

IBM Tivoli Directory Server



# Problem Determination Guide

*Version 6.1*



IBM Tivoli Directory Server



# Problem Determination Guide

*Version 6.1*

**Note**

Before using this information and the product it supports, read the general information under Appendix C, "Notices," on page 109.

This edition applies to version 6, release 1, of IBM Tivoli Directory Server and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2005, 2007. All rights reserved.**

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

# Contents

## About this book . . . . . vii

Intended audience for this book . . . . .	vii
Publications . . . . .	vii
IBM Tivoli Directory Server version 6.1 library . . . . .	vii
Related publications . . . . .	viii
Accessing terminology online . . . . .	viii
Accessing publications online . . . . .	ix
Ordering publications . . . . .	ix
Accessibility . . . . .	ix
Tivoli technical training . . . . .	ix
Support information . . . . .	x
Conventions used in this book . . . . .	x
Typeface conventions . . . . .	x
Operating system-dependent variables and paths . . . . .	xi

## Chapter 1. Introduction to problem determination . . . . . 1

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

## Chapter 2. Logging utilities . . . . . 3

## Chapter 3. Other diagnostic tools . . . . . 9

Generating core files . . . . .	9
For Windows operating systems (Dr. Watson debugger) . . . . .	9
For Linux operating systems . . . . .	9
For AIX operating systems . . . . .	10
For Solaris operating systems . . . . .	10
For HP-UX operating systems . . . . .	10
IBM Tivoli Directory Server Support Tool . . . . .	11
Data collected by the Support Tool . . . . .	11
Using the Support Tool . . . . .	13
Location of Support Tool log and collected data . . . . .	14
Server debug mode . . . . .	14

## Chapter 4. Troubleshooting installation and uninstallation . . . . . 17

Product installation overview . . . . .	17
Prerequisite software . . . . .	17
Failures when installing prerequisite software . . . . .	18
The idslsap user and group . . . . .	18
Installing the IBM Tivoli Directory Integrator Server and Administration and Monitoring Console on AIX, Linux, Solaris, and HP-UX systems . . . . .	20
Installable image subdirectories and files . . . . .	21
Installing the Deployment Engine . . . . .	21
Installing the IBM Tivoli Directory Integrator Server installable unit . . . . .	22
Installing the IBM Tivoli Directory Integrator Administration and Monitoring Console installable unit . . . . .	23

After you install . . . . .	24
Uninstalling the IBM Tivoli Directory Integrator installable units . . . . .	25
Installation logs . . . . .	25
Logs for Embedded WebSphere Application Server . . . . .	25
DB2 logs on Windows . . . . .	26
DB2 logs on AIX, Linux, Solaris, and HP-UX . . . . .	27
Tivoli Directory Integrator logs . . . . .	27
idslink log on AIX, Linux, and Solaris operating systems . . . . .	27
GSKit logs on Windows operating systems . . . . .	27
Log files generated for native packages on UNIX operating system . . . . .	28
Troubleshooting . . . . .	29
InstallShield GUI installation . . . . .	29
Operating system utility installation . . . . .	32
InstallShield GUI uninstallation . . . . .	32
If default instance creation fails during the Typical installation . . . . .	32
On Windows operating system, installation might fail giving out message such as "DB2 Install was NOT successful" . . . . .	32
Recovering from a failed InstallShield GUI installation on HP-UX systems . . . . .	33

## Chapter 5. Troubleshooting migration 35

Migration log files . . . . .	35
Kerberos service name change . . . . .	35
Database instance or database in configuration file but no longer on system . . . . .	35
Format of backed-up schema files incorrect . . . . .	36
ibm-slapdPlugin entry in configuration file changed . . . . .	36
If migration fails using ISMP . . . . .	36

## Chapter 6. Troubleshooting instance creation and configuration . . . . . 37

Instance creation overview and common errors . . . . .	37
Instance creation overview . . . . .	37
Common instance creation errors . . . . .	38
Configuration overview and common errors . . . . .	39
Overview . . . . .	39
Common errors . . . . .	40

## Chapter 7. Troubleshooting DB2. . . . . 45

DB2 license file expired . . . . .	45
Database performance is poor . . . . .	46
Recovering from migration failure in DB2 9.1 . . . . .	46

## Chapter 8. Troubleshooting the Web Administration Tool and the application server . . . . . 47

Troubleshooting the Web Administration Tool . . . . .	47
---	----

Corruption of data entered in the Web Administration Tool . . . . .	47
Migrating files when patching or migrating the Web Administration Tool . . . . .	48
Additional login panels fail . . . . .	48
idsldapmodify command puts Web Administration Tool into inconsistent state . . . . .	48
Web Administration Tool loses connections on HP-UX . . . . .	49
Web Administration Tool tabs, table headers, and static list boxes are displayed in incorrect language . . . . .	49
Microsoft Internet Explorer browser problems . . . . .	50
HTML special characters are not displayed correctly . . . . .	51
Web Administration Tool requires IBM JDK on a Domino server . . . . .	51
Web Administration Tool does not save templates created with an object class that has no attributes. . . . .	51
Using CTRL+L to view links makes non-editable fields appear editable . . . . .	51
Internet browser Back and Forward buttons not supported for Web Administration Tool . . . . .	51
Logging on to the Web Administration Tool console on Internet Explorer . . . . .	51
Difficulties encountered using the Web Administration GUI console on the Windows Server 2003 platform . . . . .	52
Users do not have options to enable auditing for LDAP extended operations using Web Administration Tool . . . . .	53
When using Web Administration Tool on HP-UX, the Manage topology panel in Replication management might not get displayed because of locale issue . . . . .	53
A new user might fail to logon to Web Administration Tool for the first time, if the password policy is enabled and "User must change password after reset (pwdMustChange)" in set . . . . .	54
Troubleshooting the embedded version of WebSphere Application Server - Express . . . . .	54
Error when starting the embedded version of WebSphere Application Server - Express on AIX . . . . .	54

## Chapter 9. Troubleshooting replication 55

Replication overview . . . . .	55
Diagnosing replication errors . . . . .	55
Sample replication topology . . . . .	55
Monitoring replication status using idsldapsearch . . . . .	56
Viewing replication errors using the Web Administration Tool . . . . .	58
Viewing replication errors using the idsldapsearch command . . . . .	59
Lost and found log . . . . .	60
Replication Troubleshooting . . . . .	61
Replicated suffix must have ibm-replicationcontext object class . . . . .	61
Verify that suffixes and replication agreements exist using idsldapsearch . . . . .	61

Peer to peer replication returns error "No such object occurred for replica" . . . . .	62
Replication returns error "Insufficient access" . . . . .	62
Replication topology extended operation returns result code 80. . . . .	63
Replication command-line interface error (Windows systems only) . . . . .	63
Entries in LDIF file are not replicated. . . . .	64
Master server can become unstable or stop when serving to large number of replica servers . . . . .	64

## Chapter 10. Troubleshooting performance 67

Identifying performance problem areas . . . . .	67
Audit log . . . . .	67
idsldapd trace. . . . .	67
Adding memory after installation on Solaris systems . . . . .	67
Setting the SLAPD_OCHANDLERS environment variable on Windows . . . . .	68
DB2 rollbacks and isolation levels . . . . .	68
Default value of LOGFILSIZ needs to be increased . . . . .	68
Auditing for performance profiling . . . . .	69

## Chapter 11. Troubleshooting scenarios 73

Server is not responding . . . . .	73
Memory leak suspected . . . . .	73
SSL communications returning errors. . . . .	74
Attribute encryption should be avoided in an environment that includes versions of Tivoli Directory Server earlier than V6.1 . . . . .	74

## Chapter 12. Interoperability . . . . . 77

Interoperability with Novell eDirectory Server. . . . .	77
When performing simple bind using Tivoli Directory Server client utilities against Novell eDirectory Server, error message such as "ldap_bind: Confidentiality required" might get displayed . . . . .	77
Interoperability with Microsoft Active Directory . . . . .	77
Making Tivoli Directory Server configured over SSL using serverClientAuth authentication to work with Microsoft Active Directory client LDP.exe. . . . .	77

## Chapter 13. Known limitations and general troubleshooting . . . . . 81

Known limitations . . . . .	81
Command line utilities allow an option to be entered more than once . . . . .	81
Some types of invalid data entered on command line utilities do not produce an error . . . . .	81
No locking mechanism for conflicting commands on the same directory instance . . . . .	81
Using CTRL+C with a client command that takes passwords can cause an error . . . . .	81
Unable to drop database . . . . .	81
On AIX, Linux, Solaris, and HP-UX operating systems, error messages appear after command prompt when starting the administration daemon . . . . .	82
Partial replication . . . . .	82

Replication is not initiated if the password encryption settings of a supplier are not supported by the consumer . . . . .	83	After enabling language tags, do not disable language tags. . . . .	90
Migrating from IBM Tivoli Directory Server V6.0 to 6.1 . . . . .	83	Create the key database certificate before setting up SSL . . . . .	90
In Tivoli Directory Server V6.1, alias dereferencing might not work when persistent search is run on a server with no alias entries . . . . .	83	idsbulkload appears to hang during parsing phase . . . . .	90
When both proxy and back-end servers are configured to use PKCS#11 mode and need to communicate with a remote nCipher crypto hardware for SSL operation, the operation times out . . . . .	84	Tivoli Directory Server may crash if the size of any log file exceeds the system file size limit . . . . .	91
Tivoli Directory Server V6.1 instance stops when nCipher crypto hardware client is restarted. . . . .	84	Not able to connect to Tivoli Directory Server over SSL while copying an instance using the idsxinst tool . . . . .	91
Querying an entry of large size using the idsldapdiff tool might throw an exception . . . . .	84	Tivoli Directory Server fails to start or displays error when performing ldap operations after bulkload is done. . . . .	91
The idsadsrun utility might fail when synchronizing a large number of entries with size-limit, time-limit like exceptions . . . . .	84	Migration fails if Tivoli Directory Server V6.0 is configured with DB2 v8 and the environment variables are set for a different version of DB2. . . . .	92
The idsadsrun utility fails if a Tivoli Directory Server instance is run on a different port using the -p option. . . . .	85	The idsadsrun tool might fail for some instances when run simultaneously for multiple instances on the same machine . . . . .	92
Operations error is displayed when null based search is performed against a proxy server . . . . .	85	On windows operating system, Tivoli Directory Server startup messages might get displayed in two different locales when a language other than English is specified for Tivoli Directory Server. . . . .	93
When installing using ISMP, a change in disk space on the system does not get refreshed on the tool. . . . .	85	Unable to open a new connection for an LDAP client to connect to Tivoli Directory Server running on a Linux or Solaris operating system . . . . .	93
When the pwdLockout attribute is set to true, user account might get locked even if the number of invalid bind attempts is less than the pwdMaxFailure value . . . . .	85	When deploying a replica or a peer in a replication environment using the idsideploy tool, if the tool detects more than one entry with same replica serverID and ibm-replicationServerIsMaster=true, the tool throws an error . . . . .	93
Description attribute for groups is not syncing from Active Directory to Tivoli Directory Server . . . . .	86	The idsadscfg, idssnmp, and idslogmgmt tools might throw error if the environment variable values contain spaces . . . . .	94
When configuring Tivoli Directory Server V6.1 over SSL to use PKCS#11 SYMMETRIC acceleration support, there are chances for memory leak . . . . .	86	Platform specific problems . . . . .	94
Importing Tivoli Directory Server V6.1 LDIF files to versions earlier than V6.1. . . . .	87	For AIX only . . . . .	94
General troubleshooting . . . . .	87	For Windows 2000, Windows Server 2003 Enterprise and Windows XP client only . . . . .	96
Instance owner unable to access core file for core that occurred during server initialization . . . . .	87	For HP only . . . . .	99
Key label in .kdb file and ibmslapd.conf file do not match. . . . .	87	<b>Appendix A. Common Base Event (CBE) features . . . . .</b>	<b>101</b>
GSKit certificate error . . . . .	87	CBE related scenarios. . . . .	102
Server instance fails to start because of incorrect file permissions . . . . .	88	Log archiving and CBE activity interference . . . . .	102
Server instance fails to start because localhost hostname is set incorrectly . . . . .	88	Log activity overlapping cycles . . . . .	103
Server instance cannot be started except by instance owner . . . . .	88	<b>Appendix B. Support information . . . . .</b>	<b>105</b>
Error opening slapd.cat file on Windows systems . . . . .	88	Searching knowledge bases. . . . .	105
DSML file client produces error. . . . .	89	Search the information center on your local system or network. . . . .	105
Non default log files need valid path. . . . .	89	Search the Internet . . . . .	105
Null searches retrieve entries of deleted suffixes . . . . .	89	Obtaining fixes . . . . .	105
Fixing an "SQL0964C Transaction log for database is full" error . . . . .	89	Contacting IBM Software Support . . . . .	106
idsldapsearch command with -h option gives error with the DIGEST-MD5 mechanism. . . . .	90	Determine the business impact of your problem . . . . .	107
		Describe your problem and gather background information . . . . .	107
		Submit your problem to IBM Software Support . . . . .	107

**Appendix C. Notices . . . . . 109**  
Trademarks . . . . . 110

**Index . . . . . 113**



---

## About this book

IBM® Tivoli® Directory Server is the IBM implementation of Lightweight Directory Access Protocol for supported Windows®, AIX®, Linux® (xSeries®, zSeries®, pSeries®, and iSeries™), Solaris, and Hewlett-Packard UNIX® (HP-UX) operating systems.

*IBM Tivoli Directory Server Version 6.1 Problem Determination Guide* contains information about possible limitations, problems, and corrective actions that can be attempted before contacting IBM Software Support. This guide also includes information about tools you can use for determining problems with IBM Tivoli Directory Server 6.1.

---

## Intended audience for this book

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

---

## Publications

This section lists publications in the IBM Tivoli Directory Server version 6.1 library and related documents. The section also describes how to access Tivoli publications online and how to order Tivoli publications.

### IBM Tivoli Directory Server version 6.1 library

The following documents are available in the IBM Tivoli Directory Server version 6.1 library:

- *IBM Tivoli Directory Server Version 6.1 What's New for This Release*, SC23-6539-00  
Provides information about the new features in the IBM Tivoli Directory Server Version 6.1 release.
- *IBM Tivoli Directory Server Version 6.1 Quick Start Guide*, GI11-8172-00  
Provides help for getting started with IBM Tivoli Directory Server 6.1. Includes a short product description and architecture diagram, as well as a pointer to the product Information Center and installation instructions.
- *IBM Tivoli Directory Server Version 6.1 System Requirements*, SC23-7835-00  
Contains the minimum hardware and software requirements for installing and using IBM Tivoli Directory Server 6.1 and its related software. Also lists the supported versions of corequisite products such as DB2® and GSKit.
- *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide*, GC32-1560-00  
Contains complete information for installing, configuring, and uninstalling IBM Tivoli Directory Server. Includes information about upgrading from a previous version of IBM Tivoli Directory Server.
- *IBM Tivoli Directory Server Version 6.1 Administration Guide*, GC32-1564-00  
Contains instructions for performing administrator tasks through the Web Administration Tool and the command line.
- *IBM Tivoli Directory Server Version 6.1 Command Reference*, SC23-7834-00

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

- *IBM Tivoli Directory Server Version 6.1 Server Plug-ins Reference*, GC32-1565-00  
Contains information about writing server plug-ins.
- *IBM Tivoli Directory Server Version 6.1 Programming Reference*, SC23-7836-00  
Contains information about writing Lightweight Directory Access Protocol (LDAP) client applications in C and Java™.
- *IBM Tivoli Directory Server Version 6.1 Performance Tuning and Capacity Planning Guide*, SC23-7836-00  
Contains information about tuning the directory server for better performance. Describes disk requirements and other hardware needs for directories of different sizes and with various read and write rates. Describes known working scenarios for each of these levels of directory and the disk and memory used; also suggests rough rules of thumb.
- *IBM Tivoli Directory Server Version 6.1 Problem Determination Guide*, GC32-1568-00  
Contains information about possible problems and corrective actions that can be tried before contacting IBM Software Support.
- *IBM Tivoli Directory Server Version 6.1 Messages Guide*, GC32-1567-00  
Contains a list of all informational, warning, and error messages associated with IBM Tivoli Directory Server 6.1.
- *IBM Tivoli Directory Server Version 6.1 White Pages*, SC23-7837-00  
Describes the Directory White Pages application, which is provided with IBM Tivoli Directory Server 6.1. Contains information about installing, configuring, and using the application for both administrators and users.

## Related publications

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

- *Java Naming and Directory Interface™ 1.2.1 Specification* on the Sun Microsystems Web site at <http://java.sun.com/products/jndi/1.2/javadoc/index.html>.  
IBM Tivoli Directory Server Version 6.1 uses the Java Naming and Directory Interface (JNDI) client from Sun Microsystems. See this document for information about the JNDI client.
- The Tivoli Software Library provides a variety of Tivoli publications such as white papers, datasheets, demonstrations, redbooks, and announcement letters. The Tivoli Software Library is available on the Web at: <http://publib.boulder.ibm.com/tividd/td/link/tdprodlist.html>
- The DB2 documentation library is located at <http://www.ibm.com/software/data/db2/library/>.

## Accessing terminology online

The *Tivoli Software Glossary* includes definitions for many of the technical terms related to Tivoli software. The *Tivoli Software Glossary* is available at the following Tivoli software library Web site:

<http://publib.boulder.ibm.com/tividd/glossary/tivoliglossarymst.htm>

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

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

## Accessing publications online

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

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

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

## Ordering publications

You can order many Tivoli publications online at the following Web site:

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

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

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

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

1. Go to <http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>.
2. Select your country from the list and click **Go**.
3. Click **About this site** in the main panel to see an information page that includes the telephone number of your local representative.

