<code>InitRecvTimeout</code>	<code>120</code>
<code>MaxConnections</code>	<code>5000</code>
<code>MaxConnIdleTime</code>	<code>900</code>
<code>MaxPageSize</code>	<code>1000000</code>
<code>MaxTempTableSize</code>	<code>10000</code>
<code>MaxResultSetSize</code>	<code>262144</code>
<code>MaxNotificationPerConn</code>	<code>5</code>
<code>MaxValRange</code>	<code>1500</code>

The `idscfgdb` command fails to configure a database

The `idscfgdb` command might fail to configure a database for a directory server instance on Red Hat Enterprise Linux (RHEL) v4, 64-bit operating system. You can update the kernel parameter to resolve this problem.

When you configure a database for a directory server instance by using the `idscfgdb` command on RHEL 4 64-bit operating system, the tool might exit with the following error messages:

```
GLPCTL026I Creating database: 'mydata'.
GLPCTL028E Failed to create database: 'mydata'. The failure might
have occurred because the system was not set up correctly before using the tool.
GLPCTL011I Stopping database manager for the database instance: 'mydata'.
GLPCTL012I Stopped database manager for the database instance: 'mydata'.
GLPCDB004E Failed to add database 'mydata' to directory server instance:
'mydata'.
GLPCDB026W The program did not complete successfully. View earlier error messages
for information about the exact error.
```

In this situation, the `db2diag.log` file might contain the following information:

```
2008-03-20-11.33.32.471455+330 I5410E982LEVEL: Error
PID: 3214TID: 182960860768 PROC : db2fm
INSTANCE: mydataNODE : 000
FUNCTION: DB2 Common, Generic Control Facility, gcf_stop, probe:30
MESSAGE : ECF=0x9000036D=-1879047315=ECF_FM_DB2FMD_PROCESS_NOT_EXIST
There is no fault monitor daemon running
CALLED: OS, -, open
RETCODE : ECF=0x9000001A=-1879048166=ECF_FILE_DOESNT_EXIST
File doesn't exist
CALLSTCK:
[0] 0x0000002A956FF982 /opt/ibm/db2/V9.5/lib64/libdb2osse.so.1 + 0x1A7982
[1] 0x0000002A956FF82F ossLogRC + 0x6B
[2] 0x0000002A9BCBDB01 gcf_stop + 0x42B
[3] 0x0000002A95DB9559 _ZN9GcfCaller4stopEP12GCF_PartInfomP11GCF_RetInfo + 0x105
[4] 0x000000000405708 main + 0x17F0
[5] 0x0000003049D1C3FB __libc_start_main + 0xDB
[6] 0x000000000403E7A __gxx_personality_v0 + 0x9A
[7] 0x0000000000000000 ?unknown + 0x0
```

```

[8] 0x0000000000000000 ?unknown + 0x0
[9] 0x0000000000000000 ?unknown + 0x0

2008-03-20-11.33.32.472141+330 I6393E948LEVEL: Error
PID: 3214TID: 182960860768 PROC : db2fm
INSTANCE: mydataNODE : 000
FUNCTION: DB2 Common, Fault Monitor Facility, db2fm, probe:170
MESSAGE : ECF=0x90000349=-1879047351=ECF_FM_FAIL_TO_STOP_GCF_FM
Failed to stop the GCF fm module
CALLED: DB2 Common, Generic Control Facility, GcfCaller::stop
DATA #1 : signed integer, 8 bytes
0
DATA #2 : unsigned integer, 8 bytes
1CALLSTCK:
[0] 0x0000002A956FF982 /opt/ibm/db2/V9.5/lib64/libdb2osse.so.1 + 0x1A7982
[1] 0x0000002A956FF883 ossLogRC + 0xBF
[2] 0x0000000000405772 main + 0x185A
[3] 0x0000003049D1C3FB __libc_start_main + 0xDB
[4] 0x0000000000403E7A __gxx_personality_v0 + 0x9A
[5] 0x0000000000000000 ?unknown + 0x0
[6] 0x0000000000000000 ?unknown + 0x0
[7] 0x0000000000000000 ?unknown + 0x0
[8] 0x0000000000000000 ?unknown + 0x0
[9] 0x0000000000000000 ?unknown + 0x0

```

To resolve this problem:

1. Update the kernel parameter, **kernel.shmmax**. For example:
`kernel.shmmax = 3221225472`
2. Run the **idscfgdb** tool again.

The idscfgdb command fails with error code GLPCTL028E

The **idscfgdb** command might fail with error code GLPCTL028E while it is creating a database. You might be required to tune the kernel parameters to resolve this issue.

On AIX, Linux, and Solaris systems, the **idscfgdb** command might fail while it is creating a database.

An example of the db2cli.log file with the information logged:

```

retcode = 1478; state = "01626"; message = "SQL1478W
The defined buffer pools could not be started.
Instead, one small buffer pool for each page size supported by DB2 has been started.
SQLSTATE=01626

```

An example of the db2diag.log file with the information logged:

```

MESSAGE : ZRC=0x850F0005=-2062614523=SQL0_NOSEG
          "No Storage Available for allocation"
          DIA8305C Memory allocation failure occurred.
DATA #1 :
Unable to attach 3 segments totalling 2478440448 bytes starting at address
0x0000000000000000. One possible cause may be an improper setting for the
shmmax Linux kernel tuneable.

