# IBM Security Directory Suite
# 8.0.1

*Troubleshooting Guide*

**Edition notice**

**Note:** This edition applies to version 8.0.1.x of *IBM Security Directory Suite* (product number 5725-Y17 ) and to all subsequent releases and modifications until otherwise indicated in new editions.

# Contents

# About this publication

IBM® Security Directory Suite, previously known as IBM Security Directory Server or IBM Tivoli® Directory Server, is an IBM implementation of the Lightweight Directory Access Protocol.

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

## Accessibility

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

For more information, see "Accessibility features for IBM Security Directory Suite" in the IBM Knowledge Center.

## Statement of Good Security Practices

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

# Chapter 1. Virtual appliance troubleshooting

Use the information that is provided to resolve, work around, or troubleshoot issues that are related to IBM Security Directory Suite virtual appliance and local management interface (LMI).

For other known issues see the technical note and troubleshooting information, IBM Security Directory Suite 8.0.1.x Known Issues.

## Error occurs with future date or time setting

When the virtual appliance system time is changed to a future date or time an error occurs. Use the solution that is provided to work around this problem.

### Problem

If you set the system time to a future date or time, and then sync it with an NTP Server or set it back to a previous time, the following error occurs:

```
javax.net.ssl.SSLHandshakeException: Client requested protocol SSLv3 not enabled or not
supported
```

Also, the local management interface (LMI) is no longer accessible.

### Solution

To resolve this issue, take the following actions:

1. Run the following command from the virtual appliance command-line interface:

   ```
   lmi reset_lmi_cert
   ```

2. Enter YES to confirm.
3. You can now access the LMI successfully.

## Port is not cleaned up after remote DB2 configuration

### Problem

This issue occurs when the following steps are done:

1. Create a remote DB2 instance on a Linux system by using the **idscfgremotedb** script.
2. Clean up the instance by using the **-r** option.

The entry in the `/etc/services` file is not cleaned up for the port that was in use. Due to this issue, if an instance is created again by using the **idscfgremotedb** script, more than one entry for the same service port might be created. When you start the instance, there might be further issues.

### Solution

Check the DB2 logs to ensure that the service port that you specify when you create a remote DB2 instance is not already in use. It must not be listed in the `/etc/services` file.

## Message on Application Interface page is not translated

When you ping a server to test the connectivity on the **Application Interface**, the message that is returned is not translated.

### Problem

This issue is noticed when the following steps are done:

1. On the virtual appliance console **Application Interfaces** page, click **Test**.
2. In the **Ping Server** window, specify the IP address of the server.
3. Click **Test**.

If the server is not reachable, then the response appears in English, even if the locale is a different language.

This issue is a known issue. The messages are standard message output from the command line and hence they are not translated.

## Values on Fix Packs page are not translated

After you apply a shell patch, the values that are shown under each column on the Fix Packs page appear in English, even if you have selected a different locale.

This issue is a known issue.

## Invalid characters are observed in Log Retrieval and View window when log file is cleared

In the **Log Retrieval and View** window of virtual appliance console, sometimes invalid characters are displayed after you clear the log file.

This issue is seen because the process is still fetching the logs.

This issue is a known limitation. No loss in functionality occurs and logging continues to happen correctly.

## Last Successful Authentication Time Stamp plug-in changes `modifyTimeStamp` attribute

If the Last Successful Authentication Time Stamp plug-in is enabled, the attributes **`ibm-latestBindTimestamp`** or **`ibm-prevBindTimestamp`** are populated.

When these attributes are updated, the **`modifyTimeStamp`** attribute of the user is also changed each time, though it must not be changed unless an explicit **`ldapmodify`** operation is run on a user.

This issue is a known limitation.

## Delay when retrieving logs of large file size

When you use the virtual appliance console option, **Manage** > **Maintenance** > **Log Retrieval and View** to fetch large log files, you might experience a delay.

To work around this issue, set the log size threshold to the default size instead of unlimited file size.

Also, you can use the **Manage** > **System Settings** > **Support Files** option to download rotated files.

## Snapshots information lost after firmware update

After you install a firmware update to upgrade the IBM Security Directory Suite virtual appliance, all of the snapshot information is lost.

This issue is a known limitation.

To work around this issue, download the snapshot before a firmware update, so that the saved snapshot is available to upload after the firmware update, if required.

To download or upload a snapshot, log in to the IBM Security Directory Suite virtual appliance console, and take the following actions:

1. From the top-level menu of the virtual appliance console, click **Manage** > **System Settings** > **Snapshots**.
2. On the **Snapshots** page, do one of the following steps:
   - To save the snapshot, click **Download**, browse to the location where you want to save the snapshot, and save the file.
   - To upload a snapshot, click **Upload**, browse to select the snapshot that you want to upload, type a comment to describe the snapshot, and click **Save Configuration**.

## Directory Server instance fails to start

If the host name is not set when virtual appliance is configured, the Directory Server instance might fail when you try to start it.

The following error appears in the `traceibmslapd.log`:

```
2015-11-04T14:27:01.251576-5:00 T-1333217504
     Error - initialize_db: Either DBXAllocEnv() or DBXSetEnvAttr() failed with -103, rc=1
2015-11-04T14:27:01.251587-5:00 T-1333217504
     Error - finish_read_config: initialize_db failed rc=1
2015-11-04T14:27:01.251594-5:00 T-1333217504
     Error - finish_read_config: serious error encountered rc=1, rolling back any pending
transactions
2015-11-04T14:27:01.251601-5:00 T-1333217504 finish_read_config: returning rc=80
2015-11-04T14:27:01.251618-5:00 T-1333217504 rdbm_back_config32: returning rc=80
2015-11-04T14:27:01.251624-5:00 T-1333217504 ldap_getenv: LDAP_DEBUG_TIME=NULL
```

To resolve this issue, you must set the host name. To set the host name, log in to the virtual appliance command-line interface. Use the **management** > **hostname** command:

```
hostname set hostname
```

## Message about unconfigured appliance

When virtual appliance is in an unconfigured state, if anyone tries to access it using the local management interface (LMI), a message appears that you must use hypervisor console for appliance configuration.

If this message appears, you must reconfigure the initial virtual appliance settings. Follow the steps in the topic, Setting up the virtual appliance.

## Fix pack installation might result in failure

Fix pack installation might fail if processes are running.

If processes are running when you apply a fix pack, the fix pack installation might fail, or it can put your appliance into an inconsistent state.

Before you apply a fix pack by using the virtual appliance console option, **Manage** > **Firmware and Fix Pack** > **Fix Packs**, or through the virtual appliance command-line interface **fixpacks** > **install** command, you must take the following precautionary actions:

1. Back up the existing partition by using the firmware setting option, in case something goes wrong during the fix pack installation. See Managing the firmware settings.
2. Stop all running processes, such as Directory Server, Directory Administration Server, Web Administration Tool, Federated Directory Server, SCIM Target, and SCIM Service. See Managing servers with the Server Control widget.
3. Optionally, you can also take a snapshot by using the virtual appliance command-line interface, so that you can restore the configuration settings back to what they were before the installation. See Snapshots commands.

## Stop all processes before upgrade and backup operations

As a precautionary measure to avoid issues, you must ensure that all processes are in the stopped state before you proceed with upgrade and backup operations.

Stop all processes before you start the following operations:

- Installing a fix pack
- Creating or applying snapshots
- Creating a partition backup by using firmware settings
- Upgrading firmware

## Virtual Directory troubleshooting

Use the explanations and workarounds to troubleshoot Virtual Directory issues.

**Error message "Incorrect data" is displayed**

This error might be displayed if you entered a duplicate attribute mapping. which is not allowed.

**Error is displayed in banner area of virtual appliance console**

The error message that is displayed in the banner area of the virtual appliance console might not specify the exact cause. The error might be due to one of the following issues:

- The Directory Server is not started. Ensure that the Directory Server is started in ConfigOnly or normal mode.

**Incorrect page appears after you log in to Directory Server**

When you access the **Cluster Configuration** or **Virtual Directory View Configuration** page for the first time or after the Admin DN password is changed, you are required to log in to Directory Server. After you log in, you might be redirected to the **Server Configuration** page. This issue is a known limitation.

Select the page that you required from the **Configure** > **Virtual Directory** menu again.

**Fields in the Virtual Directory configuration pages are blank**

After you change the Admin DN password, when you open any of the Virtual Directory configuration pages, the fields are blank.

You must restart the local management interface (LMI). Run the following command from the virtual appliance command-line interface:

```
lmi restart
```

After the LMI is restarted, open any of the Virtual Directory configuration pages. A window appears, requesting you to log in to the Directory Server.

After you log in to Directory Server with the correct Admin DN credentials, the fields are populated.

**Server does not start properly**

In the virtual views configuration, you must not select the same Virtual Directory suffix that is also used in cluster configuration.

# Use an Admin account that is not Directory Server admin to engage the SCIM interface

Use the solution that is provided to work around this problem.

### About this task

Map any user name with `cn=root` admin of LDAP and use it when you are accessing the SCIM interface.

### Procedure

1. Log in to the IBM Security Directory Suite Virtual Appliance LMI.
2. Click **Configure Directory Suite** > **Advanced Configuration** > **Update Property** > **All properties**.
3. Click on **SCIM Service Property Files** and select the **SCIM.properties** file.
4. Add New Property to Set `mapTenantNames = true`.
5. Also add new properties to define any user to map to existing `cn=root` (or admin user of LDAP) as

```
newadmin.access = all
newadmin.ldapName = cn=root
newadmin.password = <new_admin_password>
newadmin.ldapPassword = <Password_for_cn=root>
```

6. Restart the SCIM service from LMI dashboard.
7. Access SCIM interface link: https://<SDS_VA_IP_OR_HOSTNAME>:8070/USERS
8. Provide the username as `newadmin` and password that was set in `<new_admin_password>` field for the `newadmin.password` property.

   **Note:** For more information, see Authentication of SCIM requests.

# Avoid GLGUP1012E events in the ISDS Virtual Appliance

There is a periodic check, even if the virtual appliance is configured not to connect over the internet. Use the solution that is provided to work around this problem.

### Procedure

1. In the virtual appliance dashboard, click **Manage** > **System Settings** > **Advanced Tuning Parameters**.
2. In the **Advanced Tuning Parameters** panel, click **New**.
3. Create the `update.disable.remote.discovery` key and the value to 1.