You can also see the following Web site for a list of telephone numbers:<http://www.ibm.com/software/tivoli/order-lit/>

---

## 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 additional information, see the Accessibility Appendix in the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide*.

---

## Tivoli technical training

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

---

## Support information

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

- IBM Support Assistant: You can search across a large collection of known problems and workarounds, Technotes, and other information at <http://www.ibm.com/software/support/isa>.
- Obtaining fixes: You can locate the latest fixes that are already available for your product.
- Contacting IBM Software Support: If you still cannot solve your problem, and you need to work with someone from IBM, you can use a variety of ways to contact IBM Software Support.

For more information about resolving problems, see Appendix B, "Support information," on page 105.

---

## Conventions used in this book

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

### Typeface conventions

This book uses the following typeface conventions:

#### **Bold**

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

#### *Italic*

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

#### Monospace

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

## Operating system-dependent variables and paths

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

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

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



---

## Chapter 1. Introduction to problem determination

Problem determination, or troubleshooting, is the process of determining why a product is malfunctioning or not functioning as you expect it to. This chapter introduces problem determination as it relates to IBM Tivoli Directory Server Version 6.1.

---

### IBM Tivoli Directory Server overview

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

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

---

### Built-in troubleshooting features

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

#### Core file generation

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

You must have core file generation enabled in order for core file information to be generated. See "Generating core files" on page 9 for more information about core files and for instructions for enabling core file generation.

#### Support Tool (idssupport)

The Support Tool collects relevant data about the directory server and the system (such as logs, directory listing, schema files, and core files) and packages this information in a compressed file archive that you can send to IBM Software Support for help in diagnosing your problem. See "IBM Tivoli Directory Server Support Tool" on page 11 for more information.

To use the Support Tool, you must have the proxy server and IBM Tivoli Directory Integrator 6.1.1 installed. The *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* contains instructions for installing the proxy server.

#### Error logs

Error logs record error messages that occur during directory server processing. IBM Tivoli Directory Server detects and saves these errors in a text file. See Chapter 2, "Logging utilities," on page 3 for more information.

#### Audit logs

Audit logs record suspicious patterns of activity in order to detect security violations. If security is violated, the audit log can be used to determine

how and when the problem occurred. IBM Tivoli Directory Server detects and saves these errors in a text file. See Chapter 2, “Logging utilities,” on page 3 for more information.

---

## Using the Messages Guide to resolve errors

The *IBM Tivoli Directory Server Version 6.1 Messages Guide* contains a list of messages you might encounter in the IBM Tivoli Directory Server logs, graphical user interfaces, and the command line. Use the unique message ID associated with a message to locate detailed explanations and suggested operator responses in the *IBM Tivoli Directory Server Version 6.1 Messages Guide*.

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

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

You can search for “GLPL2D014E” in the *IBM Tivoli Directory Server Version 6.1 Messages Guide* for information about why the error occurred and how to resolve it.

DB2 error log messages, lost and found log messages, admin audit log messages, and server audit log messages are not contained in the *IBM Tivoli Directory Server Version 6.1 Messages Guide*.

---

## Troubleshooting topics

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

- Installation and uninstallation: See Chapter 4, “Troubleshooting installation and uninstallation,” on page 17 for more information.
- Migration: See Chapter 5, “Troubleshooting migration,” on page 35 for more information.
- Instance Creation: See Chapter 6, “Troubleshooting instance creation and configuration,” on page 37 for more information.
- Configuration: See Chapter 6, “Troubleshooting instance creation and configuration,” on page 37 for more information.
- DB2: See Chapter 7, “Troubleshooting DB2,” on page 45 for more information.
- Web Administration Tool and application server: See Chapter 8, “Troubleshooting the Web Administration Tool and the application server,” on page 47 for more information.
- Replication: See Chapter 9, “Troubleshooting replication,” on page 55 for more information.
- Performance: See Chapter 10, “Troubleshooting performance,” on page 67 for more information.
- Scenarios: See Chapter 11, “Troubleshooting scenarios,” on page 73 for more information.
- General troubleshooting: See Chapter 13, “Known limitations and general troubleshooting,” on page 81 for more information.



---

## Chapter 2. Logging utilities

IBM Tivoli Directory Server Version 6.1 provides several logs that can be viewed either through the Web Administration Tool or the system command line. See the *IBM Tivoli Directory Server Version 6.1 Administration Guide* for information about viewing the logs. See “Using the Messages Guide to resolve errors” on page 2 for information about resolving error messages that you find in the logs.

By default, all the logs listed in this section are in the `directory_server_instance_name/logs` (or `directory_server_instance_name\logs` on Windows) directory. The file names shown are the defaults, but you can change both the paths and the file names for the logs. See the *IBM Tivoli Directory Server Version 6.1 Administration Guide* for information. The IBM Tivoli Directory Server logs are:

### Administration daemon error log (`ibmdiradm.log`)

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

A sample of the log looks like this:

```
Jan 18 09:56:29 2006 GLPCOM003I Non-SSL port initialized to 3538.
Jan 18 09:56:29 2006 GLPADM028I Admin Daemon audit logging is started.
Jan 18 09:56:34 2006 GLPADM004I IBM Tivoli Directory (SSL), Version 6.1
ibmdiradm started
```

### Administration daemon audit log (`adminaudit.log`)

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

Since the Admin daemon 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 Admin daemon must receive the dynamic update configuration request or you must restart the Admin daemon.

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

A sample of the log looks like this:

```
2006-01-15-19:59:17.130-06:00GLPADM028I Admin Daemon audit logging is started.
AuditV3--2006-01-16-22:04:50.93986-06:00--V3 Bind--bindDN: CN=ROOT--client:
127.0.0.1:3665--connectionID: 0--received:
2006-01-16-22:04:50.93986-06:00--Success
AuditV3--2006-01-16-22:04:50.93986-06:00--V3 Search--bindDN: CN=ROOT--client:
127.0.0.1:3665--connectionID: 0--received:
2006-01-16-22:04:50.93986-06:00--Success
AuditV3--2006-01-16-22:04:50.93986-06:00--V3 Unbind--bindDN: CN=ROOT--client:
127.0.0.1:3665--connectionID: 0--received:
```

```

2006-01-16-22:04:50.93986-06:00--Success
AuditV3--2006-01-16-22:08:09.94185-06:00--V3 Bind--bindDN: CN=OT--client:
127.0.0.1:3678--connectionID: 1--received:
2006-01-16-22:08:09.94185-06:00--Invalid credentials
AuditV3--2006-01-16-22:08:09.94185-06:00--V3 Unbind--bindDN: --client:
127.0.0.1:3678--connectionID: 1--received:
2006-01-16-22:08:09.94185-06:00--Success

```

### Audit log (audit.log)

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

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

```

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

```

See the section "Creating the administrative group" in IBM Tivoli Directory Server Version 6.1 Administration Guide to know more about administrative roles and permissions required to access various objects.

The following is a sample of the audit log:

```

2006-01-16-17:38:15.484-06:00--GLPSRV023I Audit logging started. The audit
configuration options are:
  ibm-slapdLog = C:\idsslapd-ldaptest\logs\audit.log,ibm-auditVersion =
  ,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.
2006-01-16-17:38:15.656-06:00--GLPSRV009I IBM Tivoli Directory (SSL),
Version 6.1 Server started.
AuditV3--2006-01-16-17:39:28.468-06:00--V3 anonymous Search--bindDN:
<*CN=NULLDN*>--client:
127.0.0.1:3792--connectionID: 1--received: 2006-01-16-17:39:28.453-06:00--
No such object
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: cn=monitor
scope: wholeSubtree
derefAliases: neverDerefAliases
typesOnly: false
filter: (objectclass=*)

```

### Bulkload error log (bulkload.log)

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

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

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

The following bulkload error log resulted:

```
04/05/06 09:31:19 GLPCTL113I Largest core file size creation limit for
the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/06 09:31:19 GLPCTL114I Largest file size creation limit for
the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/06 09:31:19 GLPCTL115I Maximum data segment limit for
the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/06 09:31:19 GLPCTL116I Maximum physical memory limit for
the process (in bytes): '-1'(Soft limit) and '-1'(Hard limit).
04/05/06 09:31:19 GLPBLK072I Bulkload started.
04/05/06 09:31:19 GLPBLK050I Extracting parent DNs ...
04/05/06 09:31:19 GLPBLK116E Invalid line detected: 3
04/05/06 09:31:19 GLPBLK044I 1 errors detected during parsing phase.
04/05/06 09:31:20 GLPBLK073I Bulkload completed.
```

### Configuration tools log (`idstools.log`)

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

The following is a sample of the configuration tools log:

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

### DB2 error log (`db2cli.log`)

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

The following is a sample of the DB2 error log:

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

```
SQLSTATE=58031
```

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

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

### Lost and found log (lostandfound.log)

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

The information logged for each replaced entry includes:

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

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

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

### Server error log (ibmslapd.log)

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

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

```
Sep 13 14:31:04 2006 GLPL2D014E Suffix entry has not been created for entry
cn=Robert Dean, ou=In Flight Systems, ou=Austin, o=sample.
Sep 13 14:31:04 2006 GLPRDB002W ldif2db: 0 entries have been successfully added
out of 50 attempted.
```

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

### **Installation and uninstallation logs**

In addition, there are logs created during installation and uninstallation. The InstallShield GUI installation and uninstallation logs are: ldapinst.log, ldapuninst.log and ldaplp\_inst.log (for language packs). For more information about these logs, see Chapter 4, “Troubleshooting installation and uninstallation,” on page 17.



---

## Chapter 3. Other diagnostic tools

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

---

### Generating core files

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

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

#### For Windows operating systems (Dr. Watson debugger)

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

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

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

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

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

#### For Linux operating systems

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

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

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

## For AIX operating systems

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

```
ulimit -c unlimited
```

## For Solaris operating systems

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

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

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

```
coreadm -i core
```

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

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

You should see a core file generated.

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

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

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

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

For example:

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

You can disable core file generation using the following command:

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

## For HP-UX operating systems

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

```
ulimit -c unlimited
```



---

## IBM Tivoli Directory Server Support Tool

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

To use the Support Tool, you must have the proxy server and IBM Tivoli Directory Integrator 6.1.1 installed. The *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* contains instructions for installing the proxy server.

### Data collected by the Support Tool

The Support Tool collects the following data:

#### Configuration commands

The following configuration command data is collected:

- The configuration tools log (idstools.log). This log contains information about the following configuration commands:
  - idscfgdb
  - idsucfgdb
  - idscfgchglog
  - idsucfgchglog
  - idscfgsuf
  - idsucfgsuf
  - idscfgsch
  - idsucfgsch
  - idsdnpw
  - idsxcfg
- Bulkload data. The following bulkload data is collected:
  - Bulkload log (bulkload.log), including archived bulkload logs
  - db2load.log
  - Recovery information
- ldif2db and db2ldif data

#### ibmdiradm

The following ibmdiradm information is collected:

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

#### ibmslapd

The following ibmslapd information is collected:

- The ibmslapd version
- Server error log (ibmslapd.log)
- Archived server error logs.
- Audit log (audit.log)
- Archived audit logs

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

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

#### **Administrative commands**

The following instance administration information is collected:

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

#### **Serviceability commands**

The following serviceability information is collected:

- idssupport.log
- idslogmgmt.log

#### **DB2 data**

The following DB2 information is collected:

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

#### **Operating system data**

The following operating system information is collected:

- Platform name
- Kernel mode
- Home directory permission of the DB2 user
- Directory listing of the installation directory

- Directory listing of the directory server instance owner home directory
- Installed package listing of all installed products on UNIX and Linux platforms

#### System hardware data

The following system hardware information is collected:

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

#### GSKit data

The following GSKit information is collected:

- GSKit version

**Embedded WebSphere® Application Server and WebSphere Application Server data** The following Embedded WebSphere Application Server information is collected:

- Embedded WebSphere Application Server version

#### Installation data

The following installation information is collected:

- InstallShield GUI installation and uninstallation logs: ldapinst.log, ldapuninst.log and ldaplp\_inst.log

## Using the Support Tool

To use this tool you must have a server and IBM Tivoli Directory Integrator 6.1 with Fix Pack 1 installed. The ibmslapd.conf file must also be located in its default directory. If these conditions are not met, the Support Tool will not start.

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

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

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

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

#### Command options:

All options are case sensitive and optional.

**-I** <directory\_server\_instance\_name>

Specifies the directory server instance to collect data from. If this name is not specified then data is collected from all directory server instances. If no directory server instances exist then only global data is collected.

**-l** <data\_collection\_path>

Specifies the location to save the collected data. /<timestamp> is appended to this path. The data collection path specified must be a

full path and not a relative path. This value overrides the default value. For the default location, see “Location of Support Tool log and collected data.”

- n core** The latest IBM Tivoli Directory Server core file will not be collected unless this option and value are specified.
- q** Prevents log messages from displaying on the screen. This is an optional parameter.
- v** Displays the version number of the Support Tool.
- ?** Displays the command usage.

## Location of Support Tool log and collected data

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

- On Windows operating systems:  
`<install_home>\var\idssupport\<timestamp>`
- On AIX, Linux, Solaris, and HP-UX operating systems:  
`/var/idsldap/V6.1/idssupport/<timestamp>/`

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

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

- On Windows operating systems:  
`<install_home>\var\idssupport\<timestamp>\idssupport.log`
- On AIX, Linux, Solaris, and HP-UX operating systems:  
`/var/idsldap/V6.1/idssupport/<timestamp>/idssupport.log`

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

---

## Server debug mode

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

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

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

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

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

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

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

*Table 1. Debug categories*

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

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

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

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

```
ldtrc off
```

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

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

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

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

**Note:** The **idsldaptrace** tracing utility can be used to dynamically activate or deactivate tracing of the directory server. See the *IBM Tivoli Directory Server Version 6.1 Command Reference* for information about the **idsldaptrace** utility.

---

## Chapter 4. Troubleshooting installation and uninstallation

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

---

### Product installation overview

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

- Client SDK
- Java client
- Server
  - Proxy server
  - Full server
- Web Administration Tool
- Embedded WebSphere Application Server 6.1.0.7
- IBM DB2 Enterprise Server Edition
- Global Security Kit (GSKit)
- Tivoli Directory Integrator V6.1.1

You can install these components using an InstallShield graphical user interface (GUI) or use operating-system-specific installation methods such as the command line or installation tools for the operating system. InstallShield GUI installation is not available for HP-UX operating systems.

---

### Prerequisite software

If you are installing using the InstallShield GUI, prerequisite software is available for installation as part of the IBM Tivoli Directory Server overall installation process. If you are using the operating system utilities to install, installation might fail if you do not have the prerequisite software installed. Before you install, be sure to read the "System requirements and supported software versions" section in the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide*.

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

- On Windows, in *install\_directory*\var. For example, if you installed in the default location the `ldapinst.log` file is, in `c:\Program Files\IBM\LDAP\V6.1\var`.
- On AIX, Linux, Solaris, and HP-UX systems, the `ldapinst.log` file is in `/var/idsldap/V6.1`.
- On UNIX, the `ldapinst.log` file is in *install\_directory*/var/idsldap/V6.1.

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

If you are installing language packs using the InstallShield GUI, the installation log is in the *install\_directory\LangPack\ldaplp\_inst.log* file on Windows systems or in *install\_directory/LangPack/ldaplp\_inst.log* on AIX, Linux, Solaris, and HP-UX systems.

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

## Failures when installing prerequisite software

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

- If you are installing DB2 and the full server and the DB2 installation fails, an "Installation cannot continue" message is displayed and the installation exits. This is because the full server cannot be installed without DB2. This is considered to be a critical failure and a reason for exiting the installation completely. The reason for the DB2 installation failure can vary, but the result is the same. Examples of reasons for the failure are:
  - If installing from CDs, inserting the wrong CD when prompted for the CD that contains DB2
  - If installing from downloaded and uncompressed .zip or .tar files, uncompressing the file containing DB2 in a different directory which might result in the installation program unable to find the DB2 installation software

In these case the installation exits prematurely, and the installation log is stored in a temporary location and not in the installation path, so you must search for the *ldapinst.log* file on the computer. In order for install to find the image correctly, users must download and untar all the images into the same directory.

- If you are installing other prerequisite software, such as Embedded WebSphere Application Server, and there is a similar failure because the wrong CD was inserted or the software is not present in the directory where the downloaded and uncompressed files reside, the installation can continue for the rest of the components. At the end of installation, you receive an error summary panel describing which components failed to install and the *ldapinst.log* file shows that it could not find components on the installation media.