```

To know more about tuning kernel parameters on Linux systems, see [Modifying kernel parameters \(Linux\)](#). To know more about tuning kernel parameters on Solaris, see [Modifying kernel parameters \(Solaris operating system\)](#).

DMS cooked table space size is not extended exactly

The DMS cooked table space size might not be extended exactly in the multiples of the value that you specify with the **-z** option of the **idscfgdb** command. The database manager strives to maintain consistent growth across table space containers. Hence, the actual value that is used to extend the table space might be slightly lesser or greater than the value specified.

Compatibility issue with Common Auditing and Reporting Service (CARS) 6.0.1 server

The CARS logging feature that is provided with IBM Security Directory Server, version 6.2 uses the CARS 6.1 client. Therefore, the CARS 6.1 server is required for using CARS 6.1 clients. Any version of CARS server other than 6.1 is not compatible with the CARS logging tool.

Problem with monitoring server instances on a Solaris system

On a Solaris system, you might face a problem with monitoring directory server instances with an SNMP agent. The problem might occur with an SNMP agent that tries to log on with SSH from IBM Tivoli Directory Integrator.

You must start an rsh session on the Solaris system and then try logging with rsh on to the Solaris system. After you log on to the Solaris system, you can monitor directory server instances by using an SNMP agent.

The idsdbrestore utility displays error messages

The **idsdbrestore** utility displays error messages if the `ldapdb.properties` file is modified. You must replace the `ldapdb.properties` file to resolve this issue.

The **idsdbrestore** utility refers to the `etc/ldapdb.properties` file in the IBM Security Directory Server installation location and not the instance-specific `ldapdb.properties` file in the `install-home/idsslapped-instance-name/etc` directory during a restore operation.

If a user updated or modified the **currentDB2InstallPath** parameter in the `ldapdb.properties` file to a different DB2 installation path or to a different DB2 major version after the directory server instance creation, error messages are displayed when you run a restore operation.

To resolve this problem, user can temporarily copy the `install-home/idsslapped-instance-name/etc/ldapdb.properties` file to the `etc` subdirectory in the IBM Security Directory Server installation location before you run a restore request with the **idsdbrestore** utility. After **idsdbrestore** completes the request, restore the original `ldapdb.properties` file.

The idxinst tool fails

The **idxinst** tool might fail to run or might not display the directory server instances if the file sets for the configured locale are not installed.

If the system does not have the requisite file sets or file sets for the configured locale are corrupted, one of the following problems might occur:

- The **idxinst** tool might fail to run and might generate a core dump file.
- If the **idxinst** tools starts, it might not display the directory server instances present on the system.

To resolve this problem, ensure that all the file sets required for the configured locale are installed.

File path causes backup and restore to fail

The backup and restore operations with the Configuration Tool might not function because the file path is not valid.

When you enter paths on the graphical user interface (GUI) tools, ensure that the path specified can be represented on the system. The file path string must be representable in the system's local code page as the GUI translates the Unicode input to local code page. For example, if the Unicode input for the path contains Chinese characters on a system with French locale, the translated file path is not valid.

Directory server instance starts in config-only mode

A directory server instance might start in config-only mode when you migrate from an earlier version by using Instance Administration Tool. When you provide values for migration by using Instance Administration Tool (**idsxinst**), ensure that the input values are representable in the system local code page.

The warning message GLPSRV147W is displayed

In an IBM Security Directory Server environment, a warning message with the message code GLPSRV147W might be displayed. This message might be because of the default value of write timeout that is set to 10 seconds.

If you see this error frequently for your IBM Security Directory Server environment, you must consider increasing the write timeout value by modifying the **ibm-slapdWriteTimeout** attribute under the entry DN cn=Connection Management, cn=Front End, cn=Configuration.

You can either use the Web Administration Tool or the **ldapmodify** command to change the value of **ibm-slapdWriteTimeout**. To change the value, issue the **ldapmodify** command of the following format:

```
#idsldapmodify -D adminDN -w password -i filename
```

where *filename* contains:

```
dn: cn=Connection Management,cn=Front End, cn=Configuration
changetype: modify
replace: ibm-slapdWriteTimeout
ibm-slapdWriteTimeout: 120
```

Platform-specific issues

Use the workaround and solutions to resolve issues that are specific to an operating system, such as AIX, Windows, or Solaris systems.

AIX operating system

This troubleshooting information applies only to the AIX operating system.

Problem with MALLOCTYPE=buckets

Before you set MALLOCTYPE to buckets on the AIX 5.2 operating system, ensure that you install the patch for APAR IY50668. Otherwise, the LDAP server might fail with a core file.

Verification of 64-bit AIX hardware

The server on AIX requires 64-bit hardware. Run the **bootinfo** command to verify that your 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.

Verification of 64-bit AIX kernel

You must verify that the 64-bit AIX kernel is installed and running. Run the commands to verify that your AIX system has the 64-bit kernel.

To verify that you have the 64-bit kernel (`/usr/lib/boot/unix_64`) installed and running, run the following command:

```
bootinfo -K
```

In addition, if you type the command `lsattr -El proc0`, the output of the command returns the type of processor for your server.

Note: If the hardware is 32-bit, then you can have only a 32-bit kernel. You cannot have a 64-bit kernel. If the hardware is 64-bit, then you can have either a 32 or 64-bit kernel.