   With `update.disable.remote.discovery` set to 1 (true), the appliance no longer searches the Internet for updates.

# Problems while taking Snapshots

Use the solution that is provided to work around this problem.

### Problem

The following errors can be on the virtual appliance Local Management Interface (LMI) and the Command Line Interface:

- From the LMI: `System Error`
- From the CLI:

```
sdsva.example.com >snapshots> create
Unexpected error
```

### Cause

This error might be caused by a full virtual appliance disk. While taking a snapshot, monitor the Disk usage from the LMI. I there is any "High Disk Usage" message on Dashboard, check the free space on Disk usage tab. If space is very low, the snapshot creation might fail.

### Solution

1. From the local management interface, click **Configure Directory Suite** > **Advanced Configuration** > **Custom File Management** > **CustomOut** > **archived_logs** to clean the logs.
2. From the local management interface, click **Manage** > **Maintenance** > **Core dumps** to clean the core dumps, if there are any.

## Use diagnostic tools on ISDS Virtual Appliance LMI

Diagnostic tools capture and record details about how the program operates.

### About this task

The following sources of information can help you locate errors with the product or component.

**Logs**
The virtual appliance records system events during specific transactions. Log files contain levels of information about the product processes. Log files also include information about other software that is used to complete a task. Use the information in log files to help you isolate and debug system problems.

**Traces**
Trace data provides in-depth processing information to help you focus on a particular area that you suspect is causing a problem. Trace data is more complex and detailed than message data.

To view the virtual appliance event log, see the event logs. System events are logged when the system settings are changed and when problems occur with the system. Use the Event Log management page to view and to export system events on your network.

### Procedure

1. From the top-level menu of the Appliance Dashboard, click **Monitor** > **Logs** > **Event Log**.
   The Event Log page displays system events in the **System Events** tab.
2. From the **System Events** tab, do one of the following actions:
   - Click **Pause Live Streaming** to stop the live updating of the event log.
   - Click **Start Live Streaming** to resume live updating of the event log.
   - Click **Export** to download the displayed event log data to a CSV file.

     **Note:** The default file name is `export.csv`.

   For information about viewing and configuring component-specific and virtual appliance log, trace files and IBM Security Directory Suite logs, see Retrieving log files.

# `upload_firmware_tool` does not work with Java 8.0.5.30 when the IP address is used

The **`upload_firmware_tool`** that is shipped with the virtual appliance in 8.0.1.11 or 8.0.1.12, does not work and does not upload the `.pkg` file when Java 1.8 (8.0.5.30) is used and an IP Address parameter is specified in the command.

## Problem

The command results in the following `javax.net.ssl.SSLHandshakeException: java.security.cert.CertificateException` error:

```
bash-3.2# /opt/IBM/ldap/V6.4/java/jre/bin/java -jar FileUpload.jar  9.113.51.68 admin admin
temptrust.jks WebAS 8.0.1.12-ISS-ISDS_20190910-1053_dev.pkg
javax.net.ssl.SSLHandshakeException: java.security.cert.CertificateException: No subject
alternative names present
```

```
bash-3.2# /opt/IBM/ldap/V6.4/java/jre/bin/java -version
java version "1.8.0_201"
Java(TM) SE Runtime Environment (build 8.0.5.30 - pap6480sr5fp30-20190207_01(SR5 FP30))
IBM J9 VM (build 2.9, JRE 1.8.0 AIX ppc64-64-Bit Compressed References 20190124_408237 (JIT
enabled, AOT enabled)
OpenJ9   - 9c77d86
OMR      - dad8ba7
IBM      - e2996d1)
JCL - 20190207_01 based on Oracle jdk8u201-b09
```

## Solution

To upload the package file with the tool with Java 8.0.5.30, use the fully qualified domain name of the IBM Security Directory Suite virtual machine.

For example:

```
bash-3.2# /opt/IBM/ldap/V6.4/java/jre/bin/java -jar FileUpload.jar  sdsetz068.example.com admin
admin temptrust.jks WebAS 8.0.1.12-ISS-ISDS_20190910-1053_dev.pkg
File size: 1736432400
SERVER REPLIED:
upload completed successfully.

bash-3.2# /opt/IBM/ldap/V6.4/java/jre/bin/java -version
java version "1.8.0_201"
Java(TM) SE Runtime Environment (build 8.0.5.30 - pap6480sr5fp30-20190207_01(SR5 FP30))
IBM J9 VM (build 2.9, JRE 1.8.0 AIX ppc64-64-Bit Compressed References 20190124_408237 (JIT
enabled, AOT enabled)
OpenJ9   - 9c77d86
OMR      - dad8ba7
IBM      - e2996d1)
JCL - 20190207_01 based on Oracle jdk8u201-b09
```

# Chapter 2. Directory Server troubleshooting

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

## Directory Server troubleshooting overview

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

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

### Built-in troubleshooting features

The Directory Server contains several tools in addition to the operating system tools to help you determine the source of problems you encounter.

**Core file generation**

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

You must ensure that core file generation is enabled in order for core file information to be generated. For more information about core files and for instructions for enabling core file generation, see Core file generation.

**Error logs**

Error logs record error messages that occur during Directory Server processing. Directory Server detects and saves these errors in a text file. For more information, see “Utilities for logging” on page 10.

**Server audit logs**

Server audit logs record suspicious patterns of activity to detect security violations. If security is violated, the Server audit log can be used to determine how and when the problem occurred. Directory Server detects and saves these errors in a text file. For more information, see “Utilities for logging” on page 10.

### Tools for troubleshooting a Directory Server instance

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

**IBM Support Assistant Lite**

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

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

To get an overview and to know about the features of IBM Support Assistant Lite, see http://publib.boulder.ibm.com/infocenter/ieduasst/v1r1m0/index.jsp?topic=/

com.ibm.iea.selfassist/isalite/1.3/Overview.html. To download IBM Support Assistant Lite, visit http://www-01.ibm.com/software/support/isa/download.html.

## Message Reference

The Message Reference contains a list of messages that you might encounter when you use Directory Server. It includes messages that are displayed in the directory server logs, graphical user interfaces, and the command line. Use the unique message ID associated with a message to locate detailed explanations and suggested operator responses.

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

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

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

The following messages are not contained in the *IBM Security Directory Suite Message Reference*:

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

# Utilities for logging

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

For information about viewing the logs, see the *Administering* section in the IBM Security Directory Suite documentation. For information about resolving error messages that you find in the logs, see "Message Reference" on page 10.

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

**Administration Server log (`ibmdiradm.log`)**

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

A sample of the log is shown here:

```
05/06/2013 02:05:57 PM GLPADM056I Admin server starting.
05/06/2013 02:05:58 PM GLPCOM025I The audit plug-in is successfully loaded from
  libldapaudit.so.
05/06/2013 02:05:58 PM GLPCOM022I The database plug-in is successfully loaded from
  libback-config.so.
05/06/2013 02:05:58 PM GLPADM060I The admin server backup and restore server
  configuration entry is not enabled.
05/06/2013 02:05:58 PM GLPCOM024I The extended Operation plug-in is successfully
  loaded from libloga.so.
05/06/2013 02:05:58 PM GLPCOM003I Non-SSL port initialized to 3546.
05/06/2013 02:05:58 PM GLPADM028I Admin server audit logging is started.
05/06/2013 02:05:58 PM GLPADM004I 8.0.1.x.0 ibmdiradm started
05/06/2013 02:05:58 PM GLPSRV048I Started 5 worker threads to handle client requests.
```

**Administration Server audit log (`adminaudit.log`)**

Administration Server audit log is used to improve the security of the Administration Server. The directory administrator and administrative group members can use the records in the audit log to

check for suspicious patterns of activity to detect security violations. If security is violated, the audit log can be used to determine how and when the problem occurred and perhaps the amount of damage done.

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

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

A sample of the log is shown here:

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

**Server audit log (`audit.log`)**

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

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

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

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

A sample of the Server audit log is shown here:

```
2013-01-16-17:38:15.484-06:00--GLPSRV023I Audit logging started.
The audit configuration options are:
 ibm-slapdLog = C:\idsslapd-ldaptest\logs\audit.log,
  ibm-auditVersion = true,ibm-audit = true,
    ibm-auditFailedOPonly = true,ibm-auditBind = true,
  ibm-auditUnbind = true,ibm-auditSearch = true,
  ibm-auditAdd = true,ibm-auditModify = true,
  ibm-auditDelete = true,ibm-auditModifyDN = true,
  ibm-auditExtOPEvent = true,ibm-auditExtOp = true,
  ibm-auditAttributesOnGroupEvalOp = true,ibm-auditCompare = true,
  ibm-auditGroupsOnGroupControl = true.
2013-01-16-17:38:15.656-06:00--GLPSRV009I IBM Security Directory (SSL),
    Version 8.0.1.x    Server started.
AuditV3--2009-01-16-17:39:28.468-06:00--V3 anonymous Search--bindDN:
```

```
<*CN=NULLDN*>--client: 127.0.0.1:3792--connectionID: 1
--received: 2009-01-16-17:39:28.453-06:00-- No such object
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: cn=monitor
scope: wholeSubtree
derefAliases: neverDerefAliases
typesOnly: false
filter: (objectclass=*)
```

**Bulkload error log (`bulkload.log`)**

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

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

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

The following **bulkload** error log resulted:

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

**Tools log (`idstools.log`)**

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

The following sample shows the tools log:

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

**DB2 log (`db2cli.log`)**

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

The following sample shows the DB2 log:

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

SQLSTATE=58031

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

"
2013-09-13-19:18:30.native retcode = -1026; state = "C8";
       message = "SQL1026N"
    The database manager is already active.
```

**Lost and found log (`lostandfound.log`)**

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

The information that is logged for each replaced entry includes:

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

The following sample shows the lost and found log:

```
#Entry DN: cn=t6,o=ut1,c=us
#Operation type:Add
#Corrective action:Replace
#Entry createTimestamp: 20131106211242.000000Z
#Entry modifyTimestamp: 20131030202533.000000Z
#Supplier address: 9.53.21.187
dn: cn=t6,o=ut1,c=us
objectclass: person
objectclass: top
sn: aa
cn: aa
cn: t6
description: this should not be here
ibm-entryuuid: 0c4559de-0a76-4c91-96e4-5ae81d405466
```

**Server log (`ibmslapd.log`)**

The server log contains status and error messages that are related to the server. For example, it contains error messages that correspond to events of adding an already existing entry, deleting a non-existent entry, or messages related to replication conflict resolution.

The following sample shows the server log with no errors:

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

**Installation and uninstallation logs**

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

*Table 1. The default log file location of IBM Installation Manager*

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

**Backup status file (`dbback.dat`)**

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

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

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

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

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

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

# Other diagnostic tools

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

## Server debug mode

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

You must run the server command **`tools rundebugcode`** from a command prompt to enable debug output. The syntax is as follows:

```
ldtrc on
tools rundebugcode
```

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

After you run the `ldtrc on` command, you can also use the `-d` *debug_mask* with any of the server commands.

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

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

*Table 2. Debug categories*

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

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

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

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

```
ldtrc off
```

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

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

- Set the *LDAP_DEBUG_FILE* environment variable.
- On server commands, you can use the **-b** option to specify a file. If the *LDAP_DEBUG_FILE* environment variable is set and you use the **-b** option and specify a different file, the file you specify overrides the file that is specified in the environment variable.

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

**Note:** The **idsldaptrace** tracing utility can be used to dynamically activate or deactivate tracing of the Directory Server. See the *IBM Security Directory Suite Command Reference* for information about the **idsldaptrace** utility.

# Tracing and debugging LDAP client APIs

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

## Before you begin

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

## Procedure

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

```
sds server_tools idsenvvars -a LDAP_DEBUG -v debug_level
```

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

*Table 3. Debug levels*

| Decimal | Value | Description |
|---|---|---|
| 1 | LDAP_DEBUG_TRACE | Entry and exit from routines |
| 2 | LDAP_DEBUG_PACKETS | Packet activity |
| 4 | LDAP_DEBUG_ARGS | Data arguments from requests |
| 8 | LDAP_DEBUG_CONNS | Connection activity |
| 16 | LDAP_DEBUG_BER | Encoding and decoding of data |
| 32 | LDAP_DEBUG_FILTER | Search filters |
| 64 | LDAP_DEBUG_MESSAGE | Messaging subsystem activities and events |
| 128 | LDAP_DEBUG_ACL | Access Control List activities |
| 256 | LDAP_DEBUG_STATS | Operational statistics |
| 512 | LDAP_DEBUG_THREAD | Threading statistics |
| 1024 | LDAP_DEBUG_REPL | Replication statistics |
| 2048 | LDAP_DEBUG_PARSE | Parsing activities |
| 4096 | LDAP_DEBUG_PERFORMANCE | Relational back-end performance statistics |
| 8192 | LDAP_DEBUG_RDBM | Relational back-end activities (RDBM) |
| 16384 | LDAP_DEBUG_REFERRAL | Referral activities |
| 32768 | LDAP_DEBUG_ERROR | Error conditions |
| 65535 | LDAP_DEBUG_ANY | All levels of debug |

For example, specifying a bit mask value of 65535 turns on full debug output and generates the most complete information. To know more about debug levels, see “Server debug mode” on page 14.

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

```
sds server_tools idsenvvars -a LDAP_DEBUG_FILE -v filename
```

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

3. Run the application from the same terminal where you have environment set. Re-create the problem that you want to debug.

4. The debug information is captured in the file pointed by the *LDAP_DEBUG_FILE* environment variable. You can now use the information that is captured in the file to debug the problem. Or send this file to the IBM Support team for further analysis.

## Collecting an ASCII server trace at startup

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

### Procedure

1. Stop the Directory Server instance, if it is running. Issue the command of the following format:

   ```
   ibmslapd -I instance_name -k
   ```

2. Determine whether tracing is enabled. Issue the following command:

   ```
   ldtrc info
   ```

3. If trace is disabled (in "off" mode), issue the following command to enable tracing:

   ```
   ldtrc on
   ```

4. Start the Directory Server in DEBUG mode and redirect the output to a file specified in the **LDAP_DEBUG_FILE** environment variable.

   Issue the command in the following format:

   ```
   sds server_tools ibmslapd -n -h 65535
   ```

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

   ```
   ldtrc off
   ```

## Collecting a binary server trace at startup

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

### Procedure

1. Stop the Directory Server instance, if it is running. Issue the command of the following format: `ibmslapd -I` *`instance_name`* `-k`
2. Determine whether tracing is enabled or not. Issue the following command: `ldtrc info`.
3. If trace is enabled, disable the trace. Issue the following command: `ldtrc off`.
4. Enable binary tracing. Issue the following command: `ldtrc on -l 50000000`

   In the command, the value for the buffer size is set to 50 million bytes. This buffer size stores the latest 50 million bytes of trace record data in the shared memory. It flushes the oldest data when the 50-MB value is reached. If the command fails because of what might seem to be insufficient shared memory resources, you can scale the number down. However, less than 20 million might not provide the required information.
5. Start the Directory Server instance. Issue the command of the following format: `ibmslapd -I` *`instance_name`* `-n`.
6. Point the environment variable *TRCTFIDIR* to the *INSTANCE_HOME* directory. Use the following command:

   ```
   sds server_tools idsenvvars -a TRCTFIDIR -v INSTANCE_HOME/logs
   ```

here, INSTANCE_HOME is `/home/sdsinst1/idsslapd-sdsinst1`

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

   ```
   ldtrc dump trace.raw
   ```

   where `trace.raw` is the path name and file name that is used to capture the records in shared memory.
9. Collect the format and flow of the binary trace. Issue the following commands:`ldtrc fmt trace.raw trace.fmt ldtrc flw trace.raw trace.flw`. Send the `trace.fmt` and *trace.flw* files to support.
10. Disable tracing. Issue the following command: `ldtrc off`

# Collecting performance records dynamically

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

## About this task

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

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

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

## Procedure

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

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

2. Dump the trace to a binary trace file. Issue the following command:

   ```
   ldtrc dump trace.bin
   ```

3. Format the trace. Issue the following command:

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

## What to do next

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

```
ldtrc off
```

For more information about performance profiling, see the *Administering* section in the IBM Security Directory Suite documentation.

## Collecting a dynamic ASCII server trace

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

### Procedure

1. Verify the ports that are used by your Directory Server instance. Issue the following command:

   ```
   idsilist -I instance_name -a
   ```

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

   Run the command in the following format:

   ```
   sds client_tools idsldaptrace -p port -a admin_port -D adminDN -w adminPW
   \ -h hostname -l on -t start -m 65535 -o output file
   ```

3. Re-create the problem and issue the specific operation that is failing.
4. Disable the dynamic ASCII server trace.

   Run the command in the following format:

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

   **Note:** Run **idsldaptrace -?** to see the usage information for the command.

## Collecting trace information

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

### Enabling tracing from the command-line interface

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

### Procedure

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

   ```
   ldtrc on
   ```

   OR,

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

   **Note:** You can use the **idsldaptrace** command from any system that has the directory server client package installed. The Administration Server must be running for this command to work.
2. Enable the tracing for a specific process or a command. Select a debug level for the trace.
   For example, specifying a bit mask value of 65535 turns on full debug output and generates the most complete information. For more information about debug levels, see “Server debug mode” on page 14.

   You can use one of the following options to set the debug level that is base on the process or command that you want to trace.

   - Set the *LDAP_DEBUG* environment variable to specify the debug level. Set this environment variable with a value that you want to use for debug_mask. If the *LDAP_DEBUG* environment variable is set

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

```
sds server_tools idsenvvars -a LDAP_DEBUG -v debug_level
```

To disable the environment variable, use the **idsenvvars -d LDAP_DEBUG** command.

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

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

Operation 0 modifying entry cn=Configuration
```

Restart the Directory Server instance for the changes to take effect.

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

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

To start tracing a Directory Server, issue the **idsldaptrace** command of the following format:

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

To stop tracing of a Directory Server, issue the **idsldaptrace** command of the following format:

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

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

- To stop tracing, issue the following command:

```
ldtrc off
```

OR

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

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

## Enabling tracing with Web Administration Tool

You can use the **Web Administration Tool** to enable tracing.

### About this task

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

### Procedure

1. In the **Web Administration Tool** navigation area, under **Server administration**, select **Logs**.
2. On the expanded list, select **Start/Stop server trace** to enable or disable server tracing.
3. In the **Trace debug levels** field, you can specify the debug level. For more information about debug levels, see "Server debug mode" on page 14.
4. In the **Trace debug file** field, you can specify the output file to store trace information.

## Directory Server log and configuration file locations

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

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

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

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

# Instance configuration issues

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

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

## Configuration overview and common errors

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

### Configuration overview

You can use the virtual appliance console or the virtual appliance command-line interface for configuration tasks.

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

## Existing database instance and database configuration failure

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

## DB2 is not configured properly

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

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

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

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

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

## Server does not start after configuration file attributes are changed

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

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

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

## Transaction log is full

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

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

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

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

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

where *X* is greater than the currently defined size.

You can check the current log size by using the following command:

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

## Configuration issues

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

### Database configuration

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

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

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

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

### Standard sources of output

There are several "standard" sources of information available:

- The output on the screen

  All of the configuration programs are either started from a console command-line prompt or open a background console. As the database configuration progresses, status messages (and limited error messages) are displayed in the associated console window. If a problem occurs, copy these messages to the system clipboard and then save them in a file for the IBM Software Support teams.
- DB2 log files

  If the error is a direct error from DB2DB2, then DB2 often creates message or error files (in the `/tmp` directory on AIX, Linux, and Solaris platforms). If you have a database configuration problem on an AIX, Linux, or Solaris system, examine all of the files in the `/tmp` directory that were created around the time of the attempted configuration.

  On Windows systems, examine any DB2 error logs in your DB2 installation directory. The error logs are under the directory that is named for the instance you were trying to configure. For example, if you were trying to create an instance and database named `ldapdb2`, and if your DB2 was installed in `D:\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.

# DB2 issues

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

## DB2 license file expired

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

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

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

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

```
db2start
```

If your license is correct, you see the message:

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

Otherwise, you see a message that your license is expired or is going to expire in some number of days.

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

- You have a demonstration license.

  1. You must upgrade your DB2 product from a demonstration license to a product license. Copy the license file from the DVD to the system whereDB2 is installed. You are not required to reinstall DB2.