---

## The idslldap user and group

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

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

You can use the Instance Administration Tool to create users and groups as you are creating a directory server instance, or you can use the following commands to create the user *idslldap* and the group *idslldap* 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/chgrpem -m + root idsldap
```

**On Linux systems:**

Use the following commands.

To create the `idsldap` group:

```
groupadd idsldap
```

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

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

To set the password for user `idsldap`:

```
passwd idsldap
```

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

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

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

**On Solaris systems:**

Use the following commands.

To create the `idsldap` group:

```
groupadd idsldap
```

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

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

To set the password for user `idsldap`:

```
passwd idsldap
```

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

**On HP-UX 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, use the **sam** tool or another appropriate tool.

Be sure that all these requirements are met before you install. The proxy server does not install correctly if the idsldap user exists but does not meet the requirements.

---

## Installing the IBM Tivoli Directory Integrator Server and Administration and Monitoring Console on AIX, Linux, Solaris, and HP-UX systems

If you install the proxy server or the full server, you must install the IBM Tivoli Directory Integrator (IBM Tivoli Directory Integrator) Server and the IBM Tivoli Directory Integrator Administration and Monitoring Console (AMC) if you want to do the following:

- Use the **idssupport** tool, which gathers information from your system that you can supply to IBM Support if you encounter problems.
- Use the log management tool.

You can find information about the support tool and the log management tool in the *IBM Tivoli Directory Server version 6.1 Problem Determination Guide*.

- Use Simple Network Management Protocol (SNMP). For information about SNMP, see the *IBM Tivoli Directory Server version 6.1 Administration Guide*.
- Use Active Directory Synchronization.

On AIX, Linux, Solaris, and HP-UX systems, if you installed using operating system utilities, you must install the IBM Tivoli Directory Integrator Server and the IBM Tivoli Directory Integrator Administration and Monitoring Console manually.

**Note:** Before you install the IBM Tivoli Directory Integrator Server and the IBM Tivoli Directory Integrator Administration and Monitoring Console, you must have the following installed:

- An application server: either WebSphere Application Server or Embedded WebSphere Application Server
- The Java Client provided with IBM Tivoli Directory Server 6.1

If you use the installation instructions in this README, you will have these installed.

The high-level steps are as follows:

1. Download the installable image. See “Installable image subdirectories and files” on page 21.
2. Install the Deployment Engine. See “Installing the Deployment Engine” on page 21.
3. Install the IBM Tivoli Directory Integrator Server installable unit. See “Installing the IBM Tivoli Directory Integrator Server installable unit” on page 22.
4. Install the IBM Tivoli Directory Integrator AMC installable unit. See “Installing the IBM Tivoli Directory Integrator Administration and Monitoring Console installable unit” on page 23.
5. Post -installation steps. See “After you install” on page 24.

If you want to uninstall the IBM Tivoli Directory Integrator installable units, see “Uninstalling the IBM Tivoli Directory Integrator installable units” on page 25.

## Installable image subdirectories and files

The IBM Tivoli Directory Integrator installable image contains the following subdirectories and files. (In this list, the IBM Tivoli Directory Integrator installable image is in the /tmp/tdi/6.1 directory, assuming that /tmp is the directory where you downloaded and uncompressed the corequisites package file.)

- /tmp/tdi/6.1/de\_image  
Contains the installable files for the Deployment Engine.  
This folder contains subdirectories and the si\_inst.bat and si\_inst.shfile files.
- /tmp/tdi/6.1/ITDIServerCC  
Contains the files for the Server installable unit. The files are:
  - ITDIServerCC.jar
  - media.inf
- /tmp/tdi/6.1/TDIAMCCC  
Contains the files for the Administration and Monitoring Console installable unit. The files are:
  - TDIAMCCC.jar
  - media.inf
- /tmp/tdi/6.1/ports  
Contains the ports definition file for Administration and Monitoring Console deployment. The file is:
  - TDIAMCPortDef.props
- /tmp/tdi/6.1/responsexmls  
Contains the response XML files to create and delete the Server and Administration and Monitoring Console installable units. The files are:
  - TDIAMcCreateResp.xml
  - TDIAMcDeleteResp.xml
  - TDIServerCreateResp.xml
  - TDIServerDeleteResp.xml

## Installing the Deployment Engine

Before you can install the IBM Tivoli Directory Integrator installable units, you must install the Deployment Engine, also known as the Solution Install Engine. To install the Deployment Engine:

1. Open a command prompt.
2. Go to the directory where the installable image for the Deployment Engine is located. For example, if you downloaded and uncompressed the corequisites package file in the /tmp directory, the installable image for the Deployment Engine is in the /tmp/tdi/6.1/de\_image directory. Run the following command:

```
si_inst.sh -javaHome <path to java>
```

This installs the Deployment Engine in the /usr/ibm/common/acsi/ directory.

3. Go to the /usr/ibm/common/acsi/bin directory and run the following command to verify that the Deployment Engine installed correctly.  
./listIU.sh

If the Deployment Engine installed correctly, you receive output similar to the following :

```
IU      UUID: DDCE934782398B3E81431666515AC8B5  Name:
DE Extensions Interfaces CLI IU  Version: 1.3
IU      UUID: C37109911C8A11D98E1700061BDE7AEA  Name:
Deployment Engine IU      Version: 1.3
IU      RootIU UUID: 94240D11C8B11D99F2D00061BDE7AEA
Name: Install IU      Version: 1.3
```

## Installing the IBM Tivoli Directory Integrator Server installable unit

To install the IBM Tivoli Directory Integrator Server installable unit:

1. Go to the tdi/6.1/responsexmls subdirectory of the directory where you downloaded and uncompressed the corequisites package file (for example, /tmp/tdi/6.1/responsexmls) and update the TDIServerCreateResp.xml file. A sample of the updated file follows:

```
<create>
<packageURI>/home/TDIImage/ITDIServerCC/ITDIServerCC.jar</packageURI>
<discriminant>TDSV61</discriminant>
<selections>
  <selectedFeature name="TDIServerFeature"/>
</selections>
</create>
<parameterOverrides>
  <parameter name="TDISERVER_TARGETDIR"
    value="/opt/IBM/ldap/V6.1/tdi"/>
  <parameter name="TDISERVER_JRE_TARGETDIR"
    value="/opt/IBM/ldap/V6.1/java/jre"/>
</parameterOverrides>
<create>
```

where:

- *packageURI* is the full path of the IBM Tivoli Directory Integrator Server installable unit jar file
  - *discriminant* is always TDSV61. (There can be multiple instances of the server installable unit. This is used to distinguish the IBM Tivoli Directory Server copy.)
  - TDISERVER\_TARGETDIR is the destination directory for IBM Tivoli Directory Integrator. The recommended path for Linux is "/opt/ibm/ldap/V6.1/tdi". For AIX, Solaris, and HP-UX systems, the recommended path is "/opt/IBM/ldap/V6.1/tdi".
  - TDISERVER\_JRE\_TARGETDIR is the path to Java.
2. Go to the /usr/ibm/common/aci/bin directory and run the following command:

```
./de_processreq -ifile /home/TDIImage/responsexmls/TDIServerCreateResp.xml
```

3. The output of the command is in XML format. The following is a sample of the output:

```
<changeRequestResults id="74d774f6:110c75a2828:-7fff"
  requestStatus="completed" rollbackStatus="none">
<startTime>February 15, 2007 1:41:48 PM PST</startTime>
<endTime>February 15, 2007 1:41:48 PM PST</endTime>
<successCount>1</successCount>
<warningCount>0</warningCount>
<errorCount>0</errorCount>
</changeRequestResults>
```

Check the value in the successCount and errorCount tags. If the installable unit was successfully installed, successCount is 1 and errorCount is 0.

## Installing the IBM Tivoli Directory Integrator Administration and Monitoring Console installable unit

To install the IBM Tivoli Directory Integrator Administration and Monitoring Console installable unit:

1. Go to the tdi/6.1/responsexmls subdirectory of the directory where you downloaded and uncompressed the corequisites package file (for example, /tmp/tdi/6.1/responsexmls) and update the TDIAMcCreateResp.xml file. A sample of the updated file follows:

```
<create>
  <packageURI>/home/TDIImage/TDIAMCCC/TDIAMCCC.jar</packageURI>
  <discriminant>TDSV61</discriminant>
  <selections>
    <selectedFeature name="TDIAMCFeature"/>
  </selections>
</create>
<parameterOverrides>
  <parameter name="TDIAMC_TARGETDIR" value="/opt/ibm/ldap/V6.1/tdi"/>
  <parameter name="TDIAMC_APPSRVDIR"
    value="/opt/ibm/ldap/V6.1/appsrv"/>
  <parameter name="TDIAMC_NOAPPSRV" value="false"/>
  <parameter name="TDIAMC_PROFILE" value="TDSTDIAMCProfile"/>
</parameterOverrides>
```

where:

- packageURI is the full path of the IBM Tivoli Directory Integrator AMC IU jar file.
  - *discriminant* is always TDSV61. (There can be multiple instances of the server installable unit. This is used to distinguish the IBM Tivoli Directory Server copy.)
  - TDISERVER\_TARGETDIR is the destination directory for IBM Tivoli Directory Integrator.
  - TDIAMC\_APPSRVDIR is the path where Embedded WebSphere Application Server or WebSphere Application Server is installed so that the AMC .war file can be deployed into the application server.
  - TDIAMC\_NOAPPSRV can be true or false. If true, the AMC war file will not be deployed and TDIAMC\_APPSRVDIR will be ignored. If false or not specified, the AMC .war file will be deployed into the copy of WAS specified by TDIAMC\_APPSRVDIR.
  - TDIAMC\_PROFILE is the profile name into which the AMC war file will be deployed. If TDIAMC\_NOAPPSRV is false or not specified, TDIAMC\_PROFILE must have a value.
2. Go to the /usr/ibm/common/aci/bin directory and run the following command:  

```
./de_processreq -ifile /home/TDIImage/responsexmls/TDIAMcCreateResp.xml
```
  3. The output of the command is in XML format. The following is a sample of the output:

```
<changeRequestResults id="74d774f6:110c75a2828:-7fff"
  requestStatus="completed" rollbackStatus="none">
<startTime>February 15, 2007 1:41:48 PM PST</startTime>
<endTime>February 15, 2007 1:41:48 PM PST</endTime>
```

```

<successCount>1</successCount>
<warningCount>0</warningCount>
<errorCount>0</errorCount>
</changeRequestResults>

```

Check the value in the successCount and errorCount tags. If the installable unit was successfully installed, successCount is 1 and errorCount is 0.

## After you install

After you install the installable units, you can run one of the following commands from the /usr/ibm/common/acsi/bin directory to see the entry of the installed installable units:

```
./listIU.sh
```

or

```
./de_lsrootiu.sh
```

If you use the ./listIU.sh command, the output will look similar to the following:

```

IU      UUID: DDCE934782398B3E81431666515AC8B5  Name: DE Extensions
Interfaces CLI IU  Version: 1.3
IU      UUID: C37109911C8A11D98E1700061BDE7AEA  Name: Deployment Engine
IU      Version: 1.3
IU      UUID: 613FC586233E53FAF095E0BBEBCDA76C  Name:
TDIServerSIU.613fc586233e53faf095e0bbebcda76c Version: 6.1.1.0
IU      UUID: CABEBAB4E8046C2D9FB427430FFF1D70  Name:
TDIAMCSIU.cabebab4e8046c2d9fb427430fff1d70  Version: 6.1.1.0
IU      RootIU UUID: 3513DA4A324093B08B7C8CCF36B39A7E  Name:
TDIServerRootIU.3513da4a324093b08b7c8ccf36b39a7e Version: 6.1.1.0
IU      RootIU UUID: C803B94B174CA606C13C7A13AD44AE7D  Name:
TDIAMCRootIU.c803b94b174ca606c13c7a13ad44ae7d  Version: 6.1.1.0
IU      RootIU UUID: D94240D11C8B11D99F2D00061BDE7AEA  Name: Install IU
Version: 1.3

```

If you use the ./de\_lsrootiu.sh command, the output will look similar to the following:

```

<iuInstances>

<iuInstance undoable="false" softwareLifecycleStatus="Usable">
<instanceID>InstanceID[IU,TDIServerRootIU.3513da4a324093b08b7c8ccf36b39a7e,
      mrid:MULTI_TARGET,6.1.1.0,TDSV61,
      3513DA4A324093B08B7C8CCF36B39A7E]
</instanceID>
<UUID>3513DA4A324093B08B7C8CCF36B39A7E</UUID>
<identityName>TDIServerRootIU.3513da4a324093b08b7c8ccf36b39a7e
</identityName>
<version>6.1.1.0</version>
<hostingEnv>mrid:MULTI_TARGET</hostingEnv>
<discriminant>TDSV61</discriminant>
</iuInstance>

<iuInstance undoable="false" softwareLifecycleStatus="Usable">
<instanceID>InstanceID[IU,TDIAMCRootIU.c803b94b174ca606c13c7a13ad44ae7d,
      mrid:MULTI_TARGET,6.1.1.0,TDSV61,
      C803B94B174CA606C13C7A13AD44AE7D]
</instanceID>
<UUID>C803B94B174CA606C13C7A13AD44AE7D</UUID>
<identityName>TDIAMCRootIU.c803b94b174ca606c13c7a13ad44ae7d
</identityName>
<version>6.1.1.0</version>
<hostingEnv>mrid:MULTI_TARGET</hostingEnv>
<discriminant>TDSV61</discriminant>
</iuInstance>

```

```

<iuInstance undoable="false" softwareLifeCycleStatus="Usable">
<instanceID>InstanceID[IU,Install IU,
    mrid:http://w3.ibm.com/namespaces/2003/OS_componentTypes:
    Operating_System,1.3,
    C:/Program Files/IBM/Common/acsi,
    D94240D11C8B11D99F2D00061BDE7AEA]</instanceID>
<UUID>D94240D11C8B11D99F2D00061BDE7AEA</UUID>
<identityName>Install IU</identityName>
<version>1.3</version>
<hostingEnv>mrid:http://w3.ibm.com/namespaces/2003/OS_componentTypes:
    Operating_System
</hostingEnv>
<discriminant>C:/Program Files/IBM/Common/acsi</discriminant>
</iuInstance>

</iuInstances>

```

## Uninstalling the IBM Tivoli Directory Integrator installable units

To uninstall the IBM Tivoli Directory Integrator installable units:

1. Go to the tdi/6.1/responsexmls subdirectory of the directory where you downloaded and uncompressed the corequisites package file (for example, /tmp/tdi/6.1/responsexmls) and update the TDIAMcDeleteResp.xml and TDIServerDeleteResp.xml files.

2. Update the value of the following tag in both files:

```
<instanceID>value</instanceID>
```

For *value*, substitute the value that you find in the <instanceID> tag of the output of the de\_lsrootiu.sh command. (See “After you install” on page 24 for the output of this command.)

3. In the /usr/ibm/common/acsi/bin directory, run the following commands to uninstall the installable units:

- To uninstall the IBM Tivoli Directory Integrator Server installable unit, :

```
./de_processreq -ifile /<uncompress_directory>/tdi/6.1/
responsexmls/TDIServerDeleteResp.xml
```

For example:

```
./de_processreq -ifile /tmp/tdi/6.1/responsexmls/TDIServerDeleteResp.xml
```

- To uninstall the IBM Tivoli Directory Integrator AMC installable unit:

```
./de_processreq -ifile /<uncompress_directory>/tdi/6.1/
responsexmls/TDIAMcDeleteResp.xml
```

For example:

```
./de_processreq -ifile /tmp/tdi/6.1/responsexmls/TDIAMcDeleteResp.xml
```

---

## Installation logs

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

### Logs for Embedded WebSphere Application Server

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

#### On Windows platforms

- <install\_home>\var\installApp.log
- <install\_home>\var\installAppErr.log
- <install\_home>\var\configApp.log
- <install\_home>\var\configAppErr.log
- <install\_home>\var\migrateApp.log
- <install\_home>\var\migrateAppErr.log

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

- addWebAdminSrv.log
- addWebAdminSrvErr.log

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

- startWebAdminSrv.log
- startWebAdminSrvErr.log

#### On AIX, Linux, and Solaris platforms

- /var/idsldap/V6.1/installApp.log
- /var/idsldap/V6.1/installAppErr.log
- /var/idsldap/V6.1/configApp.log
- /var/idsldap/V6.1/configAppErr.log
- /var/idsldap/V6.1/migrateApp.log
- /var/idsldap/V6.1/migrateAppErr.log

where *install\_home* is the location where you installed IBM Tivoli Directory Server.

## DB2 logs on Windows

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

#### When installing

- <install\_home>\var\DB2setup.log
- <install\_home>\var\db2inst.log
- <install\_home>\var\db2insterr.log
- <install\_home>\var\db2wi.log

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

#### When uninstalling

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

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



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

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

### When installing using the InstallShield GUI

- /var/idsldap/V6.1/db2inst.log
- /var/idsldap/V6.1/db2insterr.log
- /var/idsldap/V6.1/DB2setup.log

### When uninstalling

- /var/idsldap/V6.1/db2uninst.log
- /var/idsldap/V6.1/db2uninsterr.log

### When installing using the db2\_install utility

- /tmp/db2\_install.rc.99999
- /tmp/db2\_install.log.99999

## Tivoli Directory Integrator logs

Logs used by the InstallShield GUI when installing and uninstalling Tivoli Directory Integrator. On UNIX platform, the logs are in the /var/idsldap/V6.1 directory, and on Windows, the log files are in <install\_home>/var directory.

### When installing Tivoli Directory Integrator

- TdiAmcIU.log
- TdiAmcIUErr.log
- TdiDEInst.log
- TdiDEInstErr.log
- TdiServerIU.log
- TdiServerIUErr.log

### When deploying Tivoli Directory Integrator Administration and Monitoring Console into WebSphere Application Server

- AMCProfileCreate.log
- AMCProfileCreateErr.log

### When adding Embedded WebSphere Application Server Tivoli Directory Integrator Administration and Monitoring Console as a Windows service

- addAMCSrv.log
- addAMCSrvErr.log

### When starting Embedded WebSphere Application Server Tivoli Directory Integrator Administration and Monitoring Console as a Windows service

- startAMCSrv.log
- startAMCSrvErr.log

## idslink log on AIX, Linux, and Solaris operating systems

The **idslink** script should be run manually during InstallShield GUI and operating system utility installation of the client, the proxy server, and full server. The **idslink.log** and **idslink.preview** files are located in the /var/idsldap/V6.1/ directory.

## GSKit logs on Windows operating systems

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

- <install\_home>\var\gsksetup.log
- <install\_home>\var\gskitinst.log
- <install\_home>\var\gskitinsterr.log

## Log files generated for native packages on UNIX operating system

On UNIX platform, two logs are generated for each native package that is installed. These logs give information about the native packages. The log files are created in the /var/idsldap/V6.1 directory. Users can refer to these logs to determine the reason why an install failed. These log files are of importance since ISMP installs the native packages in the background.

The various log files that are created during install are:

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

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

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

On AIX systems, additional logs are generated for SSL packages. The log files are created in the /var/idsldap/V6.1 directory.

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

On UNIX, when native packages are uninstalled log files are created in the /var/idsldap/V6.1/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
- clientuninst64Bit.log, clientuninst64BitErr.log
- clientBase.log, clientBaseErr.log
- clientBaseuninst.log, clientBaseuninstErr.log
- engMsg.log, engMsgErr.log
- entitle.log, entitleErr.log
- Gskit.log, GskitErr.log

- javaClient.log, javaClientErr.log
- javaClientuninst.log, javaClientuninstErr.log
- proxy.log, proxyErr.log
- server.log, serverErr.log
- webAdmin.log, webAdminErr.log
- webAdminuninst.log, webAdminuninstErr.log

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

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

---

## Troubleshooting

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

### InstallShield GUI installation

The following items relate to InstallShield GUI installation.

#### Installation failure due to lack of disk space

One reason for an installation failure is lack of disk space. IBM Tivoli Directory Server attempts to verify that there is enough space and generates messages if the required disk space is not found, but sometimes the InstallShield GUI cannot progress far enough to issue a message. Before installing, make sure you have the required free disk space available that is specified in the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide*. All platforms use temporary space. In addition, AIX, Linux, and Solaris platforms use the /var directory. When installation is first run, the JVM is installed to the installation directory, so be sure that your installation destination directory has enough space.

#### Recovering from a failed installation

The first step to recovering from a failed installation is to run the InstallShield Uninstall GUI to clean up any registry entries that might have been made by the installation process. If you do not run the InstallShield Uninstall GUI, the InstallShield GUI might fail the next time you try to use it to install IBM Tivoli Directory Server. See the following sections for information organized by operating system. See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for information about uninstalling using the InstallShield GUI.

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

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

1. Correct any problems listed in the ldapinst.log file. See “Prerequisite software” on page 17 for more information about the ldapinst.log file.

2. Uninstall IBM Tivoli Directory Server using the InstallShield GUI. See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for more information about uninstalling IBM Tivoli Directory Server.
3. Remove the IBM Tivoli Directory Server installation directory. The default directory is C:\Program Files\IBM\LDAP\V6.1.
4. Use **regedit** to remove the LDAP entry in the registry:  
HKEY\_LOCAL\_MACHINE\SOFTWARE\IBM\IDSLDAP\6.1

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

1. Correct any problems listed in the ldapinst.log file. See “Prerequisite software” on page 17 for more information about the ldapinst.log file.
2. Uninstall IBM Tivoli Directory Server using the InstallShield GUI. See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for more information about uninstalling IBM Tivoli Directory Server.
3. Type the following at a command prompt:  
lslpp -l |grep -i ids1
4. If any packages that were installed by IBM Tivoli Directory Server were left on the system, use **installp** to uninstall them, as follows:  
installp -u *packagename*  
See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for information about package names for IBM Tivoli Directory Server.
5. Remove the /opt/IBM/ldap/V6.1 directory.

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

1. Correct any problems listed in the ldapinst.log file. See “Prerequisite software” on page 17 for more information about the ldapinst.log file.
2. Uninstall IBM Tivoli Directory Server using the InstallShield GUI. See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for more information about uninstalling IBM Tivoli Directory Server.
3. Type the following at a command prompt:  
rpm -qa | grep -i ids1

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