To switch between a 32-bit and 64-bit mode at the operating system level on AIX 5.3:

When you install the operating system, go to **Additional features** and specify 64-bit mode. (The default is 32-bit mode.) To switch from 32-bit mode to 64-bit mode, use the following commands:

```
# ln -sf /usr/lib/boot/unix_64 /unix
# ln -sf /usr/lib/boot/unix_64 /usr/lib/boot/unix
# bosboot -ad /dev/ipldevice
# shutdown -Fr
# bootinfo -K
```

The kernel is now in 64-bit mode.

To switch from 64-bit mode to 32-bit mode, use the following commands:

```
# ln -sf /usr/lib/boot/unix_mp /unix
# ln -sf /usr/lib/boot/unix_mp /usr/lib/boot/unix
# bosboot -ad /dev/ipldevice
# shutdown -Fr
# bootinfo -K
```

The kernel is now in 32-bit mode.

Error occurs when you run db2start

You might get an error when you run **db2start** on an AIX system. You must turn on asynchronous I/O to resolve this error.

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 error 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. Take one of the following actions:
 - Restart your system.
 - Run **smitty aio** and move the cursor to **Configure defined Asynchronous I/O**. Then, press Enter.

The **db2start** command now works.

Kerberos service fails on AIX 6.1 with workload partitions (WPARs)

On AIX 6.1 that is configured with workload partitions (WPARs), a Kerberos service might fail to start. This problem is because of the IBM Network Authentication Service version.

When WPARs are configured on an AIX 6.1 system, a Kerberos service might fail to start and might display the following message:

```

Starting krb5kdc...
Unable to bind server socket on port 88.
Unable to initialize network.
Status 0x44 - The socket name is not available on this system..
krb5kdc could not be started.

```

This problem is because of a limitation with IBM Network Authentication Service (NAS) 1.4. To resolve this problem, use versions that are later than IBM NAS 1.4.

Utilities like **ldif2db** and **bulkload** fail

IBM Security Directory Server utilities like **ldif2db** and **bulkload** might stop or exit on an AIX 6.1 system. You can resolve this problem by installing AIX 6.1 fix pack 1 or later.

When the utility stops or exits, an error similar to the following trace of the **ldif2db** utility might be logged:

```

253:20:40:16 T1 K2244835 do_iconv_open: local_codepage=NULL
253:20:40:16 T1 K2244835 xlate_utf8_to_localcp: inlen=8
253:20:40:16 T1 K2244835 xlate_utf8_to_localcp: rc=0
GLPRDB002W ldif2db: 50 entries have been successfully added out of 50 attempted.
253:20:40:16 T1 K2244835 vPrintMessage: catid=2, level=2, num=2.
253:20:40:16 T1 K2244835 do_iconv_open: local_codepage=NULL

```

```

253:20:40:16 T1 K2244835 xlate_utf8_to_localcp: inlen=8
253:20:40:16 T1 K2244835 xlate_utf8_to_localcp: rc=0
253:20:40:16 T1 K2244835 close_one_backend: calling be->be_close
253:20:40:16 T1 K2244835 close_one_backend: calling be->be_close
253:20:40:16 T1 K2244835 calling config_close...
253:20:40:16 T1 K2244835 close_one_backend: calling be->be_close
253:20:40:16 T1 K2244835 calling rdbm_close...
253:20:40:16 T1 K2244835 leaving rdbm_close...
./ldif2db[1040]: 880864 Segmentation fault(coredump)

```

This is a problem with AIX 6.1 base level. To resolve this error, use AIX 6.1 with FP1 or later fix levels.

The idsidrop command generates Java core

The **idsidrop** command might generate Java core when a directory server instance is dropped from IBM POWER7 system with AIX 7.1. Follow the steps to work around this problem.

IBM Security Directory Server, Version 6.3 uses IBM DB2, version 9.7 fix pack 2, which includes JRE that is used by IBM Tivoli Monitoring Agent for DB2. This JRE version is not supported on IBM POWER7 system.

Workaround

If you already installed DB2 9.7 manually on the system for use with a directory server, then you must uninstall the IBM Tivoli Monitoring Agent for DB2.

Procedure

1.

Using the DB2 instance owner credentials stop all the Monitoring Agent for DB2 processes. Run the following command:

```
DB2_Directory/itma/bin/itmcmd agent -o instance stop ud
```

where, *DB2_Directory* is the directory where DB2 copy of Tivoli Monitoring Agent is installed.

Note: If multiple DB2 instances are being monitored, then there can be multiple kuddb2 processes that must be stopped.

2. On AIX and Linux operating systems, uninstall Tivoli Monitoring Agent for DB2. Run the following command:

```
DB2_Directory/itma/bin/uninstall.sh REMOVE EVERYTHING
```

For more information about uninstalling IBM Tivoli Monitoring for Databases, see the IBM DB2 documentation at <http://pic.dhe.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.doc/welcome.html>.

Windows operating system

This troubleshooting information applies only to the following Windows operating systems: Windows 2000, Windows Server 2003 Enterprise, Windows Server 2003 R2 Datacenter Edition, Windows Server 2008, and Windows XP and Windows 7 client operating systems.

System environment variables for non-English installation wizard

You must set the *LANG* and *LC_ALL* system environment variables for the installation wizard in non-English languages.