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

     – If you have a DVD: *mount_point*/*db2_installable_directory*/license/db2ese_o.lic (or *cdrom_drive:**db2_installable_directory*\license\db2ese_o.lic for Windows)
     – If you downloaded a .zip file for installation: *directory_where_file_was_unzipped*\sdsV8.0.1.x*db2_installable_directory*\license\db2ese_o.lic
     – If you downloaded a .tar file for installation: *directory_where_file_was_untarred*/sdsV8.0.1.x/*db2_installable_directory*/license/db2ese_o.lic

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

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

     ```
     db2licm -a license_filename
     ```

- You purchased a different DB2 product.

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

  ```
  db2licm -a license_filename
  ```

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

## Recovery from migration failure in DB2

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

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

See the section *Overview of online backup and restore procedures for IBM Security Directory Suite* in *Administering* section in the IBM Security Directory Suite documentation to know more about DB2 backup and restore.

## DB2 diagnostic information in db2diag.log

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

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

**For DB2, Version 10.5.0.3**

**On Windows platforms**
The DB2 diagnostic error logs are written to the `Documents and Settings\All Users\Application Data\IBM\DB2\`*`Copy Name`*`\`*`instance`*, where *Copy Name* is the name of DB2 copy.

**On AIX, Linux, and Solaris platforms**
The DB2 diagnostic error logs are written to *INSTHOME*`/sqllib/db2dump`, where *INSTHOME* is the home directory of the instance.

## SQL0964C error when large amount of data is loaded

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

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

```
SQL0964C SQLSTATE=57011
```

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

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

   ```
   su - db2instownername
   ```

   where, *db2instownername* is the name of the DB2 instance owner.
2. Determine the current log file size setting by issuing the command:

   ```
   db2 get db config for db2instancename | grep -i logfilsiz
   ```
3. Increase the size of the log file size setting by issuing the command:

   ```
   db2 UPDATE db cfg for db2instancename using LOGFILSIZ new_value
   ```
4. Stop the slapd process.
5. To apply the changes that are related to database, issue the following command:

   ```
   db2 force applications all
   ```

6. Restart the slapd process.

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

## Instance starts in config-only mode after DB2 fix pack

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

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

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

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

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

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

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

## Remote DB2 with virtual appliance limitations and issues

Some issues that you might encounter with a remote DB2 configured with virtual appliance are described here with steps to resolve or work around the issues.

### UUID check issue

If you encounter issues with the UUID check, you can delete the UUID table:

```
drop table SYSTOOLS.ALLOW_SINGLE_SDS
```

The next server that starts up with the remote database instance automatically recreates the table and its entries.

### Running database utilities on remote DB2

The following utilities are not supported when IBM Security Directory Suite virtual appliance is configured to a remote DB2 instance:

- **bulkload**
- **dbrestore**
- **dbback**
- **perftune**

However, these utilities are available for use with the embedded DB2 that comes with the IBM Security Directory Suite virtual appliance.

## The idscfgdb command fails for a remote DB2

When you configure remote DB2, even though correct parameters are specified, the **idscfgdb** command might fail and gives the following error:

```
GLPCTL020I Updating the database manager: 'sdsinst1'.
Failed to attach to database instance or node: 'idsrnode'
GLPCTL022E Failed to update the database manager: 'sdsinst1'.
GLPCTL014I Uncataloging database instance node: 'sdsinst1'.
```

The trace or db2cli.log shows the following error message:

```
2015-10-30T19:45:06.420059-5:00native retcode = -30081; state = "08001";
message = "SQL30081N  A communication error has been detected.
Communication protocol being used: "TCP/IP".  Communication API being
used: "SOCKETS".  Location where the error was detected:
"192.168.184.128".  Communication function detecting the error: "connect".
Protocol specific error code(s): "113", "*", "*".  SQLSTATE=08001"
```

To resolve this problem, check the following causes and take the corresponding actions:

- A firewall might be enabled on the remote DB2 instance machine, which does not allow the IBM Security Directory Suite virtual appliance machine to communicate to the DB2 instance.
- The DB2 instance might not be running. Run the db2start command on the remote DB2 instance machine to start the instance.

## Search operations with remote DB2 might result in discrepancies

This issue occurs if you have IBM Security Directory Server 6.4 or earlier configured with DB2, and then configure IBM Security Directory Suite virtual appliance to use the existing DB2 as a remote DB. When you perform operations such as add, modify, and delete, on both the IBM Security Directory Server instance and the virtual appliance, the subsequent search operations might yield different results. To avoid this discrepancy, you must not run operations on the software stack instance after you configure virtual appliance to use remote DB.

## Suffixes are not accessible with remote DB2

If you have IBM Security Directory Server 6.4 or earlier configured with DB2, and then configure IBM Security Directory Suite virtual appliance to use the existing DB2 as a remote DB without doing migration, the existing suffixes are not accessible. To resolve this problem, you must do the migration or use the **idscfgsuf** command to add the existing suffixes to the virtual appliance.

## Configuring remote DB2 by using idscfgdb command fails

Configuring remote DB2 with **idscfgdb** command fails with the error message, "Failed to update the KDB or STASH file in the DB2 configuration manager."

To resolve this issue, ensure that the host name of the virtual appliance is not the default. To set the host name, log in to the virtual appliance command-line interface. Use the **management** > **hostname** command:

```
hostname set hostname
```

## Reconfiguring a different remote DB2 over SSL fails

You might encounter this issue in the following scenario:

IBM Security Directory Suite virtual appliance is already configured with a remote DB2 over SSL. When you attempt to run the **idscfgdb** command to reconfigure virtual appliance to a different remote DB2 over SSL, an error occurs stating that the remote DB2 is already configured. After that, the Directory Server instance fails to start.

To resolve this issue, you must unconfigure the remote DB2 by using the **idsucfgdb** command. Then, reconfigure it to a new remote DB2.

### Directory Server startup fails after you configure a remote change log that was created for an existing Directory Server DB2 instance with the idscfgremotedb script

This issue might occur in the following scenario: You have IBM Security Directory Server 6.4 or earlier configured with the shipped DB2. If you create a change log database by using the **idscfgremotedb** script and attempt to configure this as the remote change log for Directory Server virtual appliance, then the Directory Server startup fails.

To resolve this issue, you must create the change log database for an existing Directory Server instance by using the **idscfgchglg** command that was included with that version of IBM Security Directory Server. The **idscfgremotedb** script must be used to create a change log database only for the Directory Server database that was created by using the **idscfgremotedb** script.

### Older database and change log entries are still seen after reconfiguration of remote DB2 and change log

This issue might occur in the following scenario: IBM Security Directory Suite virtual appliance is already configured with a remote DB2 along with change log. When you unconfigure the remote DB2 and remote change log and then reconfigure to the same remote DB2 and change log, you might see older entries for Directory Server and change log.

The reason is because unconfiguring the remote DB2 or the remote change log does not remove the database. To resolve this problem, you must remove the Directory Server database or change log database or their entries when required.

### Port is not cleaned up after remote DB2 configuration

This issue occurs when the following steps are done:

1. Create a remote DB2 instance on a Linux system by using the **idscfgremotedb** script.
2. Clean up the instance by using the **-r** option.

The entry in the `/etc/services` file is not cleaned up for the port that was in use. Due to this issue, if an instance is created again by using the **idscfgremotedb** script, more than one entry for the same service port might be created. When you start the instance, there might be further issues.

Check the DB2 logs to ensure that the service port that you specify when you create a remote DB2 instance is not already in use. It must not be listed in the `/etc/services` file.

### DB2 messages are displayed during remote DB2 configuration over SSL

While configuring virtual appliance with a remote DB2 over SSL, a few DB2 messages are displayed. These information messages are a result of DB2 limitations that are related to setting the SSL parameters.

## Web Administration Tool and application server issues

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

## Corruption of data entered in the Web Administration Tool

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

If data that you enter in non-English languages in the **Web Administration Tool** is damaged, take the following actions on the WebSphere Application Server administrative console tree:

- Select **Servers**.
- Select **Application Server**.
- Select the server that you want; for example, `server1`.
- Click **Process Definition**.
- Click **Java Virtual Machine**.
- Click **Custom Properties**.
- Click to create a property.
- In the **Name** field, type `client.encoding.override`.
- In the **Value** column, type `UTF-8`.
- Click **Apply**.
- Stop and restart the WebSphere Application Server.

## Migration of files before you patch or migrate Web Administration Tool

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

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

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

## Additional login panels fail

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

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

**For AIX, Linux, and Solaris systems:**

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

```
mozilla &
```

**For Windows systems:**

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

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

## Web Administration Tool in inconsistent state

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

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

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

## Incorrect language is displayed in Web Administration Tool

The tabs, table headers, and static list boxes in the **Web Administration Tool** are sometimes displayed in an incorrect language. This problem might be encountered only on the AIX operating systems. However, Solaris and Linux systems might encounter the same problem. Follow the steps to resolve this problem.

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

```
export LC_ALL=new language
export LANG=new language
```

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

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

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

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

```
export DISPLAY=valid_computer_name:0
```

On the *valid_computer_name* issue the command:

```
xhost +
```

## Microsoft Internet Explorer browser problems

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

Make the following changes to the cache setup:

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

## HTML special characters are not displayed correctly

Special characters in read-only data that comes from the server are not displayed correctly in the **HTML** page. This display problem is because of the way that the HTML is rendered by the web browsers.

For example:

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

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

## Web Administration Tool requires IBM SDK Java Technology Edition on Domino server

If you want to use the **Web Administration Tool** with a Domino® server, you must use the IBM SDK Java Technology Edition, Version 8.0.2.10 or later. Using the Java from Oracle results in communication exceptions.

The following limitations apply to the Domino server:

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

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

## Templates with object class that has no attributes

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

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

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

## Non-editable fields are displayed as editable

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

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

## Back and Forward buttons not supported

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

## Log on issues in Internet Explorer

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

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

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

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

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

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

## Web Administration Tool logon fails for new user

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

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

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

## Web Administration Tool backup creates another folder

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

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

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

## WebSphere Application Server on AIX

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

Starting the WebSphere Application Server 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 WebSphere Application Server from 9090 to 9091, which is the port that is used on AIX computers, edit the *WAS_inst_path*`/config/cells/DefaultNode/virtualhosts.xml` file and change 9090 to 9091. Do the same thing in the *WAS_inst_path*`/config/cells/DefaultNode/nodes/DefaultNode/servers/server1/server.xml` file. *WAS_inst_path* is the path where the WebSphere Application Server is installed.