```
rpm -ev packagename
```

See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for information about package names for IBM Tivoli Directory Server.

4. If an **rpm** command hangs, try running the command with the **noscripts** option:  
rpm -ev --noscripts *packagename*
5. Remove the /opt/ibm/ldap/V6.1 directory.

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

1. Correct any problems listed in the ldapinst.log file. See “Prerequisite software” on page 17 for more information about the ldapinst.log file.
2. Uninstall IBM Tivoli Directory Server using the InstallShield GUI. See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for more information about uninstalling IBM Tivoli Directory Server.

3. Type the following at a command prompt:  

```
pkginfo | grep -i ids1
```
4. If any packages that were installed by IBM Tivoli Directory Server were left on the system, use **pkgrm** to uninstall them:  

```
pkgrm packagenames
```

See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for information about package names for IBM Tivoli Directory Server.

**Note:** If you encounter problems removing these packages, try to remove the directories containing the packages from `/var/sadm/pkg`
5. Remove the `/opt/IBM/ldap/V6.1` directory, and any other directories left from the installation, such as a language directory.

### Missing files after server installation

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

If you experience this situation:

1. Type `id idsldap` at a command prompt.
2. If the results do not show that the `idsldap` user is a member of the `idsldap` group, do one of the following:
  - Modify the `idsldap` user so that it belongs to the `idsldap` group.
  - Delete the `idsldap` user and the `idsldap` group and then do one of the following:
    - Recreate the `idsldap` user and group as described in the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide*.
    - Do not recreate the `idsldap` user and group, but let the proxy server installation recreate them (when you do step 3.)
3. Reinstall the base server feature.
4. Reinstall the full server feature if it was not installed and you want the full server.

The base server package gets installed with the proxy server or with the full server package. If user is using ISMP install, this is a hidden feature and is installed for the user whenever they choose Proxy or Server.

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

### Operating system utility uninstallation after InstallShield GUI installation

The *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* instructs you to use the InstallShield GUI to uninstall the IBM Tivoli Directory Server if the InstallShield GUI was used to install. If, however, you perform an operating system uninstallation after an InstallShield GUI installation, you must clean up any registry entries that might have been made by the installation process. For instructions for cleaning up the registry entries, see “Recovering from a failed installation” on page 29.

## Operating system utility installation

The following items relate to operating system utility installation.

### InstallShield GUI uninstallation

The following items relate to InstallShield GUI uninstallation.

#### Product directories still exist after uninstallation

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

### If default instance creation fails during the Typical installation

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

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

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

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

#### Password must meet complexity requirements

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

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

- Not contain all or part of the user’s account name
- Be at least six characters in length
- Contain characters from three of the following four categories:
  - English uppercase characters (A through Z)
  - English lowercase characters (a through z)
  - Base 10 digits (0 through 9)
  - Non-alphabetic characters (for example, !, \$, #, %)
- Complexity requirements are enforced when passwords are changed or created

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

```
Found echo string in C:\Program Files\IBM\LDAP\V6.1\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.

## Recovering from a failed InstallShield GUI installation on HP-UX systems

To recover from a failed InstallShield GUI installation on HP-UX systems:

1. Correct any problems listed in the ldapinst.log file. See “Prerequisite software” on page 17 for more information about the ldapinst.log file.
2. Uninstall IBM Tivoli Directory Server using the InstallShield GUI. See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for more information about uninstalling IBM Tivoli Directory Server.
3. Type swremove at a command prompt.
4. If any packages that were installed by IBM Tivoli Directory Server were left on the system, select the IBM Tivoli Directory Server packages you want to remove. (See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for information about package names for IBM Tivoli Directory Server.)
5. Click **Actions -> Mark For Remove**.
6. Click **Actions -> Remove**.
7. Click **OK** when analysis is complete.
8. When removal is complete, click **Done**.
9. Click **File -> Exit**.

**Note:** The subdirectories of /opt/IBM/ldap/V6.1/ and /var/idslldap/V6.1/ might not be removed when you uninstall IBM Tivoli Directory Server. You can use the `rm -rf` command to remove these directories.





---

## Chapter 5. Troubleshooting migration

Migration refers to the process of installing IBM Tivoli Directory Server version 6.1

to replace an earlier version while preserving changes that were made to the data, schema definitions, and directory server configuration from the earlier version. The following sections contain troubleshooting information for migration.

---

### Migration log files

Check the following log files for information about migration processes:

**On AIX, Linux, Solaris, and HP-UX systems:**

Errors that occurred during migration are logged in the `/var/idsldap/V6.1/idsadm.log` file.

**On Windows platforms:**

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

---

### Kerberos service name change

Before IBM Directory Server 4.1, the LDAP server uses **LDAP** as its Kerberos service name, (**LDAP/ldaphost.austin.ibm.com**; ldaphost is the hostname of the computer where the LDAP server is located) to communicate with its client and the Kerberos KDC. For Version 4.1 and above, a lower case service name is used (**ldap/ldapname.austin.ibm.com**). Because of this change, a 6.0 and later version servers might not start after migrating from a 3.2.2 server that is configured to use Kerberos. This is because the 6.1 server is looking for **ldap** in the keytab file in which an **LDAP** service name was located and used by the previous 3.2.2 server. To correct this situation you can do either of the following:

- Generate a keytab file by adding a lower case LDAP Kerberos service name and start using the new keytab file to communicate.
- Set the environment variable `LDAP_KRB_SERVICE_NAME` to **LDAP** before starting the server. This environment variable causes the LDAP server to continue using the upper case LDAP server service name in the keytab file and to communicate with its clients. In the latter case, the environment variable needs to be set on the client side as well to continue using the upper case LDAP service name to communicate with its server.

---

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

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

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

- Comment out the database information from the configuration file and migrate using the Instance Administration Tool.
- Use the **idsimigr** command-line utility for migration. When you use this command-line tool, if the database instance from the `ibm-slapdDbInstance` attribute is no longer on the system, the information in the configuration file is ignored and information for a new database instance is inserted instead. If it was the database that could not be found, the information is removed from the configuration file. You must then run **idscfgdb** to configure a database.

---

## Format of backed-up schema files incorrect

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

---

## ibm-slapdPlugin entry in configuration file changed

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

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

and the line was changed to:

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

The line in the second example is not removed by the migration tool and the server will not be able to load `/usr/ldap/lib/libback-rdbm.so` at startup because the path is not a valid path for 6.1.

---

## If migration fails using ISMP

If ISMP is unable to uninstall earlier versions of Tivoli Directory Server, migration might fail. In such cases, users must uninstall the earlier versions of Tivoli Directory Server manually using the uninstall instruction for that particular release. For example, if Tivoli Directory Server V5.1 uninstall fails, users should uninstall Tivoli Directory Server V5.1 manually using the uninstall instructions provided in *IBM Tivoli Directory Server Version 5.1 Installation and Configuration Guide*.

The log files that you can use to debug migration are in the `/var/idsldap/V6.1` directory for Linux, Solaris, HP-UX, and UNIX. For Windows platform, the log files are stored in `<install_location>\var` directory. The following are the associated log files:

- `migbkup.log`
- `migbkupErr.log`
- `pre60idsimigrErr.log`
- `pre60idsimigr.log`

---

## Chapter 6. Troubleshooting instance creation and configuration

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

---

### Instance creation overview and common errors

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

#### Instance creation overview

After you install a server, you must create a directory server instance. You can use either the Instance Administration Tool (**idsxinst**), which has a GUI, or the **idsicrt** command-line utility to create this instance. When you create a directory server instance, a database instance is also created if the full server package is installed on the computer. By default, the directory server instance and the database instance have the same name. The name must match the name of an existing user on the system that meets certain qualifications. See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for information about the necessary qualifications.

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

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

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

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

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

- Edit the TCP/IP settings for an instance
- View all instances on the computer
- View details about a particular instance
- Delete an instance
- Migrate a server from a previous release to an IBM Tivoli Directory Server 6.1 instance

## Common instance creation errors

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

### Cannot create additional instance because of invalid IP address

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

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

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

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

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

The problem might be one of the following:

- The Host IP addresses file does not have the correct entry for the second IP address. For example, on Linux systems, the `/etc/hosts` file must have the second IP entry in the correct format. For example:  
9.48.181.173 mymachine.mylocation.ibm.com mymachine
- The system settings must be such that the system first checks the Host IP addresses file instead of performing a DNS lookup. The setting in the operating system Name service switch file must be changed to perform Host IP resolution lookup before going to the DNS. For example, on Linux systems, the `/etc/host.conf` file must have the line `multi on` to allow Host IP address file lookup first.

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

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

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

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

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

**Note:** Users can provide any number of IP addresses.

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

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

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

---

## Configuration overview and common errors

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

### Overview

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

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

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

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

- **idsdnpw** sets the administrator DN and password
- **idscfgdb** configures the database for a directory server instance
- **idsucfgdb** unconfigures the database
- **idscfgchglg** configures the change log for a directory server instance
- **idsucfgchglg** unconfigures the change log for a directory server instance
- **idscfgsuf** configures a suffix for a directory server instance
- **idsucfgsuf** unconfigures a suffix for a directory server instance
- **idscfgsch** configures a schema file for a directory server instance
- **idsucfgsch** unconfigures a schema file for a directory server instance
- **idsldif2db** or **bulkload** imports LDIF data
- **idsdb2ldif** exports LDIF data
- **idsdbback** backs up the database
- **idsdbrestore** restores the database
- **idsrunstats** optimizes the database

## Common errors

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

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

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

```
Task is already running.
```

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

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

### Failure when configuring an existing database instance and database

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

### Error when starting the Configuration Tool on AIX

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

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

If this error occurs, check the following:

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

### Configuration programs terminate on AIX

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

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

## DB2 does not configure properly

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

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

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

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

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

## Server does not start after making changes to configuration file attributes

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

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

## Transaction log is full

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

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

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

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

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

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

## Problems in Configuration Tool windows

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

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

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

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

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

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

## Debugging configuration

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

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

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

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

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

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

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

- The output on the screen

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

- DB2 log files



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

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

**Creating advanced debug output:** See “Server debug mode” on page 14 for information about using debugging tools that are provided.



---

## Chapter 7. Troubleshooting DB2

This chapter contains information about problems related to DB2.

---

### DB2 license file expired

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

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

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

```
db2start
```

If your license is correct, you see the message:

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

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

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

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

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

- If you have a CD: `<cdrom_mount_point>/db2/db2/license/db2ese_t.lic` (or `<cdrom_drive:>\db2\db2\license\db2ese_lic` for Windows)
- If you downloaded a zip file for installation:  
`directory_where_file_was_unzipped\tdsv6.1\db2\db2\license\db2ese_t.lic`
- If you downloaded a tar file for installation:  
`directory_where_file_was_untarred\tdsv6.1/db2/db2/license/db2ese_t.lic`

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

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

```
db2licm -a license_filename
```

- You have purchased a different DB2 product.

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

```
db2licm -a license_filename
```

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

---

## Database performance is poor

For detailed information about improving performance (including information about buffer pools), see the *IBM Tivoli Directory Server Version 6.1 Performance Tuning and Capacity Planning Guide*.

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

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

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

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

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

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

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

**Note:** For more detailed performance information, see the *IBM Directory Server Version 6.1 Performance Tuning and Capacity Planning Guide*.

---

## Recovering from migration failure in DB2 9.1

The idsdbmigr tool uses DB2 backup and DB2 restore mechanism to recover from migration failure since direct recovery from migration failure is not available with DB2 9.1.

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

See the section “Overview of online backup and restore procedures for Tivoli Directory Server” in *IBM Tivoli Directory Server Version 6.1 Administration Guide* to know more about DB2 backup and restore.

---

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

The IBM Tivoli Directory Server Version 6.1 Web Administration Tool is installed on an application server, such as Embedded WebSphere Application Server, which is included with the IBM Tivoli Directory Server and administered through a console. WebSphere Application Server (WAS) can also be used as the application server. This chapter explains how to troubleshoot the IBM Tivoli Directory Server Web Administration Tool and application server.

---

### Troubleshooting the Web Administration Tool

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

#### Corruption of data entered in the Web Administration Tool

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

##### On the embedded version of WebSphere Application Server - Express

Edit the server.xml file in the following directory:

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

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

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

##### On WebSphere Application Server

On the WebSphere Administrative Console tree:

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

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

## Migrating files when patching or migrating the Web Administration Tool

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

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

## Additional login panels fail

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

### For AIX, Linux, Solaris, and HP-UX systems:

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

```
mozilla &
```

### For Windows systems:

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

**Note:** Netscape browsers are no longer supported.

## idsldapmodify command puts Web Administration Tool into inconsistent state

If you are logged into the Web Administration Tool and you change your password using the command line (`idsldapmodify` command), the Web Administration Tool changes the server status to stopped. This occurs because the Web Administration Tool opens new connections to the server every time it launches a task. The Web Administration Tool tries to connect to the server with

the old password because it is unaware that the password has been changed; consequently the connection fails. You must log out and log back in using the new password.

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

## Web Administration Tool loses connections on HP-UX

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

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

Table 2. HP-UX operating system kernel configuration parameters

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

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

Use this procedure to set the kernel configuration parameters:

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

See the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for additional parameter settings.

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

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

The environment variables **LC\_ALL** and **LANG** must be set to a native locale supported by Java; for example en\_US.iso88591. They must not be set to either POSIX or C.

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

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

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

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

On the computer to which you want to export the **DISPLAY**, issue the command:  
export DISPLAY=<valid\_computer\_name>:0

On the <valid\_computer\_name> issue the command:  
xhost +

## Microsoft Internet Explorer browser problems

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

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



## HTML special characters are not displayed correctly

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

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

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

## Web Administration Tool requires IBM JDK on a Domino server

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

The following are limitations on the Domino server:

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

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

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

You can create object classes for IBM Tivoli Directory Server 6.1 that have no MAY or MUST attributes. Such object classes can be used to create entries using other auxiliary object classes. However, if you attempt to create a template through the Web Administration Tool using such an object class, you are unable to save the template.

**Note:** All of the object classes included with IBM Tivoli Directory Server 6.1 contain MAY and MUST attributes. They can be used to create templates.

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

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

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

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

## Logging on to the Web Administration Tool console on Internet Explorer

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

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

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

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

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

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

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

- The Web Administration Tool is installed locally.
- The Web Administration Tool runs on a locally installed version of Microsoft Internet Explorer.
- The Web Administration Tool uses the locally installed Embedded WebSphere Application Server V6.1.0.7.
- An IP address or hostname is part of the URL used to access the Web Administration Tool.