For the installation wizard to use 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.

For a list of Microsoft locale values, see the Microsoft website at <http://www.microsoft.com>.

UTF-8 supplementary characters do not display correctly

Some UTF-8 supplementary characters might not be displayed correctly. You can install one of the East Asian language kits to resolve this problem.

IBM Security Directory Server supports UTF-8 (Unicode Transformation Format, 8-bit form) to use Unicode characters. UTF-8 contains MS932 (Shift JIS) characters plus supplementary characters that are not defined in MS932. Supplementary characters might be displayed as square box in Internet Explorer running on Windows 2000. See Figure 1.

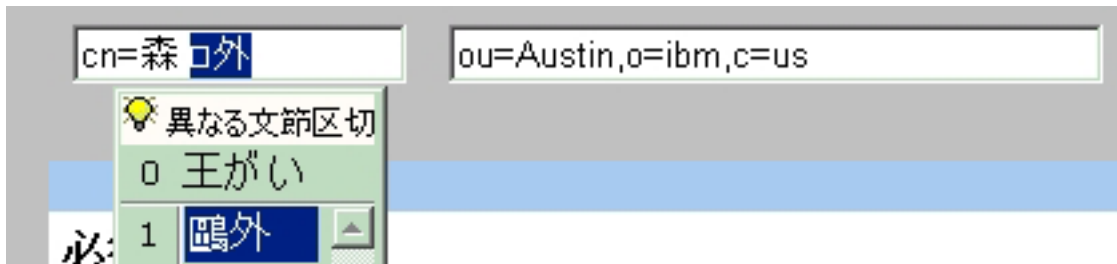


Figure 1. Unicode code point U+9DD7 displayed as a square

If this problem 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 was not observed in Windows XP.

Communications error: Exceeding 64 connections/OCH

On Windows systems, if clients are generating many connections to the server and the connections are being refused, the server might log error messages. Follow the steps to resolve these errors.

The error might be similar to the following log in the `ibmslapd.log` file:

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

If you see these errors, complete the following steps:

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.

Error starting server at operating system startup

In IBM Security 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 might occur when you restart the computer.

An error occurs 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 must be run after all the services are started so that DB2 is 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\ibmdirectl [-h hostname] [-D adminDN] [-w password]
[-p portnumber] start -- [ibmslapd options]
```

Note: Be sure that the **Startup Type** for the **IBM Security Directory Administration server** entry in the Services window is set to Automatic. If it is not, the administration server control program (**ibmdirectl**) does not work.

Backup and restore from a remote drive fails

On a Windows operating system, when you back up, restore, or load by using DB2 utilities, the operation might fail. An error message that the path is not valid is displayed when you specify a directory that is mapped as remote drive. Understand the Windows access restrictions and use the troubleshooting information to resolve this problem.

For example, you might try to back up, restore, or load to or from a directory that is mapped as remote drive such as `H:\MyFolder\test`. The error message that is displayed is "SQL2036N The path for the file or device "*file_or_devicename*:" is not valid".

There are two different causes of this error:

1. You specified an invalid shared drive in the command.
`db2 backup db mydata to H:\`

The following error message is displayed: "SQL2036N The path for the file or device "H:\" is not valid."

2. You specified a valid UNC name for the mapped drive but did not use the right user ID.
`db2 backup db mydata to \\MyFolder\test`

The following error message is displayed: "SQL2036N The path for the file or device "\\MyFolder\\test "is not valid."

If the DB2 backup fails, the db2diag.log looks like:

```
database_utilities sqlubcka Probe:0 Database:mydata
```

Starting a full database backup.

```
2006-06-10-10.42.09.175000 Instance:DB2 Node:000
PID:2404(db2syscs.exe) TID:2500 Appid:none
database_utilities sqlUMCTestDevType4Backup Probe:60
```

```
Media controller -- invaliddevice path: H:\MyFolder\test
```

```
2006-06-10-10.42.09.253000 Instance:DB2 Node:000
PID:2404(db2syscs.exe) TID:2576 Appid:*LOCAL.DB2.030610154019
database_utilities sqlubcka Probe:0 Database:mydata
```

Backup terminated.

This error is because of Windows restriction. Do the following steps to resolve this error:

- You must start the DB2 server by using an existing user ID instead of the default "Local System Account".
- To specify read and write permissions on the network drives, select **Services** in Administrative Tools under Control Panel. Next, in the Services window, select **Properties** from the **Action** menu. Select **Log On** tab and update "Log on as" to indicate a specific user who has the read and write permissions on the network drives.
- In Windows 2000 and Windows XP, you must specify the full qualified UNC name. Do not specify the network share when you run the backup, restore, or load action.
- Use the command:

```
db2 backup db mydata to \\ MyFolder\test
```

instead of

```
net use H: \\ MyFolder\test
db2 backup db mydata to H:
```

The idsicrt tool fails to catalog instance node and exits

On Microsoft Windows Server 2003 R2 Datacenter Edition SP1 or later versions, the **idsicrt** tool fails to catalog instance node and exits. You must set the *DB2INSTPROF* or install the specified DB2 fix pack to resolve this problem.

Before you create a directory server instance, on a system with Microsoft Windows Server 2003 R2 Datacenter Edition SP1 or later versions, set the *DB2INSTPROF* environment variable. For example:

```
SET DB2INSTPROF=PATHNAME
```

where, *PATHNAME*, specifies the absolute path of the directory where the profile of the instance node must be created.

Note: The length of the absolute path must not exceed 70 characters.

You can also download DB2, version 9.5, Fix Pack 2 that has the fix to this problem.

DB2 diagnostic path error in idsdbmaint tool

On Microsoft Windows operating systems, if DB2 v9.5 is installed with the installation wizard, then the **idsdbmaint** tool might display an error that the DB2 diagnostic path cannot be determined. Run the specified DB2 commands to work around this problem.

On Microsoft Windows operating systems, if DB2 v9.5 is installed with the installation wizard then the **idsdbmaint** tool is unable to fetch DB2 diagnostic path. This problem does not occur if DB2 is installed separately (without using the installation wizard). If the **idsdbmaint** tool is not able to fetch DB2 diagnostic path then the index reorganization operation does not have the required impact. This problem is because the **idsdbmaint** tool requires DB2 diagnostic path for the index reorganization and row compression features to function. The **idsdbmaint** tool might display the following error message when this problem is encountered.

GLPDBA048E The diagpath for database instance 'dsrdbm01' could not be determined.

The following is observed when the **idsdbmaint** command is run:

```
c:\>idsdbmaint -I dsrdbm01 -r
GLPWRP123I The program 'C:\Program Files\IBM\LDAP\V.2\sbin\32\dbmaint.exe'
is used with the following arguments '-I dsrdbm01 -r'.
GLPSRV200I Initializing primary database and its connections.
GLPDBA037I Row compression task will be performed.
GLPDBA048E The diagpath for database instance 'dsrdbm01' could not be determined.
```

Workaround

To resolve the issue of DB2 diagnostic path when you install DB2 v9.5 using the installation wizard on Windows operating system, issue the following commands:

```
db2start
db2 connect to instance name
db2 update dbm cfg using DIAPATH C:\instancename
db2 terminate
db2stop
```

Problems with opening Web Administration Tool in Internet Explorer

On Microsoft Windows Server 2008, Web Administration Tool might not start properly with Internet Explorer. You must configure the URL as trusted sites in internet explorer to resolve this problem.

On Windows Server 2008, Web Administration Tool panels might not work properly with Internet Explorer.

To resolve this issue, you must make the URL of Web Administration Tool as part of trusted sites of the browser:

1. Open Internet Explorer.
2. On the **Tools** menu, click **Internet Options**.
3. On the Internet Options dialog box, select the **Security** tab.
4. On the Security tab, select **Trusted sites** and then click **Sites**.
5. In the Add this website to the zone box, enter the URL for Web Administration Tool and then click **Add**.
6. Click **OK**.

Open Web Administration Tool with the Internet Explorer.

IBM Security Directory Server clients on Windows PowerShell

Windows PowerShell is not a supported shell for IBM Security Directory Server. Therefore, when IBM Security Directory Server clients are run on PowerShell, they might not function as expected.

IBM Security Directory Server clients function properly when run on command prompt.

For example, when the following command is run on command prompt, it returns root DSE search results. However, when run on PowerShell, it displays the command usage.

```
idsldapsearch.cmd -p 389 -D cn=root -w root -s base -b "" objectclass=*
```

Log management tool fails to start

On Windows operating systems, the log management tool might fail to start from Web Administration Tool. This problem is related to user privileges.

On Windows systems, the correct way to start the directory server and administration server is through services. However, processes that are started through services are always run with the SYSTEM user's privileges, regardless of the login credentials of the user.

For example, if you create a directory server instance, `myinst1`, log in as the `myinst1` user, and start the server and administration server by using the services, then the processes are run with the SYSTEM user's privileges. In this scenario, the log management tool, **idslogmgmt**, does not work if it is started or stopped from Web Administration Tool. The reason is because log management requires Web Administration Tool to be started or stopped by using the instance owner's credentials. The start or stop log management request from Web Administration Tool goes to the administration server, which is running as the SYSTEM user. Therefore, starting or stopping the log management tool through the administration server fails.

If the administration server is started from the command line by using the credentials of `myinst1`, then the process would run with the `myinst1`'s privileges. Therefore, starting or stopping the log management tool by using Web Administration Tool works.

Error during uninstallation of client packages

Uninstallation of IBM Security Directory Server client packages from a Windows 7 system as a non-privileged user might fail. An error message that access is denied might be displayed. You must run the uninstallation as an administrator.

To uninstall the packages as an administrator, select the "Run as administrator" option from the User Account Control shield. The **Run as administrator** option prompts for the administrator user name and password.

For more information, search with the User Account Control keyword on the Microsoft MSDN Library website at <http://msdn.microsoft.com/en-us/library/ms123401.aspx>.

Large number of schema updates causes server instance to fail

A directory server instance of version 6.3 might fail when it is stressed with too many schema updates. This problem happens when schema updates are made continuously for prolonged period. This problem was observed only on 64-bit Windows Server 2008 systems.

Solaris operating system

This troubleshooting information applies only to Solaris operating systems.

Removal of server components fails

On Solaris 10 that is configured with zones, the removal of server components might fail if they are configured on a different zone. Ensure that you install all dependent client and server packages from the same zone.

When you install IBM Security Directory Server, version 6.2 on Solaris big-zone, all dependent client and server installation packages must also be installed and run from big-zone.

The propagated installation of clients and `srvbase` from Solaris global-zone, and the installation of server components on big-zone does function. However, if the installed IBM Security Directory Server components (clients and `srvbase`) are uninstalled from global-zone, it results in the removal of these components from big-zone and small-zone. Hence, the removal of server component from big-zone might fail.

Instance Administration Tool or Configuration Tool might show error on console

Instance Administration Tool and Configuration Tool requires IBM Java Development Kit to open.

IBM Security Directory Server, version 6.3, fix pack 4 or later include IBM Java 1.6 SR9, which is required by GUI utilities, such as Instance Administration Tool (**idsxinst**) and Configuration Tool (**idsxcfg**). On Solaris 10 operating system on X64 architecture, Instance Administration Tool or Configuration Tool might show the generated error messages on console if user provides wrong data in the fields.

DB2 utility `db2osconf`

On Solaris systems, you can use the DB2 utility **db2osconf** to obtain suggested values for kernel parameters.

The DB2 utility, **db2osconf**, makes recommendations for kernel parameter values based on the size of a system. This command is available only for DB2 on Solaris SPARC systems with 64-bit instances. If zones are configured on a Solaris SPARC system, then the **db2osconf** utility is available only in global zone. On Solaris non-global zones and Solaris x86-64 systems, the **db2osconf** utility is not available. On these systems, you can use the **projmod** command to set the values for kernel parameters, such as limits for shared memory, semaphore ids, and total shared memory.

On Solaris SPARC global zone, use **db2osconf** to obtain suggested values for kernel parameters. An example of the **db2osconf** command and its output is as follows:

```
#db2osconf
set msgsys:msginfo_msgmni = 6144
set semsys:seminfo_semmni = 7168
set shmsys:shminfo_shmmax = 9578697523
set shmsys:shminfo_shmmni = 7168
```

```
Total kernel space for IPC:
0.98MB (shm) + 1.71MB (sem) + 2.08MB (msg) == 4.77MB (total)
```

Here is an example of the **projmod** command on Solaris SPARC that uses the values that are generated by **db2osconf**:


```
projmod -s -K "project.max-shm-memory=(privileged,9578697523,deny)" user.db2inst1
projmod -s -K "project.max-shm-ids=(privileged,7168,deny)" user.db2inst1
projmod -s -K "project.max-msg-ids=(privileged,6144,deny)" user.db2inst1
projmod -s -K "project.max-sem-ids=(privileged,7168,deny)" user.db2inst1
```

Here is an example of the **projmod** command on Solaris x86-64. You can use values that are suitable for your environment.

```
projmod -a -K "project.max-shm-ids=(priv,4k,deny)" user.db2inst1
projmod -a -K "project.max-sem-ids=(priv,4k,deny)" user.db2inst1
projmod -a -K "project.max-shm-memory=(priv,4G,deny)" user.db2inst1
projmod -a -K "project.max-msg-ids=(priv,4k,deny)" user.db2inst1
```

The values of these limits must be set in accordance with the available system resources in your environment. For more information, see Memory management and related concepts in the DB2 v9.x information center.

The idsadsrun utility does not exit when error occurs

On a system with Solaris 10 as the operating system, the **idsadsrun** utility fails to exit when errors are encountered. In such cases, you must use Ctrl C to stop the process.

Migration of server instance fails

The migration of an IBM Security Directory Server, Version 6.1 instance with DB2, Version 9.1 to IBM Security Directory Server, Version 6.3 with DB2 Version 9.7 fails. This failure happens on an AMD Opteron system with Solaris 10 operating system. You must install the specified DB2 fix pack to resolve this problem.

When migration of the directory server is initiated with the **idsimigr** command, it gives the following error message.