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

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

# Replication issues

When you use the Directory Server, you might encounter errors during replication. Use the explanations and information to troubleshoot and resolve issues that are related to replication.

## Replication overview

Directory Servers use replication to improve performance, availability, and reliability. Replication keeps the data in multiple Directory Servers synchronized.

Replication provides three main benefits:

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

For more information about replication, see the *Administering* section in the IBM Security Directory Suite documentation.

## Diagnosis of replication errors

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

### Sample replication topology

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

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

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

```
version: 1

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

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

dn: o=sample
```

```
objectclass: organization
objectclass: ibm-replicationContext

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

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

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

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

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

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

```
db2ldif -o out
```

The server configuration file for `masterhost` must contain:

```
dn: cn=Configuration

ibm-slapdServerId: masterhost-389
```

The configuration file for `replicahost` must contain:

```
dn: cn=Configuration

ibm-slapdServerId: replicahost-389
```

and the following entry

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

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

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

## Replication status

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

**Note:** The following **idsldapsearch** examples are based on the sample replication topology that is provided in the topic, “Sample replication topology” on page 33.

There are many operational attributes that provide replication status information when explicitly requested on a search. One of these attributes is associated with the entry that is the base of the replicated subtree, that is, the entry that the `ibm-replicationContext` object class was added to.

If you do a base search of that entry and request that the **ibm-replicationIsQuiesced** attribute is returned, the return attribute indicates whether the subtree was quiesced; for example:

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

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

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

The available attributes are:

**ibm-replicationLastActivationTime**
The time that the last replication session started between this supplier and consumer.

**ibm-replicationLastFinishTime**
The time that the last replication session finished between this supplier and consumer.

**ibm-replicationLastChangeId**
The change ID of the last update that is sent to this consumer.

**ibm-replicationLastGlobalChangeId**
The change ID of the last update to a global entry sent to this consumer. Global entries are things like `cn=schema` or `cn=pwdpolicy` that apply to the entire contents of a DIT.

This attribute is deprecated.

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

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

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

**Waiting**
Waiting for next scheduled replication time.

**Binding**
In the process of binding to the consumer.

**Connecting**
In the process of connecting to the consumer.

**OnHold**
This replication agreement is suspended or "held".

**Error log full**
The replication errors that occurred are more than can be logged. The number of errors that can be logged is based on the configured value for `ibm-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 that is returned from the consumer for the last update. This attribute is available only if the replication method is single threaded.

**ibm-replicationPendingChangeCount**
The number of updates queued to be replicated to this consumer.

**ibm-replicationPendingChanges**
Each value of this attribute gives information about one of the pending changes in the form:

```
change ID operation entry DN
```

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

**ibm-replicationChangeLDIF**
Gives the full details of the last failing update in LDIF. This attribute is available only if the replication method is single threaded.

**ibm-replicationFailedChangeCount**
Indicates the total number of failed changes that are logged for the specified replication agreement.

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

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

## Viewing replication errors with the Web Administration Tool

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

### Procedure

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

   From the **View errors** panel, you can:

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

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

   a. Select the replication error that you want to view from the **Replication error management** table.
   b. Click **View details** on the toolbar. The **Replication error details** table contains the following information about the selected error.

   **Supplier**
   The host name or IP address of the supplier.

   **Consumer**
   The host name or IP address of the consumer.

   **Change ID**
   The unique ID of the failed update that is sent to the consumer.

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

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

**Details**
The LDIF representation of the entry that is associated with the failed update, including all the operational attributes.

**Controls**
The controls that are used during the update.

## Viewing replication errors with the `idsldapsearch` command

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

### About this task

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

- **ibm-replicationFailedChanges**
- **ibm-replicationFailedChangeCount**

### Procedure

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

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

   This command returns an output similar to the following output:

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

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

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

   This command returns an output similar to the following output:

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

3. Use the **idsldapexop** command to try the update again:

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

   This command returns an output similar to the following output:

```
Operation completed successfully.
```

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

4. Run the **idsldapsearch** command again:

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

This command returns an output similar to the following output:

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

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

5. 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 returns an output similar to the following output:

```
Operation completed successfully.
```

6. Run the **idsldapsearch** command again:

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

This command returns an output similar to the following output:

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

It is also possible to try and delete all failures again 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 that are replaced because of replication conflict resolution. You can use the log of these entries to recover the data in the replaced entries, if necessary.

The information that is logged for each replaced entry includes:

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

## Write and replicated write messages

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

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

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

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

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

## `ibm-replicaSubentry` object class in a replication topology

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

When a Directory Server that is in a replication environment starts, it compares its server ID against the server IDs in the `replicaSubEntry` entry. If the server ID matches, then as per the attribute value of **`ibm-replicationServerIsMaster`**, the server either plays the role of a supplier or consumer. If the server ID does not match, the server assumes that it is a consumer in a replication topology.

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

For example:

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

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

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

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

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

## Command-line utilities to view replication status

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

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

*Table 4. Replication-related operational attributes*

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

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

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

For example,

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

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

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

For example,

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

To know more about replication status, see Monitoring replication status in the *Administering* section in the IBM Security Directory Suite documentation.

## *IBMSLAPD_REPL_UPDATE_EXTRA_SECS* environment variable

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

The default timeout for any change to be completed through replication is 60 seconds. The replication updates might involve many changes, such as adding a large group entry or adding or modifying entries that contains large objects such as credentials. In such cases, the update operation might require more than 60 seconds to finish. Update operations include `add`, `delete`, `modify`, or `modifydn`

operations. If any such single update operation through replication takes more than 60 seconds, the supplier server times out that update operation. It tries again to send the same update through replication. To extend the timeout duration for update operations in replication, you can use the *IBMSLAPD_REPL_UPDATE_EXTRA_SECS* environment variable.

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

- Minimum: 1
- Maximum: 2147483647

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

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

**Adding the variable to configuration file**

**Using the LDAP client utility:**

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

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

2. Restart the Directory Server instance.

**Using the Web Administration Tool:**

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

**From command prompt**

Set the environment variable *IBMSLAPD_REPL_UPDATE_EXTRA_SECS*.

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

```
export IBMSLAPD_REPL_UPDATE_EXTRA_SECS=180
```

**On Windows systems:**

```
set IBMSLAPD_REPL_UPDATE_EXTRA_SECS=180
```

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

# Information for troubleshooting replication

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

## Replicated suffix

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

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

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

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

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

where the `mod.ldif` file contains:

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

## Verify that suffixes and replication agreements exist

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

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

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

where *o=sample* is the replication context.

If this command does not return any results, the suffix is not configured to be replicated. You must configure the suffix to be replicated. See the *Administering* section in the IBM Security Directory Suite documentation for instructions about configuring a suffix for replication.

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

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

where *replctxt* is the location where the replication agreements for a replication context are stored; for example, o=sample. If the command does not return results, the replication agreement might not exist. To replicate correctly, the correct replication agreements must exist. See the *Administering* section in the IBM Security Directory Suite documentation for instructions for adding replication agreements.

## Peer-to-peer replication error

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

You might see an error similar to the following error:

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

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

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

To correct this error, complete the following steps:

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

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

   where:

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

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

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

## Replication topology extended operation fails with result code 80

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

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

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

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

If this error occurs, check for the following situations:

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

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

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

## Replication command-line interface error

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

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

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

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

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

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

## Entries in LDIF file are not replicated

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

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

Ways in which problems can occur include the following situations:

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

## Problem with cn=ibmpolicies subtree

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

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

To remove the entries, complete the following steps:

1. Search the Directory Servers to get the entries. Issue the following command:

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

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

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

dn: ibm-replicaserverid=ac1156c0-a214-1029-934c-cd9424fd6984,\
ibm-replicagroup=
default,cn=ibmpolicies
objectclass: top
objectclass: ibm-replicasubentry
ibm-replicationserverismaster: TRUE
cn: V8.0.1.x Migration
ibm-replicaServerId: ac1156c0-a214-1029-934c-cd9424fd6984
```

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

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

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

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

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

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

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

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

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

## Master server becomes unstable or stops

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

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

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

```
objectclass: top
objectclass: ibm-slapdConfigUlimit
objectclass: ibm-slapdConfigEntry
```

Then, configure the system `ulimit` values to:

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

Restart the servers for the changes to take effect.

## Stopping a multithreaded replication supplier

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

### About this task

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

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

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

### Procedure

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

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

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

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

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

3. Find all the replication agreements for the server.

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

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

4. Verify the status for a specified replication agreement.

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

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

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

5. Suspend replication for the agreement.

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

Operation completed successfully.
```

6. Verify the status of replication agreement.

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

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

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

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

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

   **c**
   Indicates the connection number. In the output of step 6, there are four connections. The first connection shows the most traffic. The workload determines how often the other connections are used.

   **l**
   Indicates the size limit for each queue. In the example, the value is 10. For each connection, there are two queues of same length. One queue is for updates to be sent on the connection, which is called the `send` queue. The other queue is for updates that were sent but no response was received from the consumer, which is called the `receive` queue. When updates are sent, they are moved from `send` queue to the `receive` queue. When the `receive` queue reaches its size limit, no more updates are sent until some responses from the consumer are received. When the `send` queue reaches its size limit, no more updates are assigned to the connection. When the size limit for all the `send` queues of connections are reached, the supplier waits for the consumer to process the backlog.

   **op**
   Specifies the replication ID of the last operation that is assigned to the `send` queue of the connection. Replication IDs are assigned to all updates that are to be replicated to a consumer. The process of assigning replication IDs must not stop even if replication is suspended.

   **q**
   Specifies the current size of the `send` queue. This value must not change if replication is suspended.

   **d**
   Specifies the count of dependent updates. An `add` request for an entry followed by a `modify` request of the same entry is counted as a dependency. All dependent updates must be assigned to the same connection so that they can be applied in correct order.

   **ws**
   Indicates the number of times the `send` queue reached its size limit.

   **ds**
   Specifies the number of dependent updates sent.

   **wd**
   Specifies the number of times that the `send` queue waited for a dependent update before it sent more updates.

   **wr**
   Indicates the number of times the `receive` queue reached its size limit.

   **r**
   Indicates the number of replicated updates that is waiting for a response from the consumer.

   **e**
   Specifies the number of replication errors reported by the consumer.

   **ss**
   Indicates the session count of the sender thread. It is incremented when a connection to the consumer is established.

   **rs**
   Indicates the session count of the receiver thread.

   The value indicated by `r` might show that the number of replicated updates that are waiting for a response is 0. In this case, it is safe to stop the supplier for this consumer server. The value of `r` varies between 0 and the value of `l`, the size limit of the queue, which defaults to 10. If the value is

not 0 for `r`, you must wait for it to be 0. The value of `r` depends on the size and type of the replication update and the workload on the consumer. After this value becomes 0, the supplier sends the status of updates to the consumer. On restarting the supplier, it replicates only the updates that were not sent before.

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

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

### Results

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

## Synchronizing Directory Servers in a replicated environment

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

### About this task

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

### Procedure

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