To avoid these errors:

1. If Embedded WebSphere Application Server V6.1.0.7 is running locally, add **http://localhost** to the list of trusted sites.
2. If Embedded WebSphere Application Server V6.1.0.7 is running on a remote machine, add the IP address or hostname of the computer on which the Web application server is running to the list of trusted sites. **http://<IP address>** or **http://<hostname>**

To add a Web address to the Trusted Site list:

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

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

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

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

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

## Users do not have options to enable auditing for LDAP extended operations using Web Administration Tool

You can use command line tools to enable auditing for LDAP extended operations. Issue the following command at the command prompt:

```
idsldapmodify -D <adminDN> -w <password> -i audit.ldif
```

where, audit.ldif has the following data:

```
dn: cn=Audit, cn=Log Management, cn=Configuration
ibm-audit: true
ibm-auditExtOp: true
```

If the auditing for LDAP extended operations is not enabled, the Server log shows:

```
GLPSRV098I Directory server audit logging started.
```

The Audit log should have an entry like:

```
2007-06-20-11:17:50.711+00:00DST--GLPSRV023I Directory server audit logging started.
The audit configuration options are:
ibm-auditVersion = c:\idsslapd-idsinst\logs\audit.log,
ibm-audit = 3,ibm-audit = true,ibm-auditFailedOPonly = false,
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,
ibm-auditPerformance = false,ibm-auditPTABindInfo = true.
AuditV3--2007-06-20-11:17:50.731+00:00DST--V3 Unbind--bindDN:
cn=admin--client: 127.0.0.1:1336--connectionID: 6--received:
2007-06-20-11:17:50.731+00:00DST--Success
```

To disable auditing for LDAP extended operations, change the value of `ibm-auditExtOp` to `false` in the input LDIF file and run the `idsldapmodify` command.

## When using Web Administration Tool on HP-UX, the Manage topology panel in Replication management might not get displayed because of locale issue

To resolve the problem, user must set the locale of the system to a desired language. User can list all the available locales by issuing the following command:

```
# locale -a
```

Next, user must set the locale from one of the available locales by specifying language, country, and encoding. The encoding must be `utf8`.

The following command will set the locale to US English with `utf8` encoding.

```
# export LC_ALL=en_US.utf8
```

The eWAS server should to be restarted after this, if it is already running issue the following commands.

```
# /opt/IBM/ldap/V6.1/appsrv/bin/stopServer.sh server1
# /opt/IBM/ldap/V6.1/appsrv/bin/startServer.sh server1
```

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

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

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

---

## **Troubleshooting the embedded version of WebSphere Application Server - Express**

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

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

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

To change the port number for the embedded version of IBM WebSphere Application Server - Express from 9090 to 9091, which is the port used on AIX computers, edit the `WAS_inst_path/config/cells/DefaultNode/virtualhosts.xml` file and change 9090 to 9091. Do the same thing in the `WAS_inst_path/config/cells/DefaultNode/nodes/DefaultNode/servers/server1/server.xml` file. `WAS_inst_path` is the path where the embedded version of IBM WebSphere Application Server - Express is installed.

**Note:** This path does have two subdirectories named `DefaultNode`.

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

---

## Chapter 9. Troubleshooting replication

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

---

### Replication overview

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

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

See the replication chapter in the *IBM Tivoli Directory Server Version 6.1 Administration Guide* for a more detailed overview of replication.

---

### Diagnosing replication errors

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

#### Sample replication topology

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

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

```
version: 1

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

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

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

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

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

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

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

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

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

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

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

```
dn: cn=Configuration  
  
ibm-slapdServerId: masterhost-389
```

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

```
dn: cn=Configuration  
  
ibm-slapdServerId: replicahost-389
```

and the following entry

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

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

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

## Monitoring replication status using `idsldapsearch`

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

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

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

The remainder of the status-related operational attributes are all associated with a replication agreement object. These attributes are only returned when explicitly requested on the search; for example, the following `idsldapsearch` example requests replication agreement status information indicating the replication state for all the replication agreements:

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

The available attributes are:

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

This attribute is deprecated in version 6.0.

- **ibm-replicationState:** The current state of replication with this consumer. Possible values are:

**Ready**

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

**Retry**

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

**Waiting**

Waiting for next scheduled replication time.

**Binding**

In the process of binding to the consumer.

**Connecting**

In the process of connecting to the consumer.

**OnHold**

This replication agreement has been suspended or "held".

**Error log full**

More replication errors have occurred than can be logged. The amount of errors that can be logged is based on the configured value for `ibm-slapdReplMaxErrors`.

- **ibm-replicationLastResult** The results of the last attempted update to this consumer, in the form:

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

This attribute is available only if the replication method is single threaded.

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

<change id> <operation> <entry DN>

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

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

## Viewing replication errors using the Web Administration Tool

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

To view replication errors:

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

From the "View errors" panel you can:

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

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

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

### Supplier

The host name or IP address of the supplier

### Consumer

The host name or IP address of the consumer

### Change ID

The unique ID of the failed update sent to the consumer

### Update DN

The DN of the entry on which the update was attempted

### Operation type

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



### Details

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

### Controls

The controls used during the update

## Viewing replication errors using the `idsldapsearch` command

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

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

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

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

This command can return output similar to the following:

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

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

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

This command can return output similar to the following:

```
dn: entry-85,o=sample
cn: entry-85
objectclass: person
objectclass: eperson
objectclass: organizationalperson
objectclass: inetorgperson
objectclass: top
userpassword: {AES256}tD09yQT540xpp7ZMIg95mA==
sn: user
ibm-entryuuid: bf201fcb-758e-41dc-bdea-1855fe0b860b
control: 1.3.6.1.4.1.42.2.27.8.5.1 false
control: 1.3.18.0.2.10.19 false::
    MIQAAADJMIQAAAAnCgEAMIQAAAAeBAXjcmVhdG9yc05hbWUxhAAAAAECENOPUFETU10MIQAAA
AxCgEAMIQAAAAoBA9jcmVhdGVUaW1lc3RhbXAxhAAAAABEEDzIwMDUwMzMwMjMyNzQ3wjcEAAAAKA
oBADCEAAAAHwQNbW9kaWZpZXJzTmFtZTGTGEAAAACgQIQ049QURNSU4whAAAAEKAQAwhAAAAACgED2
1vZG1meVRpbWVzdGFtcDGEAAAEEQPMjAwNTAzMzAyMzI3NDda
```

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

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

can return output similar to the following:

Operation completed successfully.

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

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

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

the search can return output similar to the following:

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

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

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

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

This command can return output similar to the following:

Operation completed successfully.

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

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

the search can return output similar to the following:

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

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

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

## Lost and found log

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

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

---

## Replication Troubleshooting

The following sections contain troubleshooting information about replication

### Replicated suffix must have ibm-replicationcontext object class

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

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

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

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

where the `mod.ldif` file contains:

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

### Verify that suffixes and replication agreements exist using `idsldapsearch`

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

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

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

where "o=sample" is the replication context.

If this command does not return any results, the suffix is not configured to be replicated. You must configure the suffix to be replicated. See the *IBM Tivoli Directory Server Version 6.1 Administration Guide* for instructions for configuring a suffix for replication.

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

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

where `replctx` is the location where the replication agreements for a replication context are stored; for example, `o=sample`. If the command does not return results, the replication agreement might not exist. In order to replicate correctly, the correct replication agreements must exist. See the *IBM Tivoli Directory Server Version 6.1 Administration Guide* for instructions for adding replication agreements.

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

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

```
09/07/04 12:57:10 Error No such object occurred for replica '<CN=SERVER2>,<CN=SERVER3>,<CN=MISSING_ENTRY>' change ID 5109011.
```

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

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

To correct this error, do the following:

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

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

where:

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

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

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

## Replication returns error "Insufficient access"

When a replication topology extended operation is issued to a Release 6.0 and later servers and the server's consumer is a Release 5.1 or 5.2 (with a Fix Pack lower than 3) server, the operation fails. In the server trace, Insufficient Access can be identified as the cause of the failure.

In this release, when a replication topology extended operation is issued to a server, the server propagates all of its replication topology entries to its consumers. However, the consumers must be either Release 5.2 with at least Fix Pack 3 or Release 6.0 and later consumers. Consumers from either of the following releases are not supported for this extended operation:

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

For a consumer of either of these levels to have exactly the same replication topology entries as its supplier, import and export tools, such as **idsdb2ldif** and **idsldif2db**, can be used.

## Replication topology extended operation returns result code 80

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

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

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

If this occurs, check for the following:

- If the replication context entry exists on the consumer server, be sure that the replication context entry has an objectclass of `ibm-replicationContext`. Alternatively, you can delete the replication context entry so that the supplier can propagate all of its replication topology-related entries, including the replication context entry, to the consumer.
- After sending all the replication topology-related entries under a replication context to the consumer, the supplier of the extended operation sends the replication topology extended operation to the consumer in an effort to cascade the operation. If more than one tier of servers is involved in a replication topology, be sure that each supplier has the proper credential object to bind with its consumers.
- One of the consumer servers is down or not reachable at that instance.
- One of the consumer servers (either the first level or further downstream) is a server from a release prior to version 5.1.
- The replication context is a non-suffix entry and the consumer does not have the parent entry of the context.  
For example, suppose that `cn=johndoe,cn=people,o=sample` is the context for the topology you want to replicate. If `o=sample` is the suffix on the consumer and `cn=people,o=sample` does not exist, the operation will fail.
- The `repltopology` extended operation timed out on a heavily loaded consumer. (This results in message `GLPRPL098E`.)
- If the consumer is from release 5.1 or 5.2 and it has no suffix to which the context can "belong", then the `repltopology` extended operation will fail.  
In the previous example, if neither `o=sample` nor `cn=people,o=sample` are suffixes, the `repltopology` extended operation actually creates a suffix, `cn=darshan,cn=people,o=sample`. This is true **only** for release 6.0 and later consumers. Release 5.1 and 5.2 consumers do not have this capability and so the extended operation fails.
- Suppose that a certain set of agreements already exists on the consumer. The `repltopology` extended operation attempts to delete these agreements and before that attempts to purge the queue associated with that agreement. If the purge fails, the extended operation fails. (This results in message `GLPRPL093E`.)

## Replication command-line interface error (Windows systems only)

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

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

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

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

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

## Entries in LDIF file are not replicated

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

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

Ways in which problems can occur include the following:

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

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

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

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

And then configure the system ulimit values to:

```
core file size      (blocks, -c)    unlimited
data seg size      (kbytes, -d)    unlimited
file size          (blocks, -f)    unlimited
```

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

Restart the servers for the changes to take effect.





---

## Chapter 10. Troubleshooting performance

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

---

### Identifying performance problem areas

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

#### Audit log

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

#### idsslapd trace

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

1. `ldtrc on`
2. `idsslapd -h 4096`

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

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

```
ldtrc off
```

---

### Adding memory after installation on Solaris systems

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

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

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

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

You must restart the computer for the new settings to take effect.

2. From the command line, set the `ulimit` values for increasing process memory and file size to unlimited:

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

---

## Setting the SLAPD\_OCHANDLERS environment variable on Windows

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

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

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

If you see these errors, do the following:

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

---

## DB2 rollbacks and isolation levels

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

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

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

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

where *isolation\_level* is the isolation level you want to apply, such as `DB2_RR_TO_RS`.

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

---

## Default value of LOGFILSIZ needs to be increased

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

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

### For AIX, Linux, Solaris, and HP-UX platforms:

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

To increase this value, issue the following command:

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

### For Windows platforms:

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

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

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

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

---

## Auditing for performance profiling

In order to identify performance bottlenecks during operation execution, you can check the audit log for the summary figures indicating performance hotspots. The following hotspots are identified for auditing:

- When an operation has to wait in the worker thread queue for a long time before the worker thread actually starts executing the operation.
- The time spent for cache contention inside the backend needs to be tracked.
- The time spent in handling client I/O, that is, the time spent in receiving the request and returning the result. This value can also be used for detecting bottlenecks because of slow clients or network issues.

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

While auditing the performance profiling, the following points should be considered:

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

For each operation, performance data field in the audit records is controlled using the configuration option “`ibm- auditPerformance`”. Currently, the following performance data fields will be defined for each operation:

### **operationResponseTime**

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

### **timeOnWorkQ**

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

### **rdbmLockWaitTime**

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

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

- Resource cache
- DN cache
- Entry cache
- Filter cache
- Attribute cache
- Deadlock detector
- RDBM locks

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

### **clientIOTime**

This field represents time in milliseconds that was spent in receiving the complete operation request and returning the complete operation response. This field is implemented in the operation structure and is updated on receiving the complete BER for operation request and on successfully returning the response BER message for the operation.

An example of the audit version 3 format for search operation with `ibm-auditPerformance` enabled will look like:

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

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



---

## Chapter 11. Troubleshooting scenarios

This chapter contains some troubleshooting scenarios you might encounter and provides some solutions.

---

### Server is not responding

If the server appears to not respond, first verify whether the server is truly not responding, or simply performing very slowly.

To determine if the server is suffering from poor performance, follow the directions in the *IBM Tivoli Directory Server Version 6.1 Performance Tuning and Capacity Planning Guide* for monitoring performance. Compare the operations initiated and operations completed values, as well as the adds requested and adds completed values for a better understanding of what is happening on your system in regard to performance.

If you determine that the server is not responding, run the **idssupport** tool. This tool gathers information that you can provide to IBM Software Support to help identify the problem. See “IBM Tivoli Directory Server Support Tool” on page 11 for information about the **idssupport** tool.

---

### Memory leak suspected

If you suspect that you are experiencing a memory leak, run a script similar to the following one. This script gathers information about the memory sizes of the processes running on your system.

**Note:** This is an example for AIX. You might need to make modifications for your operating system.

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

The script is as follows:

```
#!/bin/sh
instance=ldapdb2
port=389
binpath=/opt/IBM/ldap/V6.1/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 aux command: ' | tee -a /tmp/monitor.out
  ps aux | egrep '(slapd|$instance|PID)' | grep -v grep | tee -a /tmp/monitor.out

  echo 'Memory info via vmstat: ' | tee -a /tmp/monitor.out
  #<VMSTAT command-#">
  vmstat -t 2 5 | tee -a /tmp/monitor.out

  echo 'Port activity via netstat: ' | tee -a /tmp/monitor.out
  netstat -an | grep $port | tee -a /tmp/monitor.out
  date | tee -a /tmp/monitor.out

  echo 'cn=monitor output follows....' | tee -a /tmp/monitor.out
```

```

$binpath/ldapsearch -p $port -s base -b cn=monitor objectclass=* | tee
-a /tmp/monitor.out 2>&1

date | tee -a /tmp/monitor.out

echo 'Sample LDAP query follow: ' | tee -a /tmp/monitor.out

##
date | tee -a /tmp/monitor.out
echo 'Same query but direct to db2: ' | tee -a /tmp/monitor.out
##
date | tee -a /tmp/monitor.out

sleep 600 #10minutes

done

```

---

## SSL communications returning errors

If you are experiencing errors on SSL, run the following command to verify that SSL is set up correctly.

```

ldapsearch -Z -K <keyfile> -P <keyfilepw>
-b suffix objectclass=*

```

Where

- *keyfile* is the name of the SSL database file
- *keyfilepw* is the SSL key database password
- *suffix* is the suffix being searched; for example, -b o=sample

Record and send any errors to IBM Software Support.

---

## Attribute encryption should be avoided in an environment that includes versions of Tivoli Directory Server earlier than V6.1

Attribute encryption should not be used in a Tivoli Directory server environment that include server versions earlier than v6.1. This is because storing encrypted attributes on some servers and not storing encrypted attributes on other servers defeats the purpose of encrypting attributes. In such an environment, for interoperability between servers you should not encrypt attributes. If attribute encryption is used in such an environment the following situations might arise:

- Attempts to replicate schema definitions for encrypted attributes might fail because the target server will not recognize the new `IBMAttributeTypes` keywords.
- On a server earlier than v6.1 and servers that do not have matching encryption keys, for an attribute that is defined but not encrypted, data are decoded for replication and are stored in decoded format. If the servers have the matching keys, data are not decoded during replication rendering data useless on the servers that do not have matching keys.

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



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



---

## Chapter 12. Interoperability

This chapter contains information on interoperability between Tivoli Directory Server and other directory servers.

---

### Interoperability with Novell eDirectory Server

#### **When performing simple bind using Tivoli Directory Server client utilities against Novell eDirectory Server, error message such as “ldap\_bind: Confidentiality required” might get displayed**

If you get error message such as “ldap\_bind: Confidentiality required” when performing simple bind using Tivoli Directory Server client utilities against Novell eDirectory Server, you must run the following command:

```
#ldapconfig set "Require TLS for Simple Binds with Password=no"
```

---

### Interoperability with Microsoft Active Directory

#### **Making Tivoli Directory Server configured over SSL using serverClientAuth authentication to work with Microsoft Active Directory client LDP.exe**

To make Tivoli Directory Server configured over SSL using serverClientAuth authentication to work with Microsoft Active Directory client LDP.exe perform the following steps.

1. Select Internet Information Services (IIS) Manager from Administrative Tools in Control Panel.
2. On the left navigation panel, select the Web Site node.
3. Under the Web Site node, right-click Default Web Site, and then select Properties.
4. On the Default Web Site Properties dialog box, select the Directory Security tab.
5. To request for a new certificate, click the Server Certificate button under the Secure communications area. This opens Web Server Certificate Wizard.
  - a. On the Server Certificate page in the IIS Certificate Wizard dialog box, select the Create a new certificate option and click Next.
  - b. On the Delayed or Immediate Request page, enter the required options and click Next.
  - c. On the Name and Security Settings page, in the Name field enter the host name of the machine and click Next.
  - d. On the Organization Information page, specify appropriate names and click Next.
  - e. On the Your Site's Common Name page, in the Common name field, enter the host name of the machine and click Next.

- f. On the Geographical Information page, specify appropriate values and click Next.
  - g. On the Certificate Request File Name page, in the File name field specify the path name and file name for the certificate request and click Next.
  - h. The summary of the values provide is displayed. Click next.
  - i. Click Finish.
6. Send the certificate request using the steps mentioned above to any Certificate Authority (CA) to issue a certificate.
  7. After receiving the server certificate add the certificate using IIS Certificate Wizard.
    - a. On the Pending Certificate Request page, select the Process the pending request and install the certificate option and click Next.
    - b. On the Process a Pending Request page, in the Path and file name field specify the path name and file name of the certificate. You can also use Browse to select the certificate. Click Next.
  8. Export the personal certificate to pfx or p12 format using IIS Certificate Wizard.
    - a. On the Modify the Current Certificate Assignment page, select the Export the current certificate to a .pfx file option and click Next.
    - b. On the Export Certificate page, in the Path and file name field enter the path name and file name where pfx certificate to be stored. Click Next.
    - c. On the Certificate Password page, in the Password and Confirm password fields enter the password and click Next.
    - d. On the Export Certificate Summary page, the summary of the provided values are displayed. Click Next.
    - e. Click Finish.
  9. To import the certificate, double-click the stored pfx certificate. This opens Certificate Import Wizard.
    - a. On the File to Import page, in the File name field enter the path and file name of the pfx certificate and click Next.
    - b. On the Password page, enter the password and click Next.
    - c. On the Certificate Store page, select the Place all certificate in the following store option and the click Browse and select Personal from the Select the certificate store you want to use list in the Select Certificate Store dialog box. Click Next.
    - d. Click Finish.
  10. To export the personal certificate in BER format, perform the following steps.
    - a. Open Internet Explorer, select Internet Options from the Tools menu, select the Content tag in the Internet Options dialog box, and select Certificates under the Certificates area.
    - b. On the Personal tab in the Certificates dialog box, select the certificate and click Export. This opens Certificate Export Wizard.
    - c. On the Export File Format page, select the Base-64 encoded X.509 (.CER) option and click Next.
    - d. On the File to Export page, in the File name field enter the file name you want to export and click Next.
    - e. Click Finish .
  11. On a machine on which a Tivoli Directory Server instance is running, open the Tivoli Directory Server's key database file using GSKit's key management application, gsk7ikm.

12. Add the exported certificate as a signer in the server key database.



---

## Chapter 13. Known limitations and general troubleshooting

This chapter contains miscellaneous problem determination information.

---

### Known limitations

The following sections describe known limitations in IBM Tivoli Directory Server 6.1.

#### Command line utilities allow an option to be entered more than once

You can run a command that specifies an option more than once. If an option is specified more than once, the option entered last is used. For example, if you enter the following command, the `-I inst1` option is ignored and the `-I inst2` option is used.

```
idsdnpw -p root -n -I inst1 -I inst2
```

#### Some types of invalid data entered on command line utilities do not produce an error

If you enter a command that contains invalid data after all required options have been specified, you will not receive an error message. For example, the following command contains the required options for the `idsdnpw` command, but the `'--'` characters following the required option are invalid.

```
idsdnpw -p root -n -I inst1 --
```

Even though the `'--'` characters are invalid, no error is returned.

#### No locking mechanism for conflicting commands on the same directory instance

No locking mechanism exists at this time to prevent conflicting commands from running at the same time for the same directory instance. For example, you can run a command to configure a database and drop the database at the same time.

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

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

#### Unable to drop database

On Windows systems, if all of the following are true, you might not be able to drop the database immediately after you stop a directory server instance.

- The directory server instance is started from the console and not as a service.
- You stop the directory server instance by using the `ibmslapd -k` command.
- You try to drop the database immediately after stopping the directory server instance with the `ibmslapd -k` command.

The Instance Creation Tool and the `idsidrop` and `idsucfgdb` commands are able to unconfigure the database but fail to drop it if all the listed conditions are satisfied. If you encounter this problem, you can manually delete the database directory after running the `idsidrop` or `idsucfgdb` commands. Alternatively, wait at least two minutes after stopping the server, and then drop the database.

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

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

## Partial replication

Partial replication is a replication feature that replicates only the specified entries and a subset of attributes for the specified entries within a subtree. The entries and attributes that are to be replicated are specified by the LDAP administrator. Using partial replication, an administrator can enhance the replication bandwidth depending on the deployment requirements. For instance, an administrator may choose the entries of the object class `person` with `cn`, `sn`, and `userPassword` attributes to be replicated and `description` attribute not to be replicated.

There are situations when administrator's intervention is required for the smooth running of partial replication. These scenarios are listed.

### Creating missing parent entries on the consumer

In filtered replication, an entry addition might fail displaying "No such object" error because the parent entry does not exist on the consumer. This happens because the parent entry did not match the filter and was not replicated. In such cases, if the `ibm-replicationCreateMissingEntries` attribute is set to `TRUE`, the supplier should detect this error case and then generate and submit an add request for the missing entry before retrying the add operation instead of processing this case as an error. The missing entry should have the same DN as that of the immediate parent of the entry whose add failed. The missing entry belong to the `objectclass extensibleObject` and will contain operational attributes for `create` and `modify` timestamps as present on master server, that is, the timestamps will not be modified when the entry is created on consumer. The missing entry should have ACL's as on the supplier server and should also have the `description` attribute value set to "Missing entry created by <master server>".

#### Scenario

Sometimes the method to generate and submit a request to add a missing entry will be recursive and the end condition would be either a successful add of all missing ancestors in the chain or a failure might occur while adding any of the missing ancestors (for any reason other than `NO_SUCH_OBJECT`). In case of a failure, the change cannot be replicated and administrator intervention will be required.

#### Workaround

The administrator should manually take care of handling errors when the `ibm-replicationCreateMissingEntries` attribute is set to `FALSE`. Administrators can also use error logs to identify the replication failure error messages that are logged into error logs.



## Modification in replication filter

### Scenario

In partial replication, changes to replication filter can be dynamic. When a replication filter is changed, the data on the consumer would be in sync with the supplier cannot be assured.

### Workaround

In cases where replication filter is changed, the administrator should take of such changes and reinitialize the consumer as per the new replication filter.

**Note:** The replication filter entry cannot be deleted if it is in use.

## Replication is not initiated if the password encryption settings of a supplier are not supported by the consumer

In a replication environment, if a supplier is using a password encryption setting that is not supported by the consumer, then replication will not be initiated. Also, the supplier logs a message and sets the replication state to “error xxxx” where xxxx is the id of the message that describes the problem.

## Migrating from IBM Tivoli Directory Server V6.0 to 6.1

When migrating from IBM Tivoli Directory Server V6.0 to 6.1, the existing instances can be migrated using the **idsimigr** command line utility. This tool retrieves the schema and configuration files of the instances from the standard location specific to the instances. While migrating from Tivoli Directory Server V6.0 to 6.1, certain checks need to be performed, otherwise, the tool might exit displaying error messages.

- If an instance already exists, the backup directory should not be specified. If the backup directory is specified, the tool will exit displaying appropriate error messages.
- If the Tivoli Directory Server 6.0 instance to be migrated has been dropped before running **idsimigr**, the backup directory should be specified. In such a scenario, the encryption key is not required but if the encryption key is specified, the tool will exit displaying appropriate error messages.
- During migration, the Windows service entry for each directory server and the Directory Administration daemon are migrated to Tivoli Directory Server 6.1. In such scenarios, to avoid any unforeseen errors, it is required that the user take backup of the schema, configuration, and key stash files before migration, even if, the user has not dropped the instances.

## In Tivoli Directory Server V6.1, alias dereferencing might not work when persistent search is run on a server with no alias entries

If persistent searches are run before any alias entries are added to the server, then persistent searches will not dereference aliases. That means, only if alias entries exist on the server before running persistent searches, the dereferenced aliases will be displayed.

## When both proxy and back-end servers are configured to use PKCS#11 mode and need to communicate with a remote nCipher crypto hardware for SSL operation, the operation times out

In order to increase the operation timeout duration, you need to increase the number of retries that a proxy server should attempt to establish a connection. This is because

the total time for which a proxy server waits to establish a connection =  
maximum time for which proxy waits to establish connection \*  
number of retries by a proxy server to establish a connection

To increase the number of retries, export the environment variable, `SERVER_ATTEMPT_TIME`, with the required retry count. Set the retry count to greater than 12, if the crypto hardware used for SSL operation is at a remote location.

## Tivoli Directory Server V6.1 instance stops when nCipher crypto hardware client is restarted

### Scenario

The below mentioned steps describe the situation in which a Tivoli Directory Server instance might stop.

1. Start a Tivoli Directory Server V6.1 instance configured over SSL with server client auth to use PKCS#11 in keystore and accelerator mode.
2. Perform search operation using an LDAP client in SSL mode.
3. Restart the crypto hardware used.
4. Perform search operation using an LDAP client in SSL mode.

### Reason

Users must not restart the crypto hardware if a Tivoli Directory Server V6.1 instance is configured to use PKCS#11 in keystore or accelerator mode. If crypto hardware is reset that is used by a Tivoli Directory Server instance for cryptographic operations, then the instance will stop logging appropriate messages in trace file.

## Querying an entry of large size using the `idsldapdiff` tool might throw an exception

The Java implementation of the `idsldapdiff` tool has limitation because of which it is unable to handle entries on Tivoli Directory Server that have more than 50 MB size. As a result of this, the tool might throw an Out of Memory exception when dealing with entries with more than 50 MB size.

## The `idsadsrun` utility might fail when synchronizing a large number of entries with size-limit, time-limit like exceptions

To avoid exceptions like size-limit, time-limit, you need to consider the following:

1. Before performing synchronization, configure Microsoft Active Directory setting parameters to the following:

<code>MaxPoolThreads</code>	4
<code>MaxDatagramRecv</code>	4096
<code>MaxReceiveBuffer</code>	10485760
<code>InitRecvTimeout</code>	120

MaxConnections	5000
MaxConnIdleTime	900
MaxPageSize	1000000
MaxQueryDuration	1000
MaxTempTableSize	10000
MaxResultSetSize	262144
MaxNotificationPerConn	5
MaxValRange	1500

2. In the `ibmdisrv` file, tune the JVM parameters, for example, for a machine with 1 GB RAM, the parameter values can be `Xms254m-Xmx1024m`. You can tune the parameters based on your machine configurations. For best results, use a machine with high-end configurations to run the Active Directory synchronization tool.
3. Also, synchronizing approximately 100000 entries using the “Run full synchronization of the entries from Active Directory Server to IBM Tivoli Directory Server followed by real time synchronization” mode while running the `idsassrun` tool gives best results. With the “Run real-time synchronization” mode, up to 400000 entries can be synchronized.

## The `idsadsrun` utility fails if a Tivoli Directory Server instance is run on a different port using the `-p` option

Presently, the Active Directory synchronization tool detects the Tivoli Directory Server admin DN, password, LDAP URL, and port number from the instance name. Therefore, when a Tivoli Directory Server instance is run on a different port using the `-p` option, the tool is unable to detect the port number specified using the `-p` option.

## Operations error is displayed when null based search is performed against a proxy server

Proxy server does not support null based search and gives an operations error if null based search is fired against it.

## When installing using ISMP, a change in disk space on the system does not get refreshed on the tool

When using ISMP installation, the tool does not refresh the information when there is a change in disk space on the system. If user modifies disk space allocation on a system, the changed information does not get reflected on the tool. In order to use the changed disk space allocation, user has to cancel the current installation and start a fresh ISMP installation.

## When the `pwdLockout` attribute is set to true, user account might get locked even if the number of invalid bind attempts is less than the `pwdMaxFailure` value

A user account might get locked when all the invalid bind attempts are made within a given time interval that is set in the `pwdFailureCountInterval` attribute. For example, consider the following attributes are set to:

```
ibm-pwdPolicyStartTime=20070217044605Z
pwdInHistory=0
pwdCheckSyntax=1
pwdGraceLoginLimit=0
pwdLockoutDuration=0
pwdMaxFailure=3
```

```

pwdFailureCountInterval=0
passwordMaxRepeatedChars=0
pwdMaxAge=99
pwdMinAge=0
pwdExpireWarning=0
pwdMinLength=5
passwordMinAlphaChars=0
passwordMinOtherChars=0
passwordMinDiffChars=0
ibm-pwdPolicy=true
pwdLockout=false
pwdAllowUserChange=true
pwdMustChange=false
pwdSafeModify=false
ibm-pwdGroupAndIndividualEnabled=true

```

With this setting, if a user makes three invalid bind attempts, the user can still continue with bind attempts because the `pwdLockout` attribute is set to `false`. However, `pwdFailureTime` is registered even when `pwdLockout` is `false` therefore if user has done three invalid bind attempts with `pwdLockout=false`, `pwdFailureTime` will have timestamps of the consecutive authentication failures.

Set the `pwdLockout` attribute to `true`:

```

# 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 will cause lockout of user account. This is because the invalid bind attempts made when “`pwdLockout=false`” is also taken into account depending on the number of values in the `pwdFailureTime` attribute that are younger than `pwdFailureCountInterval`.

## Description attribute for groups is not syncing from Active Directory to Tivoli Directory Server

Active Directory synchronization solution only synchronizes the user entry attributes provided with `TDSOptionalAttributes` in the `adsync_public.prop` file.

## When configuring Tivoli Directory Server V6.1 over SSL to use PKCS#11 SYMMETRIC acceleration support, there are chances for memory leak

On configuring Tivoli Directory Server V6.1 over SSL to use PKCS#11 SYMMETRIC acceleration support for performing cryptographic operations using `nFast` crypto library, memory leak is observed during operations.

**Note:** `nFast` cryptographic library is a third party library which is the cause of memory leak. It is used for PKCS#11 support provided by Tivoli Directory Server V6.1. This library is not shipped along with Tivoli Directory Server V6.1.

## Importing Tivoli Directory Server V6.1 LDIF files to versions earlier than V6.1

Tivoli Directory Server V6.1 LDIF files containing user passwords data that are encrypted using Salted SHA-1 password encryption mechanism cannot be used for LDIF file to earlier versions of Tivoli Directory Server. This is because Tivoli Directory Server versions earlier than 6.1 consider the Salted SHA-1 tag to be an unknown encryption format. In such cases, the server assumes data is in clear text, and encrypts it depending on the user password encryption option value.

---

## General troubleshooting

The following sections describe general problems and solutions in IBM Tivoli Directory Server 6.1.

### Instance owner unable to access core file for core that occurred during server initialization

If the root user starts the server and a core file is produced early during initialization of the server, the core file might not be accessible to the instance owner user. Instead, the root user has access to the core file.

If this error occurs, the root user can manually set the core file's ownership to the instance owner user if desired.

This problem occurs only on AIX, Linux, Solaris, and HP-UX operating systems.

### Key label in .kdb file and ibmslapd.conf file do not match.

If the key label in the SSL key database certificate does not match the key label in the IBM Tivoli Directory Server configuration file (ibmslapd.conf), you will receive the following error:

```
The default SSL key database certificate is incorrect in file
c:/keytabs/pd_ldapkey.kdb.
```

Check the key label in the configuration file and the SSL key database certificate. If they do not match, create a self-signed SSL key database certificate that matches the key label in the configuration file. For more detailed information about how to create a self-signed key database certificate, see the *IBM Tivoli Directory Server Version 6.1 Administration Guide*.

### GSKit certificate error

If you are trying to import a signer or personal certificate from an external certificate authority (CA) such as Entrust and the GSKIT fails with the error, An error occurred while receiving the certificate from the given file.

the problem might be that the certificate returned from Entrust is a chain certificate, not a root certificate. You must have a root certificate to start a certificate chain. A chain certificate cannot start a certificate chain.

If you do not already have a root certificate, the following is one method of obtaining one.

An example of a root certificate is GTE Cybertrust, which is included in Internet Explorer (IE) 5.5; however, it is not included by default in the GSKit kdb database. To obtain this certificate you must:

1. Export one of the GTE Cybertrust certificates (there are 3) from Internet Explorer as Base64 encoded.
2. Add the certificate as a trusted root certificate.

**Note:** In order to use the GSKit option to set a certificate as a trusted root, the certificate must be self-signed.

3. Add the chain CA certificate from Entrust.
4. Receive the SSL certificate from Entrust.

## Server instance fails to start because of incorrect file permissions

On AIX, Linux, HP-UX, and Solaris systems, file permissions are frequently altered inadvertently by the actions of copying or editing a key database file. Because these actions are generally done as the user ID **root**, file permissions are set for the user **root**. For the directory server instance to make use of this file, you must change the file permissions so that it is readable by the user ID **idsldap**. Otherwise the directory server instance fails to start.

```
chown idsldap:idsldap <mykeyring>.*
```

## Server instance fails to start because localhost hostname is set incorrectly

The localhost hostname must correspond to the local loopback address of 127.0.0.1. If localhost is renamed or the TCP/IP address has changed, the directory server instance does not start.

## Server instance cannot be started except by instance owner

On AIX, Linux, HP-UX, and Solaris systems, if a user other than the directory server instance owner cannot start the directory server instance, be sure that the following are true:

- The user who is attempting to start the directory server instance is a member of the primary group of the directory server instance owner.
- The directory server instance owner's primary group has Write access to the location where the database was created.

See "Setting up users and groups: directory server instance owner, database instance owner, and database owner" in the *IBM Tivoli Directory Server Version 6.1 Installation and Configuration Guide* for information about requirements for the directory server instance owner, database instance owner, and database owner.

## Error opening slapd.cat file on Windows systems

On Windows systems, you might receive an error message that includes the following:

```
Error opening slapd.cat
Plugin of type DATABASE is successfully loaded from
  C:/Program Files/IBM/LDAP/V6.1/bin/libback-config.dll.
Error opening rdbm.cat
```

If this occurs, check the NLSPATH environment variable. The installation program sets the NLSPATH environment variable as a system environment variable. However, if the system also has the NLSPATH variable set as a user environment variable, the user NLSPATH environment variable overrides the system setting.