```
GLPMIG041E The database name listed in the backed up configuration
file cannot be found on the system.
```

This problem is observed when DB2 Version 9.1 has fix pack level 7 or lower.

Use DB2, Version 9.1, Fix Pack 8 or higher for migration of IBM Security Directory Server, Version 6.1 instance to IBM Security Directory Server, Version 6.3.

Appendix A. Common Base Event features

To create self-managing environment, IBM took the initiative in introducing "Autonomic Computing". Autonomic computing is an open standard based architecture that allows systems to configure, heal, optimize, and protect itself.

To determine the conditions of the different components of the system, the format of the event data must be standardized. Standardized event data is required for the system to resolve its current conditions.

To standardize the format of data for the problem determination architecture, IBM introduced a common format for log and trace information. The format is called the Common Base Event format. This format creates consistency across similar fields and improves the availability to correlate across multiple logs. The Common Base Event format is based on a 3-tuple structured format, which includes:

- Component that is impacted by a situation, or the source
- Component observing a situation
- Situation data, the properties that describe the situation, including correlation information

The 3-tuple format makes it possible to write and deploy resource-independent management functions that can isolate a failing component.

To align IBM Security Directory Server to autonomic computing space, the logs that are produced must conform to the Common Base Event format. For example, logs such as error log and server audit log that are produced by IBM Security Directory Server must conform to the Common Base Event format.

The IBM Common Auditing and Reporting Service (CARS) component uses Common Base Event format. It is a common format for events that are proposed by IBM and IBM Common Event Infrastructure (CEI) technologies to provide an audit infrastructure. The purpose of Common Base Event format is to facilitate effective intercommunication among disparate components within an enterprise. To effectively process audit data, the CARS component requires the audit data to be in the Common Base Event format. CEI is an IBM strategic event infrastructure for submission, persistent storage, query, and subscription to events that conform to the Common Base Event format. The CARS component uses the CEI interfaces for submission of events. These events can be denoted as auditable by using configuration options at the CEI Server. The CEI Server stores them in a CEI XML Event store that meets the auditing requirements.

The CARS component allows staging of data from the CEI XML Event store into report tables. IBM products and customers can provide audit reports that are based on auditable events staged into report tables. The CARS component also supports managing the lifecycle of auditable events, which includes archive, restore, and audit reports on restored archives.

In IBM Security Directory Server, auditing capability is implemented by using the directory server audit plug-in. A user can implement the audit enhancements to write audited data to Common Base Event format. An example is listed:

- The audit data can be read and transformed to Common Base Event format by an external application such as, IBM Tivoli Directory Integrator. The audit data can then be sent over to CARS by using the CEI API or the CARS embeddable Java client.

To implement the example, the settings for this feature in the `ibmslapd` configuration file are retrieved. If the settings are specified, the audit data files are read periodically and converted into Common Base Event format by the log management tool. Depending on the settings, the data that conforms to Common Base Event format can be written to a file, to a CEI server, or both. The data sent to a CEI server is stored in the CEI database and CARS moves the audit data into a CARS database. The data then moves into a staging area for CARS reports or into a database archive for long-term storage.

Common Base Event scenarios

Some special case scenarios that must be considered for events that conform to the Common Base Event specification are related to attributes and log management.

There are some special case scenarios that must be considered for the Common Base Event feature.

Attribute related special case scenarios

Unspecified attribute settings

If a value is set for `ibm-slapdLogEventFileSizeThreshold` and the value of `ibm-slapdLogEventFileMaxArchives` is not specified either in the default entry or in the specific log entry, then archiving occurs but the number of archive files are unlimited.

Attribute settings that are given in wrong format

- If the value provided for `ibm-slapdLogEventFileSizeThreshold` is in the wrong format then an error message is logged and no archiving occurs.
- If the value provided for `ibm-slapdLogEventFileMaxArchives` is in the wrong format then an error message is logged but archiving occurs and the number of archive files are unlimited.
- If the value provide for `ibm-slapdLogEventFileArchivePath` is invalid then the archived file path is in the same directory as the original file's path.

Scenarios that are related to Common Base Event file and log management actions

Out of disk space

If the disk gets full, the log management activity fails. The failure is indicated by an error message that is displayed on the standard output. If possible, the message is also logged in the log.

Archive path errors

- If the archived file cannot be written to the path specified, then an error message is logged on the `idslogmgmt` log.
- If a file with the same name exists in the mentioned archive path, then an error message is logged in the `idslogmgmt` log and the archiving fails. When the next log management occurs, for the operation to succeed the timestamp must be different.

Log archiving and Common Base Event activity interference

In some scenarios, the directory server log management tool is configured to send Common Base Event data to a CEI server and log archiving is also enabled. There is a possibility that the archiving threshold is reached but the log data is not sent to the CEI server.

The log data might not be sent to the CEI server because the CEI server might be down. The reason might also be because the transmission rate to CEI server is lesser than the original log write rate. In such cases, the log archiving is suspended so that data that must be sent to CEI server is not lost. The expected behaviors when this situation occurs with the log management settings are listed.

- The Common Base Event formatted logs are enabled and the value of **ibm-slapdLogEventFileMaxArchives** is set to zero. In this scenario, the Common Base Event file that was to be deleted is kept and the file continues to grow.
- The Common Base Event formatted logs are enabled and the value of **ibm-slapdLogEventFileMaxArchives** is set to a number greater than zero. In this scenario, when the set maximum number of Common Base Event log files are reached, the log file that was to be deleted is kept. The number of Common Base Event archives continues to grow.
- The Common Base Event formatted logs are disabled and the value of **ibm-slapdLogMaxArchives** is set to zero. In this scenario, the log file that was to be deleted is kept and the file continues to grow.
- The Common Base Event formatted logs are disabled and the value of **ibm-slapdLogMaxArchives** is set to a number greater than zero. In this scenario, when the set maximum number of log files is reached, then the oldest archived file that was to be deleted is kept. The number of archives continues to grow.

Log activity of overlapping cycles

A current cycle of log activities might be running for a log when the next cycle of log activities gets triggered. In this scenario, the tool must not allow multiple cycles to overlap. The next cycle must start only after the completion of the first cycle. Thus, different log activity cycles are prevented from interfering and causing data loss.