```
idsdb2ldif -o /tmp/M1.ldif
```

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

```
idsucfgdb -I instance_name -r
```

7. Configure the database on R1 and R2. Issue the command of the following format:

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

8. Synchronize the modified schema. Copy the `V3.modifiedschema` from M1 over to R1 and R2. The modified schema, `V3.modifiedschema`, is in the *instance_home*`/idsslapd-`*instance_name*`/etc` directory.
9. Synchronize the `ibmslapddir.ksf` file. To know more about Synchronizing two-way cryptography between server instances, see the *Administering* section in the IBM Security Directory Suite documentation.

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

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

```
idsldif2db -i /tmp/Master.ldif -r no
```

11. Start the R1 and R2 servers.

### Results

**Note:** On Windows platform, change the paths accordingly.

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

## Multimaster configurations

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

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

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

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

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

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

### Options for replication filter and replication method are not available

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

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

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

# Performance issues

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

## Identification of performance problem areas

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

### Server audit log

The server audit log shows the searches that are being done and the parameters that are used in each search. The server audit log also records when a client binds and unbinds from the directory. By observing these measurements, you can identify LDAP operations that take a long time to complete.

### idsslapd trace

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

1. `ldtrc on`
2. `idsslapd –h 4096`

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

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

```
ldtrc off
```

## Setting the *SLAPD_OCHANDLERS* environment variable

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

### About this task

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

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

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

### Procedure

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

```
ibm-slapdSetenv: SLAPD_OCHANDLERS=15
```

3. Stop and restart the server.

## DB2 rollbacks and isolation levels

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

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

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

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

```
db2set isolation_level=YES
```

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

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

## Default value of LOGFILSIZ must be increased

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

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

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

**For AIX, Linux, and Solaris platforms:**

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

To increase this value, issue the following command:

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

**For Windows platforms:**

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

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

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

## Audits for performance profiling

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

The following hotspots are identified for auditing:

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

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

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

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

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

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

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

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

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

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

**Note:** Attribute cache is deprecated. Henceforth, users must avoid the use of attribute cache.

- Deadlock detector
- RDBM locks

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

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

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

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

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

# Information for troubleshooting in various scenarios

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

## Server is not responding

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

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

If you determine that the server is not responding, run the **IBM Support Assistant Lite** tool. This tool gathers information that you can provide to IBM Software Support to help identify the problem. See “Tools for troubleshooting a Directory Server instance” on page 9 for information about the **IBM Support Assistant Lite** tool.

## Memory leak is suspected

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

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

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

The script is as follows:

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

while [ true ]; do
echo | tee -a /tmp/monitor.out
echo 'Begin Monitoring.....' | tee -a /tmp/monitor.out
date | tee -a /tmp/monitor.out
echo 'Process info via ps auwx command: ' | tee -a /tmp/monitor.out
ps auwx | egrep '(slapd|$instance|PID)' | grep -v grep | tee -a /tmp/monitor.out

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

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

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

date | tee -a /tmp/monitor.out

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

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

sleep 600 #10minutes

done
```

## SSL communications return errors

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

Run the following command:

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

Where

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

Record and send any errors to IBM Software Support.

## Recovering data from a Directory Server instance where encryption seed value is lost

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

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

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

   ```
   idsadduser -u newinst -w newinst -l /home/newinst -g idsldap
   ```

2. Configure a Directory Server instance. Issue the commands of the following format:

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

   **Note:** Save the encryption seed `thisismynewencryptionseed`.

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

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

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

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

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

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

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

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

## Character sets larger than 7-bit ASCII in passwords

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

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

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

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

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

## Premature expiry of user password

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

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

## Troubleshooting the limitation in the idssethost command

You might configure a Directory Server instance to listen on a specific interface by using the **idssethost** command. The **idssethost** command configures the directory server instance to listen only on the specific IP address. Follow the steps to work around this limitation.

### About this task

An entry is added to the ibmslapd.conf file:

```
ibm-slapdIpAddress: IP_address
```

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

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

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

If ibmslapd is required to listen on a specific interface and the loopback interface, the directive to listen on loopback must be added manually. Perform the following steps:

### Procedure

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

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

where, *filename* contains:

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

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

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

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

```
n: cn=Configuration
cn: Configuration
ibm-slapdAdminDN: cn=root
ibm-slapdAdminGroupEnabled: true
ibm-slapdAdminPW: {AES256}ohtCABBYFbFo7jREOPz/zQ==
ibm-slapdCryptoSync: vDydxlFDW0xKtWBL
ibm-slapdDerefAliases: always
ibm-slapdIpAddress: 10.10.10.10
ibm-slapdIpAddress: 127.0.0.1
ibm-slapdPort: 389
#ibm-slapdPwEncryption must be one of:
# none/aes128/aes192/aes256/crypt/sha/ssha/md5/
# sha224/sha256/sha384/sha512/ssha224/ssha256/ssha384/ssha512
ibm-slapdPwEncryption: aes256
ibm-slapdServerBackend: RDBM
...
```

3. Restart the Directory Server instance.

### Results

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

## Environment with SNMP agent configured

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

Some of the likely scenarios and their workaround are listed:

- When the **idssnmp** tool is running for a long time, it is observed that the `LDAP_HOME/idstools/snmp/logs/ibmdi.log` file size grows large.

  If a user wants **idssnmp** to generate or keep less amount of log, the user can modify the `SDI_HOME/etc/log4j.properties` file and configure an appropriate log appender. For more information about the list of appenders, see *Installing and Administering* section in the IBM Security Directory Integrator documentation.
- To run the **idssnmp** tool over SSL, user must edit the `LDAP_HOME/idstools/snmp/solution.properties` file and specify the certificate information.
- If a user wants to run **idssnmp** over SSL and the `solution.properties` file is not present, the user can create the solution files that are required by `idssnmp` by running the following command:

```
SDI_InstallDirectory/ibmdisrv -s LDAP_HOME/idstools/snmp –v
```

This command creates the `solution.properties` file in the `LDAP_HOME/idstools/snmp` directory.

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

## Tombstone entries in a Directory Server

Before entries are permanently deleted from the RDBM database, a subtree is created to hold the entries to be deleted with operation attributes.

The to-be-deleted entries are moved to the tombstone subtree, `cn=Deleted Objects`, and the attribute table is updated for the entry to mark the entry as deleted by adding attribute such as **isDeleted**. This feature is supported only on the primary RDBM back-end of the Directory Server. Tombstones are not supported in configuration, schema, or change log back-end.

There might be situations where data inconsistency gets introduced by entry deletions when this feature is enabled, which requires the intervention of the directory administrator. For performance reasons, there is no check that is provided, which can possibly prevent tombstones entries with the attribute **isDeleted** set to TRUE from being accidentally created or modified under any subtrees.

You can identify these entries in an RDBM back-end database by comparing the searches. Compare the search results that are returned by a normal search with a search base to that returned by a null base search.

For example, consider an RDBM back-end database with two subtrees: `o=sample` and `cn=Deleted Objects`. where, `o=sample`, contains three entries: `cn=A`, `cn=B`, and `cn=C` (with `isDeleted=TRUE`). The subtree `cn=Deleted Objects` containing entries, `cn=X`, `cn=Y`, and `cn=Z` (without `isDeleted=TRUE`).

When searches that use a search base and null base are requested without including the return deleted object control, the following results are displayed.

- In a search with a search base `o=sample` and search filter, `objectclass=*`, all entries under the search base, including entries with `isDeleted=TRUE`, are displayed.
- In a null base search with search filter, `objectclass=*`, all entries except for those entries with `isDeleted=TRUE` are displayed.

*Table 5. The results of different search base with search filter, `objectclass=*`*

| Subtree search base | Search filter | With control | Search results | Remarks |
|---|---|---|---|---|
| `o=sample` | `objectclass=*` | No | `cn=A`<br>`cn=B`<br>`cn=C` | `cn=C` is a normal entry with `isDeleted=TRUE` |
| `null` | `objectclass=*` | No | `cn=A`<br>`cn=B`<br>`cn=Z` | List LDAP_ENTRY table with `isDeleted!=TRUE`. `cn=C` is not qualified. |

It is possible that the `isDeleted` attribute is accidentally deleted or is set to FALSE for entries under the tombstone subtree. When searches that use a search base and null base are requested with the return deleted object control, the following results are displayed.

- In a search with a search base, `cn=Deleted Objects`, and search filter, `objectclass=*`, all entries under search base are returned. However, when a search with a search base, `cn=Deleted Objects`, and search filter, `isDeleted=TRUE`, is requested, entries with `isDeleted=FALSE` are not returned.

- In a null base search with search filter, `objectclass=*`, all entries in the database are displayed. However, when a null base search with search filter, `isDeleted=TRUE`, is requested, only the entries with attribute `isDeleted=TRUE` in the database are displayed.

*Table 6. The results of different search base with different search filters*

| Subtree search base | Search filter | With control | Search results | Remarks |
|---|---|---|---|---|
| `cn=Deleted Objects` | `objectclass=*` | Yes | `cn=X`<br>`cn=Y`<br>`cn=Z` | `cn=Z` is a tombstone without `isDeleted=TRUE` |
| `cn=Deleted Objects` | `isDeleted=TRUE` | Yes | `cn=X`<br>`cn=Y` | |
| `null` | `objectclass=*` | Yes | `cn=A`<br>`cn=B`<br>`cn=C`<br>`cn=X`<br>`cn=Y`<br>`cn=Z` | List the `LDAP_ENTRY` table, including entries with `isDeleted=TRUE`. |
| `null` | `isDeleted=TRUE` | Yes | `cn=C`<br>`cn=X`<br>`cn=Y` | |

**Note:** Deletion of schema attributes might fail because some of tombstone entries still reference them. A delete, rename, or restore of a tombstone entry is not replicated. It might result in data inconsistencies on replicas such as rename and restore cases.