To correct this, append the NLSPATH information from the system environment variable to the information in the user environment variable.

## DSML file client produces error

The DSML file client produces the following error when it is set up to communicate using SSL and a user tries to connect to an LDAP server that does not use SSL.

```
SSL IS ON
javax.naming.CommunicationException: 9.182.21.228:389. Root exception is javax.
net.ssl.SSLProtocolException: end of file
    at com.ibm.jsse.bd.a(Unknown Source)
    at com.ibm.jsse.b.a(Unknown Source)
    at com.ibm.jsse.b.write(Unknown Source)
    at com.sun.jndi.ldap.Connection.<init>(Connection.java:226)
    at com.sun.jndi.ldap.LdapClient.<init>(LdapClient.java:127)
    at com.sun.jndi.ldap.LdapCtx.connect(LdapCtx.java:2398)
    at com.sun.jndi.ldap.LdapCtx.<init>(LdapCtx.java:258)
    at com.sun.jndi.ldap.LdapCtxFactory.getInitialContext(LdapCtxFactory.java:91)
    at javax.naming.spi.NamingManager.getInitialContext(NamingManager.java:674)
    at javax.naming.InitialContext.getDefaultInitCtx(InitialContext.java:255)
    at javax.naming.InitialContext.init(InitialContext.java:231)
    at javax.naming.InitialContext.<init>(InitialContext.java:207)
    at javax.naming.directory.InitialDirContext.<init>(InitialDirContext.java:92)
    at com.ibm.ldap.dsml.DsmlRequest.processRequests(DsmlRequest.java:767)
    at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:253)
    at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:402)
    at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:373)
    at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:296)
    at com.ibm.ldap.dsmlClient.DsmlFileClient.main(DsmlFileClient.java:203)
```

The error is not fatal and the output XML file is generated.

## Non default log files need valid path

If you want to store your log files in a nondefault path, you must ensure that the path exists and is valid. You must create the directory before you can configure the log files.

## Null searches retrieve entries of deleted suffixes

A null search (`ldapsearch -s sub -b "" objectclass=*`) returns all the entries found in the database. If you have deleted a suffix without first removing its entries from the database, those entries are returned by the null search even though the suffix no longer exists.

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

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

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

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

1. Determine the current log file size setting by issuing the command:  
`db2 get db config for ldapdb2 | grep -i logfilsiz`
2. Increase the size of the log file size setting by issuing the command:  
`db2 update db cfg for ldapdb2 using LOGFILSIZ <new value>`
3. Stop the `idsslapd` process.

4. Issue the command:  
`db2 force applications all`
5. Restart the **idsslapd** process.

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

## **idsldapsearch command with -h option gives error with the DIGEST-MD5 mechanism**

The DIGEST-MD5 SASL bind mechanism requires that the client be able to resolve the fully-qualified host name of the server. If the client cannot resolve the server's fully-qualified hostname the bind fails with an LDAP\_PROTOCOL\_ERROR. To correctly resolve the host name, you might need to make system changes or make DNS configuration changes, such as enabling reverse DNS mapping.

For example, AIX, Linux, Solaris, and HP-UX systems have lines in the `/etc/hosts` file with the syntax:

```
<IP address><fully qualified distinguished name><alias>
```

This syntax is used to define the local hostname to the IP address mappings.

If the syntax is something like:

```
127.0.0.1 localhost
```

when localhost is resolved, it is seen as the fully qualified distinguished name of the system. This causes DIGEST-MD5 to fail.

For the DIGEST-MD5 mechanism to work correctly, the syntax must be something like:

```
127.0.0.1 ldap.myserver.mycompany.com localhost
```

The syntax of the line is now such that `ldap.myserver.mycompany.com` is a valid fully qualified distinguished name for the localhost system.

## **After enabling language tags, do not disable language tags**

After enabling the language tag feature, if you associate language tags with the attributes of an entry, the server returns the entry with the language tags. This occurs even if you later disable the language tag feature. Because the behavior of the server might not be what the application is expecting, to avoid potential problems, do not disable the language tag feature after it has been enabled.

## **Create the key database certificate before setting up SSL**

Before setting up SSL communications on your server, you must use the GSKit utility, **gsk7ikm**, to create the necessary certificates. See "Using gsk7ikm" and "Secure Sockets Layer" in the *IBM Tivoli Directory Server Version 6.1 Administration Guide*.

## **idsbulkload appears to hang during parsing phase**

The **idsbulkload** utility has special code to handle nested groups, and the extra processing takes time.



For example, if an LDIF file contains 50,000 nested groups with 100 membergroups in each of the nested groups, **idsbulkload** might need about 1 to 2 seconds to process each one of the nested groups during the parsing phase.

In this case, **idsbulkload** appears to hang before showing any progress.

An environment variable, `BULKLOAD_REPORT_CHUNK`, can be used to increase the frequency of progress reporting.

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

- On AIX, Linux, HP-UX, and Solaris systems: `export BULKLOAD_REPORT_CHUNK=100`
- On Windows systems: `set BULKLOAD_REPORT_CHUNK=100`

**idsbulkload** will then report parsing progress at 100 entry interval. For example:

```
...
GLPBLK061I Parsing entries ...
GPBLK004I 100 entries parsed successfully out of 100 attempts.
LPBLK004I 200 entries parsed successfully out of 200 attempts.
..
```

## **Tivoli Directory Server may crash if the size of any log file exceeds the system file size limit**

When the size of any log file grows beyond the system file size limit, Tivoli Directory Server may crash. This typically occurs when tracing is enabled on the server.

## **Not able to connect to Tivoli Directory Server over SSL while copying an instance using the `idsxinst` tool**

The reason for this problem could be incorrect configuration. To resolve this problem, perform the following steps:

1. Verify that GSKit is installed on the server.
2. Verify that the `gskikm.jar` file is present in the `<tds_ldap_home>/java/jre/lib/ext` directory.
3. In the `java.security` file under the `<tds_ldap_home>/java/jre/lib/security` directory, check if the CMS provider entry exists. If the entry does not exist, add this entry in the `java.security` file by entering the following:

```
security.provider.X=com.ibm.spi.IBMCMSPProvider
```

where, X is the next number in the order.

4. Ensure that `/lib` exists in the system path.
5. While connecting to source server over SSL, providing the 'Key name' is not mandatory and can be left blank.

## **Tivoli Directory Server fails to start or displays error when performing ldap operations after bulkload is done**

After performing bulkload, if Tivoli Directory Server fails to start or displays error when performing LDAP operations, it could be because of one of the following reasons:

- Check the log file, db2diag.log, if there is an error that states "ACCESS TABLE WHEN IN RESTRICTED STATE". This means that loading data or bulkload was not complete or was unsuccessful.
- The table is in the "Load Pending" or "Locked" state. A previous LOAD attempt on the table might have resulted in failure. Accessing the table is not allowed until the LOAD operation is restarted or terminated.

Consider the following options to rectify the problem:

- Stop or restart the failed LOAD operation on the table by issuing LOAD with the TERMINATE or RESTART option.
- Check if the bulkload\_status file is present. This file is created in the home directory of the instance. If this file is present, it means that bulkload was unsuccessful. Check the file for errors and rectify it, and try running the bulk load utility again.

## **Migration fails if Tivoli Directory Server V6.0 is configured with DB2 v8 and the environment variables are set for a different version of DB2**

Migration might fail if either one of these conditions exists:

- If Tivoli Directory Server V6.0 is configured with DB2 v8 and the environment variables are set for a different version of DB2.
- If Tivoli Directory Server V6.0 is configured with DB2 v9 and the environment variables are set for a different version of DB2.

To resolve this, you must ensure that:

- When you migrate Tivoli Directory Server V6.0 configured with DB2 v8 the environment variables set on the system are of DB2 v8.
- When you migrate Tivoli Directory Server V6.0 configured with DB2 v9 the environment variables set on the system are of DB2 v9.

The following environment variables must be updated depending on the DB2 version in use:

- PATH
- CLASSPATH
- INCLUDE
- LIB
- DB2INSTANCE

## **The idsadsrun tool might fail for some instances when run simultaneously for multiple instances on the same machine**

When running the idsadsrun tool simultaneously for multiple instances on the same machine, if the user gets the following exception:

"org.apache.derby.client.am.DisconnectException: java.net.ConnectException : Error opening socket to server <host\_name> on port <port\_number> with message : Connection refused", then user must apply the fix "TDI 6.1.1 LA0002".

To get this fix, go to Tivoli Directory Integrator support site: [http://www-306.ibm.com/software/sysmgmt/products/support/IBMDirectoryIntegrator.html?S\\_CMP=rnav](http://www-306.ibm.com/software/sysmgmt/products/support/IBMDirectoryIntegrator.html?S_CMP=rnav)

## On windows operating system, Tivoli Directory Server startup messages might get displayed in two different locales when a language other than English is specified for Tivoli Directory Server

If Tivoli Directory Server startup messages are being displayed in two different locales, the most likely reason for the problem is that the currently logged in user and the Tivoli Directory Server instance owner have different set of locale configured on the system.

You can consider one of the following ways to rectify this problem:

- Set the LANG environment variable explicitly to the language you want use. For example, set LANG=de\_DE (or any other supported language). You must then start the server from the same window.
- Modify the regional and language settings on the Regional and Language Options dialog box to ensure that both the currently logged in user and the instance owner have the same set of regional and language settings to view the server messages in the same language.

## Unable to open a new connection for an LDAP client to connect to Tivoli Directory Server running on a Linux or Solaris operating system

On Linux and Solaris operating system, there is a limit on the maximum number of file descriptors that can be opened by a process. For Linux and Solaris operating systems, the default value of the maximum number for open file descriptor is 1024 and 256, respectively.

A Tivoli Directory Server V6.1 instance uses 15 file descriptors for the purpose of logging messages. So on Linux, a Tivoli Directory Server V6.1 instance stops accepting new connections after 1009, that is 1024 – 15, concurrent client connects. Whereas on Solaris, a Tivoli Directory Server V6.1 instance stops accepting new connections after 241, that is 256 – 15, concurrent client connects. If an error is encountered while accepting new connections appropriate message is logged. This error does not affect any existing connections only new LDAP clients will fail to connect to the Directory server.

In order to increase the maximum open file descriptors, user should issue the following command and restart the server from the same command prompt.

```
#ulimit -Hn <number of connections>
```

**Note:** The performance with very high number of concurrent client connections depends on the hardware and the operations being performed. With thousands of concurrent client connections sending operations simultaneously the performance of the directory server may decrease.

## When deploying a replica or a peer in a replication environment using the idsideploy tool, if the tool detects more than one entry with same replica serverID and ibm-replicationServerIsMaster=true, the tool throws an error

When deploying a replica or a peer in a replication environment using the idsideploy tool, if the tool detects more than one replication subentry containing

the same serverID value for the attribute `ibm-replicaServerId` with the attribute `ibm-replicationServerIsMaster` set to true, the tool throws error.

For any given replication context, multiple replication subentries are not required, only one replication subentry is required. For example, if the entries are made as shown in the below example, `idsideploy` will fail.

```
dn: ibm-replicaServerId=Peer1,ibm-replicaGroup=default, ou=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=ouunit1, o=sample
objectclass: top
objectclass: ibm-replicaSubentry
ibm-replicaServerId: Peer1
ibm-replicationServerIsMaster: true
cn: Peer1_entry
description: Peer1
```

In the above example, to rectify the problem users should create only one entry.

## The `idsadscfg`, `idssnmp`, and `idslogmgmt` tools might throw error if the environment variable values contain spaces

If you use a copy of IBM Tivoli Directory Integrator that you did not install with IBM Tivoli Directory Server 6.1 and have installed Tivoli Directory Integrator in a different location:

- For the Active Directory synchronization (`idsadscfg`) and SNMP (`idssnmp`) tools, you must explicitly set the `IDS_LDAP_IDI_HOME` environment variable to the directory where you installed Tivoli Directory Integrator.
- For the Log management (`idslogmgmt`) and `idssupport` tools, you must explicitly set the `IDS_LDAP_TDI_HOME` environment variable to the directory where you installed Tivoli Directory Integrator.

For both these variables, `IDS_LDAP_IDI_HOME` and `IDS_LDAP_TDI_HOME`, if the value has space or is in double quotes, then the tools will not work properly. The tools work correctly when tilde, “~” (that is, short path or path with no spaces) is used.

---

## Platform specific problems

This information applies to the following operating systems:

### For AIX only

The following information applies only to the AIX operating system.

#### Problem with `MALLOCTYPE=buckets`

Before setting `MALLOCTYPE` to `buckets` on the AIX 5.2 operating system, ensure that you have installed the patch for APAR IY50668. Otherwise the LDAP server might fail with a core file.

## Verifying that AIX hardware is 64-bit

The server on AIX requires 64-bit hardware. To verify that your AIX hardware is 64-bit, run the following command:

```
bootinfo -y
```

If the command returns 32, your hardware is 32-bit.

In addition, if you type the command `lsattr -El proc0`, the output of the command returns the type of processor for your server. If you have any of the following, you have 64-bit hardware: RS64 I, II, III, IV, POWER3™, POWER3 II, POWER4™, or POWER5™.

## Verifying that the AIX kernel is 64-bit

To verify that you have the 64 bit kernel (`/usr/lib/boot/unix_64`) installed and running, run the following command:

```
bootinfo -K
```

In addition, if you type the command `lsattr -El proc0`, the output of the command returns the type of processor for your server. If you have any of the following, you have 64-bit hardware: RS64 I, II, III, IV, POWER3, POWER3 II, POWER4, or POWER5.

**Note:** If the hardware is 32-bit, then you can only have a 32-bit kernel. You cannot have a 64-bit kernel. If the hardware is 64-bit, then you can have either a 32 or 64-bit kernel.

To switch between a 32-bit and 64-bit mode at the operating system level on AIX 5.1, 5.2, or 5.3:

When you install the operating system, go to Additional features and specify 64-bit mode. (The default is 32-bit mode.) To switch from 32-bit mode to 64-bit mode, use the following commands:

```
# ln -sf /usr/lib/boot/unix_64 /unix
# ln -sf /usr/lib/boot/unix_64 /usr/lib/boot/unix
# bosboot -ad /dev/ipldevice
# shutdown -Fr
# bootinfo -K
```

The kernel is now in 64-bit mode.

To switch from 64-bit mode to 32-bit mode, use the following commands:

```
# ln -sf /usr/lib/boot/unix_mp /unix
# ln -sf /usr/lib/boot/unix_mp /usr/lib/boot/unix
# bosboot -ad /dev/ipldevice
# shutdown -Fr
# bootinfo -K
```

The kernel is now in 32-bit mode.

## Error on AIX when running db2start

The following error might occur when you try to run **db2start**:

```
0509-130 Symbol resolution failed for /usr/lib/threads/libc.a(aio.o)
because:
```

```
0509-136 Symbol kaio_rdwr (number 0) is not exported from
dependent module /unix.
0509-136 Symbol listio (number 1) is not exported from
dependent module /unix.
0509-136 Symbol acancel (number 2) is not exported from
```

```
dependent module /unix.  
0509-136 Symbol iosuspend (number 3) is not exported from  
dependent module /unix.  
0509-136 Symbol aio_nwait (number 4) is not exported from  
dependent module /unix.  
0509-192 Examine .loader section symbols with the  
'dump -Tv' command.
```

If this occurs on AIX, you have asynchronous I/O turned off.

To turn on asynchronous I/O:

1. Run **smitty chgaio** and set **STATE to be configured at system restart** from **defined to available**.
2. Press Enter.
3. Do **one** of the following:
  - Restart your system.
  - Run **smitty aio** and move the cursor to **Configure defined Asynchronous I/O**. Then press Enter.

The **db2start** command now works.

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

The following sections apply only to the Windows 2000, Windows Server 2003 Enterprise and Windows XP client platforms.

### Setting LANG and LC\_ALL system environment variables for nonEnglish InstallShield GUI installation

For the InstallShield GUI installation to bring up the same language that the operating system is using, two variables must be set in the system environment

- LANG = <locale>
- LC\_ALL = <locale>

where <locale> is the locale that the operating system is using.

Go to <http://www.microsoft.com/globaldev/> for a list of Microsoft locale values.

### Certain UTF-8 supplementary characters do not display correctly

IBM Directory Server supports UTF-8 (Unicode Transformation Format, 8-bit form) to use Unicode characters, which contains MS932 (Shift JIS) characters plus supplementary characters not defined in MS932. Supplementary characters might be displayed as square box in Internet Explorer running on Windows 2000. See Figure 1 on page 97.




Figure 1. Unicode Code Point U+9DD7 displayed as a square

If this occurs, install one of the East Asian language kits. Depending on your environment, install the Japanese, Korean, Simplified Chinese or Traditional Chinese language kit which is included in your Windows CDs. For example, Unicode code point U+9DD7 is one of the supplementary characters in the Japanese environment. With the correct language kit installed, the supplementary character is displayed correctly. See Figure 2.



Figure 2. U+9DD7 displayed correctly

**Note:** This problem has not been observed in Windows XP.

### Communications error: Exceeding 64 connections/OCH

On Windows, if you have clients that are generating many connections to the server and the connections are being refused, the server might log error messages similar to the following in the `ibmslapd.log` file:

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

If you see these errors, do the following:

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

```
ibm-slapdSetenv: SLAPD_OCHANDLERS=5
```
4. Restart your server.

If you continue to receive error messages, increase the value of the `SLAPD_OCHANDLERS` environment variable by 5 until you stop receiving error messages.