Appendix B. Support information

To obtain support for IBM products, you can use one or more of the several options such as, knowledge bases and information centers.

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.

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 and most complete information.

To search multiple Internet resources for your product, expand the product folder in the navigation frame to the left and select **Web search**. You can search various resources that include:

- IBM technotes
- IBM downloads
- IBM Redbooks®
- IBM developerWorks
- Forums and newsgroups
- Google

Product 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 website.

The following instructions can help you identify the product fix that might help you resolve your problem:

1. Go to the IBM Software Support website (<http://www.ibm.com/software/support>).
2. Under **Products A - Z**, select your product name. A product-specific support site is opened.
3. Under **Self help**, follow the link to **All Updates**, where you can 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 email 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 are already registered, skip to the next step. If you are not already 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 email 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>).

Contact IBM Software Support

IBM Software Support assists with product defects. Before your contact 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, enroll in Passport Advantage®. These products include, but are not limited to, Tivoli, Lotus®, and Rational® products. They also include DB2 and WebSphere products that run on Windows, AIX, Linux, Solaris, and HP-UX operating systems. You can enroll 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 website (<http://techsupport.services.ibm.com/guides/contacts.html>) and click the name of your geographic region.

- For IBM eServer™ software products, you can purchase a software maintenance agreement by working directly with an IBM marketing representative or an IBM Business Partner. These products include, but are not limited to, DB2 and WebSphere products that run in System z, System p, and System i environments). 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. 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>). Click the name of your geographic region for phone numbers of people who provide support for your location.

Follow the steps to contact 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 must understand and assess the business impact of the problem you are reporting.

Use the following criteria to determine the business impact of your problem:

Table 12. Severity level and their description

Severity	Business impact
Severity 1	Critical business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	Significant business impact: The program is usable but is severely limited.
Severity 3	Some business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	Minimal business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem is implemented.

Describe your problem and gather background information

When you explain a problem to IBM, you must 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, and networking software.)
- Are you currently using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submit your problem to IBM Software Support

You can submit your problem to IBM Software Support either online or by phone.

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. Other users who experience the same problem can also benefit from the same resolutions.

For more information about problem resolution, see [Searching knowledge bases](#) and [Obtaining fixes](#).

Notices

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

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

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

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

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

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

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

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

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

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

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

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

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

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to

IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Index

A

- accessibility xii
- adminaudit.log 5
- administration server
 - audit log 5
 - log 5
- Auditing for performance 83

B

- backup status file
 - dbback.dat 5
- bulkload error log 5
- bulkload.log 5

C

- CBE features 139
- configuration
 - Configuration Tool 40
- core files
 - AIX operating systems 11
 - description 11
 - Linux operating systems 11
 - Solaris operating systems 11
 - Windows operating systems 11
- customer support
 - see Software Support 144

D

- DB2 log 5
- DB2 rollbacks 82
- DB2 troubleshooting 47
- db2cli.log 5
- debugging
 - advanced output 44
 - description 13
 - ldtrc command 13
 - server debug mode 13
- directory server, installation
 - mounting DVD 23
- directory server, media
 - mount DVD 23

E

- education xiii

F

- fixes, obtaining 143

I

- IBM
 - Software Support xiii
 - Support Assistant xiii

- ibmdiradm.log 5
- ibmslapd.log 5
- idsldap group 24
- idsldap user
 - requirements 24
- idsldaptrace utility 13
- idslink.log 27
- idslink.preview 27
- idsslapd trace 81
- idstools.log 5
- information centers, searching to find
 - software problem resolution 143
- installation
 - overview 21
 - prerequisite software 21
 - installation logs 5, 21, 25, 26, 27
 - installation media, directory server
 - DVD mount issue 23
 - installation troubleshooting
 - installation wizard 29
- instance creation
 - idsicrt 37
 - Instance Administration Tool 37
 - troubleshooting 38
- Internet, searching to find software
 - problem resolution 143
- isolation levels 82

K

- knowledge bases, searching to find
 - software problem resolution 143
- Known limitations
 - Partial replication 104

L

- LDAP_DEBUG 13
- LDAP_DEBUG_FILE 13
- ldtrc command 13
- LOGFILSIZ
 - modifying 83

logs

- administration server audit log 5
- administration server log 5
- bulkload error log 5
- DB2 installation
 - AIX 26
 - Linux 26
 - Solaris 26
 - Windows 26
- DB2 log 5
- DB2 uninstallation
 - Windows 26
- GSKit installation
 - Windows 27
- idslink 27
- installation 5, 21
 - AIX 25

logs (continued)

- installation (continued)
 - Linux 25
 - Solaris 25
 - Windows 25
- lost and found log 5
- native packages 27
- overview 5
- server audit log 5
- server log 5
- tools log 5
- lost and found log 5
- lostandfound.log 5

M

- memory leak 87
- memory, adding on Solaris 81
- messages, resolving 2
- migration troubleshooting 33

O

- online
 - publications xi
 - terminology xi

P

- performance troubleshooting 81
- problem determination
 - describing problem for IBM Software Support 145
 - determining business impact for IBM Software Support 145
 - submitting problem to IBM Software Support 145
- problem-determination xiii
- publications
 - accessing online xi
 - list of for this product xi

R

- replication
 - overview 59
 - troubleshooting 59

S

- Secure Sockets Layer (SSL) 88
- server audit log 81
- Server audit log 5
- server log 5
- SLAPD_OCHANDLERS environment variable 82
- Software Support
 - contacting 144

Software Support (*continued*)
 describing problem for IBM Software
 Support 145
 determining business impact for IBM
 Software Support 145
 submitting problem to IBM Software
 Support 145

T

terminology xi
thread stacks 87
tools log 5
trace
 idsslapd 81
training xiii
troubleshooting xiii
troubleshooting features, overview 1

U

uninstallation logs 26
uninstallation troubleshooting 32

W

Web Administration Tool
 troubleshooting 51



Printed in USA

GC27-2752-01