## Directory Server instance backup

You can take multiple backups of the Directory Server instance, both offline and online, at multiple locations at different points in time. Understand how to work with the Directory Server instance backup in different scenarios.

For example, if `myinst1` is the Directory Server instance and *`instance-location`*`/idsslapd-myinst1/backup1`, *`instance-location`*`/idsslapd-myinst1/backup2`, and *`instance-location`*`/idsslapd-myinst1/backup3` are the locations where backups are stored at T1, T2, and T3 time (where, T1<T2<T3).

When the instance, `myinst1`, is dropped or the database instance for the instance is unconfigured, the database configuration details (**dbbackuponline** and **clbackuponline**) are set to FALSE in the *`instance-location`*`/idsslapd-myinst1/backup3/dbback.dat` file.

**Scenario 1**

If the instance, `myinst1`, is re-created and configured with the backup location set to *`instance-location`*`/idsslapd-myinst1/backup2`, where the offline backup of the previous instance was stored before it was dropped. In this case, the results of the monitor search, `"cn=backup,cn=monitor"`, will be consistent for searches that are done before the directory server instance is started (server state as stopped) and after the Directory Server instance is started (server state as running).

**Scenario 2**

If the instance, `myinst1`, is re-created and configured with the backup location set to *`instance-location`*`/idsslapd-myinst1/backup1`, where the online backup of the previous instance was stored before it was dropped. In this case, the results of the monitor search,

"cn=backup,cn=monitor", are not consistent for searches that are done the Directory Server instance is started (server state as stopped) and after the directory server instance is started (server state as running).

The reason is when the server is in stopped state monitor search refers *instance-location*/idsslapd-myinst1/backup1/dbback.dat file for the online status of database and change log (which is true since the previous online backup was stored at *instance-location*/idsslapd-myinst1/backup1), and when the server is in running state monitor search refers directory server for the online status of database and change log (which is false as database is not configured for online backup). Suppose that a user starts the Directory Server. Then, the user does online backup that is based on the monitor search results that the user received when the Directory Server was in stopped state. In this scenario, the user would get unexpected behavior because the database is not configured for online backup.

If a user intends to restore from an existing backup for a re-created Directory Server instance, it is alright to configure to a previous backup location. However, suppose that the user intends to back up the re-created Directory Server instance. Then, to avoid the situation as mentioned in scenario 2, it is advisable to remove previous backup files from the location. Or else the user must specify a location that does not have any backup image.

## Configuration of preaudit records for serviceability

Sometimes, when a Directory Server locks up or stops, the audit log might not have a record of the operation that causes the problem. You can configure the auditing of preaudit records to record operations that were not completed.

The audit log does not record the operation that causes the problem because the audit logs are updated after the Directory Server back-end completes the operation. So, any problems that occur before the audit records get updated are not logged and the result of the operation is unknown.

You can configure auditing of preaudit records to record operations that were not completed. When preaudit records are enabled, the audit plug-in is called to update audit records before the operation completes. When preaudit is enabled, the thread ID is also logged in the audit header. To enable pre-auditing, you must set the value of the *IBMSLAPD_PREOP_AUDIT* environment variable to YES. You can set this value by accessing the environment variable or by using the **ldapmodify** command with the following format:

```
ldapmodify -D adminDN -w adminPW
dn: cn=Front End, cn=configuration
changetype: modify
add: ibm-slapdSetEnv
ibm-slapdSetEnv: IBMSLAPD_PREOP_AUDIT=YES
```

An example of a pair of diagnostic audit records when preaudit is enabled, where the sequence identifier is 3: *PREOP: 3* and *POSTOP: 3*, is as follows:

```
AuditV3--2007-08-29-11:44:32.912-06:00DST--V3 PREOP: 3 threadId:  1161116592
Search--bindDN: cn=root--client: 127.0.0.1:1044--connectionID:
3--received: 2007-08-29-11:44:32.912-06:00DST--Success
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: o=sample
scope: baseObject
derefAliases: neverDerefAliases
typesOnly: false
filter: (objectclass=*)

AuditV3--2007-08-29-11:44:33.092-06:00DST--V3 POSTOP: 3 threadId:  1161116592
Search--bindDN: cn=root--client: 127.0.0.1:1044--connectionID:
3--received: 2007-08-29-11:44:32.912-06:00DST--Success
controlType: 1.3.6.1.4.1.42.2.27.8.5.1
criticality: false
base: o=sample
scope: baseObject
derefAliases: neverDerefAliases
typesOnly: false
filter: (objectclass=*)
```

## Entries that are displayed to root and anonymous users

No entries are displayed to root and anonymous users when logged on to a Proxy Server instance with the **Web Administration Tool**. You must change the page control restrictions for such users.

Root and anonymous users are not able to view entries with the **Web Administration Tool** when logged on to a Proxy Server, even with the **ibm-slapdAllowAnon** attribute under the DN entry `cn=Connection Management, cn=Front End, cn=Configuration` set to `true`. The reason is because the **Web Administration Tool** uses page control to browse through entries. If paging is enabled only for administrators by setting the **ibm-slapdPagedResAllowNonAdmin** attribute to `false` under the DN entry `cn=ProxyDB, cn=Proxy Backends, cn=IBM Directory, cn=Schemas, cn=Configuration`, then non-administrators are not allowed to browse through entries with the **Web Administration Tool**. This restriction is also applicable in the case of RDBM servers.

For the root and anonymous users to view the entries with the **Web Administration Tool**, the following must be considered:

- Set the **ibm-slapdAllowAnon** attribute under the DN entry `cn=Connection Management, cn=Front End, cn=Configuration` to `true`.
- Set the **ibm-slapdPagedResAllowNonAdmin** attribute under the DN entry `cn=ProxyDB, cn=Proxy Backends, cn=IBM Directory, cn=Schemas, cn=Configuration` to `true`.

After you set the attributes, restart the Directory Server instance.

**Note:** The root user is an anonymous user in the case of Proxy Server for DIT-related operations.

## Directory Server instance is restored to latest consistent state

When you run the restore operation on a Directory Server instance, the Directory Server instance might get restored to latest consistent state. It is restored to this state instead of getting restored to the point when the backup was done. Follow the steps to resolve this issue.

**Observed**

1. Create and configure a Directory Server instance for online backup.
2. Stop the Directory Server instance and run the initial offline backup either with the **Web Administration Tool** or the **idsdbback -u -k** command.
3. Add the suffix, `o=sample`, and start the directory server instance.
4. Add the entry `o=sample`.
5. Verify that the database parameter `LOGARCHMETH1` is correctly set.
6. Perform a restore operation either with the **Web Administration Tool** or the **idsdbrestore -k** command.
7. Verify that the suffix, `o=sample`, is not present (since backup was done before the suffix was added).
8. Add the suffix, `o=sample`, and start the directory server.
9. Run **ldapsearch** for the entry `o=sample` and you can observe that the entry is present.

**Expected result**

The entry, `o=sample`, must not be present because only the suffix, `o=sample`, is added after restore on a clean database (no data).

**Reason**

During roll-forward, DB2 scans the current logs in the `newlogpath`location. Because of the options that are specified in the roll-forward, DB2 scans the logs until the end, and restores a database to the latest consistent state.

For example, at the time of backup, suppose that you have 100 entries in the Directory Server. After the backup operation, if you delete five entries and then run the restore operation. You might still find the 95 entries in the directory rather than the 100 entries that you backed up. This reason is because, the latest consistent state of the database was after the deletion on five entries.

However, you can modify the options in the rollforward recovery operation such that the database is restored to the point where it had 100 entries. You must specify the timestamp of the last committed change. This timestamp is the one at which the 100th entry was added and to obtain this value of timestamp is difficult.

## Online backup and restore limitation

When you change the original backup location, you might encounter an error. Follow the steps to work around this limitation with the online backup and restore feature.

During online backup and restore, suppose that the folder name (backup location) to which online backup was initially configured is changed for the subsequent backups. It is observed that no error is thrown during the backup operation (**idsdbback**), but during the restore operation (**idsdbrestore**) the following error messages might be displayed:

```
...
GLPCTL101I Restoring backup database rdsdb to configured database rdsdb.
GLPCTL103E Failed to restore backup database rdsdb to configured database rdsdb.
GLPDBR004E Failed to restore Directory Server instance 'sdsadmin'.
GLPDBR028W The program did not complete successfully. View earlier error messages
 for information about the exact error.
...
```

**Reason**

As per the **idsdbback** and **idsdbrestore** (also available as **dbback** and **dbrestore**) design for online backup, the first-time backup must be a complete offline backup while the `ibmslapd` process is in stopped state. After the first offline backup, the online backup feature can be used while the `ibmslapd` process is running.

During the first offline backup, the **idsdbback** command takes the following options:

```
idsdbback -I instance_name -k /path/backupfolder1 -u [-a /path/logarchivefolder]
```

If the optional path for `logarchive` folder is not provided, the command uses a folder inside the `backupfolder1` folder (as per the example) to configure the `logarchivefolder` and sets this value in the corresponding DB2 database configuration parameter, **LOGARCHMETH1**.

If the backup folder is changed for a subsequent online backup, **idsdbrestore** fails if the previous backup folder does not exist, since the **LOGARCHMETH1** still points to the previously configured value.

**Confirming the problem**

To confirm, verify the *LOGARCHMETH1* variable for the corresponding database configuration.

```
su - instance_name
db2 list db directory
db2 get db configuration for databasename | grep -i LOGARCHMETH1
```

**Note:** Replace the *instance_name* and *databasename* with the appropriate names.

**Resolving the problem**

If you want to change the backup location after the first offline backup, or even after subsequent online backups, complete the following procedure to update the `backup` folder and `logarchive` folder values:

1. Stop the `ibmslapd` process.

   ```
   ibmslapd -I instance_name -k
   ```

2. Use the **idsbackup** command to update both the `backup` folder and `logarchive` folder.

   ```
   idsdbback -I instance_name -k /path/backupfolder2 \
   -a /path/backupfolder2/INACTIVE_LOGS -u -n
   ```

3. Start the **ibmslapd** process.

```
ibmslapd -I instance_name -n -t
```

## Log management servers fails to stop

When you attempt to stop the log management service after you start the service with the **Web Administration Tool**, it fails to stop. Follow the steps to work around this problem.