### Starting IBM Tivoli Directory Server at operating system startup

In IBM Tivoli Directory Server, the server (the `ibmslapd` process) is started manually through the Services window or by the `ibmslapd` command. If you try to start the server automatically by updating the **Startup Type** in the Services window to **Automatic**, errors occur when you restart the computer. This is because DB2 must be running before the `ibmslapd` process can start.

If you want the server to start automatically, you can create a batch file to start the **ibmslapd** process. The batch file should be invoked after all the services are started, so that DB2 will be completely up and running before the **ibmslapd** process starts.

The following is an example of commands in a .bat file that you can add to the Startup folder to start the server:

```
@echo off
%LDAPHome%\bin\ibmdirctl [-h <hostname>] [-D <adminDN>] [-w <password>]
[-p <portnumber>] start -- [ibmslapd options]
```

**Note:** Be sure that the **Startup Type** for the **IBM Tivoli Directory Admin Daemon** entry in the Services window is set to **Automatic**. If it is not, the administration daemon control program (**ibmdirctl**) will not work.

## Backup and restore

### Scenario

On Windows platform, when performing a backup, restore, or load to or from a directory mapped as remote drive using DB2 utilities fails giving error message "SQL2036N The path for the file or device "<file\_or\_devicename>:\ " is not valid".

### Reason

When a user tries to perform a backup, restore, or load to or from a directory mapped as remote drive, for instance, H:\MyFolder\test; user gets the error.

There are two different reasons for getting this error:

1. The user is specifying an invalid shared drive in the command.

```
db2 backup db mydatabase to H:\
```

Error message displayed:

```
SQL2036N The path for the file or device "H:\ " is not valid.
```

2. The user is specifying a valid UNC name for the mapped drive but he is not using the right USERID.

```
db2 backup db mydatabase to \\MyFolder\test
```

Error message displayed:

```
SQL2036N The path for the file or device "\\MyFolder\test "is not valid.
```

For the db2 backup failing the bb2diag.log looks like:

```
database_utilities sqlubcka Probe:0 Database:mydatabase
```

```
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 -- invalidevice 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:mydatabase
```

Backup terminated.

### Workaround

This error is because of Windows restriction. You need to consider the following:

- The user must to start DB2 server by using an existing USERID instead of the default "Local System Account".
- To specify read and write permissions on the network drives, select Services in Administrative Tools under Control Panel. Next, in the Services window, select Properties from the Action menu. Select Log On tab and update "Log on as" to indicate a specific user who has the read and write permissions on the network drives.
- In Windows 2000 and Windows XP, you need to perform backup, restore, or load action by specifying the full qualified UNC name instead of the network share.
- Use the command:

```
db2 backup db mydatabase to \\ MyFolder\test
```

instead of

```
net use H: \\ MyFolder\test
db2 backup db mydatabase to H:
```

## For HP only

### On HP computer, online backup fails when backing up the database to a network mounted directory

On HP computer, in order to use database backup on network mounted directory, perform the following steps:

1. Run the idsideploy command to backup the database to a local directory.

```
idsideploy -L <directory>
```

2. Copy the backed up image to the computer where you want to deploy.
3. Run the idsideploy command to restore the database.

```
idsideploy -p
```



---

## Appendix A. Common Base Event (CBE) features

In an effort to create self-managing environment, IBM has taken initiative in introducing "Autonomic Computing". Autonomic computing is an open standard based architecture that allows systems to configure, heal, optimize, and protect itself. In order to determine the conditions of the different components of the system, it is necessary to standardize the format of the event data so that the system can resolve its current conditions.

To standardize the format of data for the problem determination architecture IBM introduced a common format for log and trace information called the Common Base Event (CBE) format. This format creates consistency across similar fields and improves the availability to correlate across multiple logs. CBE is based on a 3-tuple structured format, which includes:

- Component impacted by a situation, or the source
- Component observing a situation
- Situation data, the properties describing the situation including correlation information

The 3-tuple format makes it possible to write and deploy resource-independent management functions that can isolate a failing component.

In an effort to align IBM Tivoli Directory Server to autonomic computing space, it is a must to have the logs such as error log, audit log, and so on, produced by the Tivoli Directory Server product to provide these logs in CBE format.

The IBM Common Auditing and Reporting Service (CARS) component leverages CBE, which is a common format for events proposed by IBM, and IBM Common Event Infrastructure (CEI) technologies to provide an audit infrastructure. The purpose of CBE is to facilitate effective intercommunication among disparate components within an enterprise. In order to effectively process audit data, the CARS component requires the audit data to be in the CBE format. CEI is an IBM strategic event infrastructure for submission, persistent storage, query, and subscription of the CBE events. The CARS component uses the CEI interfaces for submission of events. These events can be denoted as auditable by using configuration options at the CEI Server that stores them in a CEI XML Event store that meets the auditing requirements.

The CARS component allows staging of data from the CEI XML Event store into report tables. IBM products and customers can provide audit reports based on auditable events staged into report tables. The CARS component also supports managing the lifecycle of auditable events, which includes archive, restore, and audit reports on restored archives.

In Tivoli Directory Server, auditing capability is implemented using the Tivoli Directory Server audit plug-in. A user can implement the audit enhancements to write audited data to CBE format. An example is listed:

- The audit data could be read and transformed to CBE format by an external application such as, IBM Tivoli Directory Integrator, and then sent over to CARS using the CEI API or the CARS embeddable Java client.

To implement the example, the settings for this feature in the `ibmslapd` configuration file are retrieved. If the settings are specified the audit data files are

read periodically and converted into CBE format by the log management tool. Depending on the settings, the CBE formatted data could be written to a file, to a CEI server, or both. The data sent to a CEI server is stored in the CEI database and CARS will move the audit data into a CARS database. The data will then move into a staging area for CARS reports or into a database archive for long term storage.

---

## CBE related scenarios

There are some special case scenarios that should be considered for the CBE feature.

### Attribute related special case scenarios

#### Unspecified attribute settings

If a value is set for `ibm-slapdLogEventFileSizeThreshold` and the value of `ibm-slapdLogEventFileMaxArchives` is not specified either in the default entry or in the specific log entry then in such case archiving will occur but the number of archive files will be unlimited.

#### Attribute settings given in wrong format

- If the value provided for `ibm-slapdLogEventFileSizeThreshold` is in the wrong format then an error message is logged and no archiving will occur.
- If the value provided for `ibm-slapdLogEventFileMaxArchives` is in the wrong format then an error message is logged but archiving will occur and the number of archive files will be unlimited.
- If the value provide for `ibm-slapdLogEventFileArchivePath` is invalid then the archived file path is in the same directory as that of the original file's path.

### CBE file and Log management actions related scenarios

#### Out of disk space

If the disk gets full the log management activity will fail this is indicated by displaying an error message on the standard output and if possible it is also logged in the log.

#### Archive path errors

- If the archived file cannot be written to the path specified than an error message is logged on the `idslogmgmt` log.
- If a file with the same name already exists in the mentioned archive path then an error message is logged in the `idslogmgmt` log and the archiving will fail. When the next log management occurs, for the operation to succeed the timestamp should be different.

---

## Log archiving and CBE activity interference

When the Tivoli Directory Server log management tool is configured to send CBE data to a CEI server and log archiving is also enabled then there is a possibility that the archiving threshold is reached but the log data has not been sent to the CEI server. The reason for this could be that the CEI server is down or the transmission rate to CEI server is lesser than the original log write rate. In such

cases the log archiving is suspended in order to not lose data that data that should be sent to CEI server. The expected behaviors when this situation occurs with the log management settings are listed.

- When CBE formatted logs are enabled and the value of `ibm-slapdLogEventFileMaxArchives` is set to zero, then the CBE file that should have been deleted is kept and the file continues to grow.
- When CBE formatted logs are enabled, the value of `ibm-slapdLogEventFileMaxArchives` is set to a number greater than zero, and the set maximum number of CBE log files have been reached, then the CBE file that should have been deleted is kept and the number of CBE archives continues to grow.
- When CBE formatted logs are disabled and the value of `ibm-slapdLogMaxArchives` is set to zero, then the log file that should have been deleted is kept and the file continues to grow.
- When CBE formatted logs are disabled, the value of `ibm-slapdLogMaxArchives` is set to a number greater than zero, and the set maximum number of log files have been reached, then the oldest archived file that should have been deleted is kept and the number of archives continues to grow.

---

## Log activity overlapping cycles

When a current cycle of log activities are running for a log and if the next cycle of log activities are triggered, the tool should not allow multiple cycles to overlap. The next cycle should only start after the completion of the first cycle. This prevents different log activity cycles from interfering and causing data loss.



---

## Appendix B. Support information

This section describes the following options for obtaining support for IBM products:

- “Searching knowledge bases”
- “Obtaining fixes”
- “Contacting IBM Software Support” on page 106

---

### Searching knowledge bases

If you have a problem with your IBM software, you want it resolved quickly. Begin by searching the available knowledge bases to determine whether the resolution to your problem is already documented.

#### Search the information center on your local system or network

IBM provides extensive documentation that can be installed on your local computer or on an intranet server. You can use the search function of this information center to query conceptual information, instructions for completing tasks, reference information, and support documents.

#### Search the Internet

If you cannot find an answer to your question in the information center, search the Internet for the latest, most complete information that might help you resolve your problem. To search multiple Internet resources for your product, expand the product folder in the navigation frame to the left and select **Web search**. From this topic, you can search a variety of resources including:

- IBM technotes
- IBM downloads
- IBM Redbooks
- IBM developerWorks
- Forums and newsgroups
- Google

---

### Obtaining fixes

A product fix might be available to resolve your problem. You can determine what fixes are available for your IBM software product by checking the product support Web site:

1. Go to the IBM Software Support Web site (<http://www.ibm.com/software/support>).
2. Under **Products A - Z**, select your product name. This opens a product-specific support site.
3. Under **Self help**, follow the link to **All Updates**, where you will find a list of fixes, fix packs, and other service updates for your product. For tips on refining your search, click **Search tips**.
4. Click the name of a fix to read the description and optionally download the fix.

To receive weekly e-mail notifications about fixes and other news about IBM products, follow these steps:

1. From the support page for any IBM product, click **My support** in the upper-right corner of the page.
2. If you have already registered, skip to the next step. If you have not registered, click register in the upper-right corner of the support page to establish your user ID and password.
3. Sign in to **My support**.
4. On the My support page, click **Edit profiles** in the left navigation pane, and scroll to **Select Mail Preferences**. Select a product family and check the appropriate boxes for the type of information you want.
5. Click **Submit**.
6. For e-mail notification for other products, repeat Steps 4 and 5.

For more information about types of fixes, see the *Software Support Handbook* (<http://techsupport.services.ibm.com/guides/handbook.html>).

---

## Contacting IBM Software Support

IBM Software Support provides assistance with product defects.

Before contacting IBM Software Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have:

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus®, and Rational® products, as well as DB2 and WebSphere products that run on Windows or UNIX operating systems), enroll in Passport Advantage® in one of the following ways:
  - **Online:** Go to the Passport Advantage Web page ([http://www.lotus.com/services/passport.nsf/WebDocs/Passport\\_Advantage\\_Home](http://www.lotus.com/services/passport.nsf/WebDocs/Passport_Advantage_Home)) and click **How to Enroll**
  - **By phone:** For the phone number to call in your country, go to the IBM Software Support Web site (<http://techsupport.services.ibm.com/guides/contacts.html>) and click the name of your geographic region.
- For IBM eServer™ software products (including, but not limited to, DB2 and WebSphere products that run in zSeries, pSeries, and iSeries environments), you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for eServer software products, go to the IBM Technical Support Advantage Web page (<http://www.ibm.com/servers/eserver/techsupport.html>).

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States or, from other countries, go to the contacts page of the IBM Software Support Handbook on the Web (<http://techsupport.services.ibm.com/guides/contacts.html>) and click the name of your geographic region for phone numbers of people who provide support for your location.

Follow the steps in this topic to contact IBM Software Support:

1. Determine the business impact of your problem.
2. Describe your problem and gather background information.
3. Submit your problem to IBM Software Support.



## Determine the business impact of your problem

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you need to understand and assess the business impact of the problem you are reporting. Use the following criteria:

Severity 1	<b>Critical</b> business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	<b>Significant</b> business impact: The program is usable but is severely limited.
Severity 3	<b>Some</b> business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	<b>Minimal</b> business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.

## Describe your problem and gather background information

When explaining a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently. To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? (For example, hardware, operating system, networking software, and so on.)
- Are you currently using a workaround for this problem? If so, please be prepared to explain it when you report the problem.

## Submit your problem to IBM Software Support

You can submit your problem in one of two ways:

- **Online:** Go to the "Submit and track problems" page on the IBM Software Support site (<http://www.ibm.com/software/support/probsub.html>). Enter your information into the appropriate problem submission tool.
- **By phone:** For the phone number to call in your country, go to the contacts page of the IBM Software Support Handbook on the Web ([techsupport.services.ibm.com/guides/contacts.html](http://techsupport.services.ibm.com/guides/contacts.html)) and click the name of your geographic region.

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround for you to implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the IBM product support Web pages daily, so that other users who experience the same problem can benefit from the same resolutions.

For more information about problem resolution, see Searching knowledge bases and Obtaining fixes.



---

## Appendix C. Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

---

## Trademarks

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

AIX  
DB2  
developerWorks  
Domino  
eServer  
IBM  
iSeries  
Lotus  
OMEGAMON

Passport Advantage  
POWER3  
POWER4  
POWER5  
pSeries  
Rational  
Redbooks  
Tivoli  
WebSphere  
xSeries  
zSeries

Adobe, the Adobe logo, PostScript<sup>®</sup>, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel<sup>®</sup>, Intel logo, Intel Inside<sup>®</sup>, Intel Inside logo, Intel Centrino<sup>™</sup>, Intel Centrino logo, Celeron<sup>®</sup>, Intel Xeon<sup>™</sup>, Intel SpeedStep<sup>®</sup>, Itanium<sup>®</sup>, and Pentium<sup>®</sup> are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Microsoft, Windows, Windows NT<sup>®</sup>, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.



---

# Index

## A

- adminaudit.log 3
- administration daemon
  - audit log 3
  - error log 3
- audit log 4, 67
- audit.log 4
- Auditing for performance 69

## B

- bulkload error log 4
- bulkload.log 4

## C

- CBE features 101
- CD, inserting incorrect 18
- configuration
  - Configuration Tool 39
  - troubleshooting 40
- configuration tools log 5
- conventions
  - typeface x
- core files
  - AIX operating systems 10
  - description 9
  - HP-UX operating systems 10
  - Linux operating systems 9
  - Solaris operating systems 10
  - Windows operating systems 9
- customer support
  - see Software Support 106

## D

- DB2 error log 5
- DB2 rollbacks 68
- DB2 troubleshooting 45
- db2cli.log 5
- debugging
  - advanced output 43
  - description 14
  - ldtrc command 14
  - server debug mode 14
- directory names, notation xi

## E

- environment variables, notation xi

## F

- fixes, obtaining 105

## I

- ibmdiradm.log 3
- ibmslapd.log 6
- idsldap group 18
- idsldap user
  - requirements 18
- idsldaptrace utility 16
- idslink log 27
- idslink.log 27
- idslink.preview 27
- idsslapd trace 67
- idstools.log 5
- Incorrect CD inserted 18
- information centers, searching to find
  - software problem resolution 105
- installation
  - overview 17
  - prerequisite software 17
- installation logs 7, 17, 25, 26, 27
- installation troubleshooting
  - InstallShield GUI 29
  - operating system utility 32
- instance creation
  - idsicrt 37
  - Instance Administration Tool 37
  - troubleshooting 38
- Internet, searching to find software
  - problem resolution 105
- isolation levels 68

## K

- knowledge bases, searching to find
  - software problem resolution 105
- Known limitations
  - Partial replication 82

## L

- LDAP\_DEBUG 14
- LDAP\_DEBUG\_FILE 15
- ldapinst.log 17
- ldaplp\_inst.log 17
- ldtrc command 14
- LOGFILSIZ, modifying 68
- logs
  - administration daemon audit log 3
  - administration daemon error log 3
  - audit log 4
  - bulkload error log 4
  - configuration tools log 5
  - DB2 error log 5
  - DB2 installation
    - AIX 27
    - Linux 27
    - Solaris 27
    - Windows 26
  - DB2 uninstallation
    - Windows 26

- logs (*continued*)

- GSKit installation
  - Windows 27
- idslink 27
- installation 7, 17
  - AIX 26
  - Linux 26
  - Solaris 26
  - Windows 25
- lost and found log 6
- native packages 28
- overview 3
- server error log 6
- lost and found log 6
- lostandfound.log 6

## M

- memory leak 73
- memory, adding on Solaris 67
- messages, resolving 2
- migration troubleshooting 35

## N

- notation
  - environment variables xi
  - path names xi
  - typeface xi

## P

- path names, notation xi
- performance troubleshooting 67
- problem determination
  - describing problem for IBM Software Support 107
  - determining business impact for IBM Software Support 107
  - submitting problem to IBM Software Support 107
- publications
  - accessing online ix
  - ordering ix
  - related viii

## R

- replication
  - overview 55
  - troubleshooting 55

## S

- Secure Sockets Layer (SSL) 74
- server error log 6
- SLAPD\_OCHANDLERS environment variable 68

- Software Support
  - contacting 106
  - describing problem for IBM Software Support 107
  - determining business impact for IBM Software Support 107
  - submitting problem to IBM Software Support 107
- Support Tool
  - data collected 11
  - location of data collected 14
  - location of log 14
  - overview 11
  - using 13

## T

- thread stacks 73
- trace, idsslapd 67
- troubleshooting features, overview 1
- typeface conventions x

## U

- uninstallation logs 26
- uninstallation troubleshooting 32

## V

- variables, notation for xi

## W

- Web Administration Tool
  - troubleshooting 47







Printed in USA

GC32-1568-00