**Observed**

The log management service gets started when you use the **Start/Stop log management** panel of the **Web Administration Tool**. However, when you attempt to stop the log management service by using the **Start/Stop log management** panel, the panel displays that the service is stopped. However, the log management service is running in the background.

**Expected behavior**

The expected result is that the log management service must stop.

**Reason**

For the log management to work, IBM Security Directory Integrator is required. The IBM Security Directory Integrator installation wizard prompts for a installation location for the IBM Security Directory Integrator Solutions Directory. If you opt for the default preselected option, "Use a subdirectory named Directory Integrator under my home directory," this problem with log management service occurs.

**Workaround**

If you opt for the **Use Install Directory** option from the IBM Security Directory Integrator installation wizard for the Solutions Directory, the mentioned problem with log management service does not occur. For example, if the IBM Security Directory Integrator is installed in the location `/opt/IBM/ldap/version/TDI`, and you opt to provide the IBM Security Directory Integrator Solutions Directory within the `/opt/IBM/ldap/version/TDI` directory, then the problem with log management service is not observed.

## Instance does not start and returns error GLPCRY007E

In certain scenarios, the Directory Server instance might not start and might return error GLPCRY007E. Follow the steps to resolve this issue.

**Scenario**

1. Create a Directory Server instance, `inst1`, configure the instance, and start the instance. The encryption seed that is used to create the instance, `inst1`, is `thisismyseed`.
2. Drop the instance, `inst1`, without dropping the database that is associated with it.
3. Re-create the instance with the encryption seed, `thisismyseed`, and configure the instance with the existing database.
4. Start the instance.

**Observed**

The instance does not start and returns error:

```
GLPCRY007E The directory key stash file is inconsistent with the
 associated encrypted data.
```

**Reason**

When a Directory Server instance is created and is started, some information from key stash file (`.ksf`), is stored in the database. Therefore, an existing database cannot be used with a key stash file that gets created when an instance is re-created.

**Workaround**

In such case, if you intend to use an existing database with a new instance, then at the time of instance creation you must use **-e** and **-g** options to specify the encryption seed and encryption salt

values for the new instance. This encryption seed and salt value must be same as the encryption seed and salt value of the dropped instance.

If you did not provide the salt value with the **-g** option for the instance that you are intending to drop, then the salt value must be determined before an instance is dropped. Issue the **idsldapsearch** command of the following format to retrieve the salt value.

```
idsldapsearch -h IP address -p port -s base -b "cn=crypto,cn=localhost" \
objectclass=* ibm-slapdCryptoSalt
```

# Interoperability

Use the information on interoperability between IBM Security Directory Suite Directory Server and other Directory Servers to troubleshoot related issues.

## Interoperability with Novell eDirectory Server

When you do a simple bind by using IBM Security Directory Suite Directory Server client utilities against Novell eDirectory Server, you might encounter an error message. Run the configuration command to resolve this problem.

You might encounter an error message such as `ldap_bind: Confidentiality required`.

Run the following command:

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

## Interoperability with Microsoft Active Directory

If IBM Security Directory Suite Directory Server is configured over SSL by using `serverClientAuth` authentication, follow the steps to make it work with Microsoft Active Directory client **LDP.exe**.

To make IBM Security Directory Suite Directory Server configured over SSL by using `serverClientAuth` authentication to work with Microsoft Active Directory client **LDP.exe**, complete the following steps.

1. Select **Internet Information Services (IIS) Manager** from **Administrative Tools** in Control Panel.
2. On the left navigation panel, select the **Web Site** node.
3. Under the website node, right-click **Default Web Site**, and then select **Properties**.
4. On the Default website Properties dialog box, select the **Directory Security** tab.
5. To request for a new certificate, click **Server Certificate** under the Secure communications area. The **Web Server Certificate Wizard** is opened.
   a. On the **Server Certificate** page in the **IIS Certificate Wizard** dialog box, select the **Create a new certificate** option and click **Next**.
   b. On the **Delayed or Immediate Request** page, enter the required options and click **Next**.
   c. On the **Name and Security Settings** page, in the **Name** field enter the host name of the system and click **Next**.
   d. On the **Organization Information** page, specify appropriate names and click **Next**.
   e. On the **Your Site's Common Name** page, in the **Common name** field, enter the host name of the system and click **Next**.
   f. On the **Geographical Information** page, specify appropriate values and click **Next**.
   g. On the **Certificate Request File Name** page, in the **File name** field specify the path name and file name for the certificate request and click Next.
   h. The summary of the values is displayed. Click **Next**.
   i. Click **Finish**.

6. Send the certificate request by using these steps to any certificate authority (CA) to issue a certificate.
7. After you receive the server certificate, add the certificate by using **IIS Certificate Wizard**.
    a. On the **Pending Certificate Request** page, select the **Process the pending request and install the certificate** option and click **Next**.
    b. On the Process a Pending Request page, in the **Path and file name** field specify the path name and file name of the certificate. You can also use Browse to select the certificate. Click **Next**.
8. Export the personal certificate to `pfx` or `p12` format by using **IIS Certificate Wizard**.
    a. On the **Modify the Current Certificate Assignment** page, select the **Export the current certificate** to a `.pfx` file option and click **Next**.
    b. On the **Export Certificate** page, in the **Path and file name** field enter the path name and file name where `pfx` certificate to be stored. Click **Next**.
    c. On the **Certificate Password** page, in the **Password** and **Confirm password** fields enter the password and click **Next**.
    d. On the **Export Certificate Summary** page, the summary of the provided values is displayed. Click **Next**.
    e. Click **Finish**.
9. To import the certificate, double-click the stored `pfx` certificate. The **Certificate Import Wizard** is opened.
    a. On the **File to Import** page, in the **File name** field enter the path and file name of the `pfx` certificate and click **Next**.
    b. On the **Password** page, enter the password and click **Next**.
    c. On the **Certificate Store** page, select the **Place all certificate in the following store** option and click **Browse** and select **Personal** from the **Select the certificate store you want to use** list in the Select Certificate Store dialog box. Click **Next**.
    d. Click **Finish**.
10. To export the personal certificate in BER format, complete the following steps.
    a. Open Internet Explorer, select **Tools** > **Internet Options**.
    b. Select the **Content** tab in the **Internet Options** dialog box, and select **Certificates** under the Certificates area.
    c. On the **Personal** tab in the **Certificates** dialog box, select the certificate and click **Export**. The **Certificate Export Wizard**.
    d. On the **Export File Format** page, select the **Base-64 encoded X.509 (.CER)** option and click **Next**.
    e. On the **File to Export** page, in the **File name** field enter the file name that you want to export and click **Next**.
    f. Click **Finish**.
11. On a system on which a Directory Server instance is running, open the Directory Server key database file by using the GSKit key management application, **`ikeyman`**.
12. Add the exported certificate as a signer in the server key database.

# Known limitations and general troubleshooting

Use the list of known limitations and guidelines for general troubleshooting to identify and resolve problems that are related IBM Security Directory Suite.

## Known limitations

Use the descriptions of the known limitations section to identify issues in Directory Server and work around the problems.

### Maximum distinguished name (DN) length

The maximum length of the DN or distinguished name is 1000 characters.

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

You can run a command that specifies an option more than once. If an option is specified more than once, the option entered last is used.

For example, if you enter the following command, the `-I inst1` option is ignored and the `-I inst2` option is used.

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

### Invalid data entered on command-line utilities

A known limitation is that some types of invalid data that are entered on command-line utilities do not produce an error.

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

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

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

### No locking mechanism for conflicting commands

No locking mechanism exists to prevent conflicting commands from running at the same time for the same directory instance.

For example, you can run a command to configure a database and drop the database at the same time.

### Unable to drop database

On Windows systems, you might not be able to drop the database immediately after you stop a Directory Server instance.

This problem occurs in a scenario where all of the following conditions are true:

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

The **Instance Creation Tool** and the **idsidrop** and **idsucfgdb** commands are able to unconfigure the database but fail to drop it if all the listed conditions are satisfied. If you encounter this problem, you can manually delete the database directory after you run the **idsidrop** or **idsucfgdb** commands. Alternatively, wait at least 2 minutes after you stop the server, and then drop the database.

## Partial replication

Partial replication is a replication feature that replicates only the specified entries and a subset of attributes for the specified entries within a subtree. The entries and attributes that are to be replicated are specified by the LDAP administrator. Using partial replication, an administrator can enhance the replication bandwidth according to deployment requirements.

For instance, an administrator might choose the entries of the object class person with **cn**, **sn**, and **userPassword** attributes to be replicated and description attribute not to be replicated.

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

**Creating missing parent entries on the consumer**

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

**Scenario**

Sometimes the method to generate and submit a request to add a missing entry is recursive. The end condition is either a successful addition of all missing ancestors in the chain or a failure. The failure might occur during addition of any of the missing ancestors (for any reason other than NO_SUCH_OBJECT). If there is a failure, the change cannot be replicated and administrator intervention is required.

**Workaround**

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

**Modification in replication filter**

**Scenario**

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

**Workaround**

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

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

## Replication is not initiated

In a replication environment, if a supplier uses a password encryption setting that is not supported by the consumer, then replication is not initiated.

Also, the supplier logs a message and sets the replication state to "error xxxx" where xxxx is the ID of the message that describes the problem.

## Alias dereferencing does not work

Alias dereferencing might not work when persistent search is run on a server with no alias entries.

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

## Operation times out

Suppose that both proxy and back-end servers are configured to use PKCS#11 mode. They are required to communicate with a remote nCipher cryptographic hardware for SSL operation. In this scenario, the operation times out. To increase the operation timeout duration, you must increase the number of times that a Proxy Server must try to attempt to establish a connection.

You must increase the number of times that a Proxy Server tries to establish a connection because:

```
the total time for which a Proxy Server waits to establish a connection =
maximum time for which proxy waits to establish connection *
number of retries by a Proxy Server to establish a connection
```

To increase the number of times that the Proxy Server must try again, export the environment variable, *SERVER_ATTEMPT_TIME*, with the required count. Set the count for trying again to greater than 12, if the cryptographic hardware used for SSL operation is at a remote location.

## Instance stops when nCipher cryptographic hardware client is restarted

Directory Server instance stops when nCipher cryptographic hardware client is restarted.

**Scenario**

The following steps describe the situation in which a directory server instance might stop.

1. Start a Directory Server instance. The instance is configured over SSL with server client authorization to use PKCS#11 in key storage and accelerator mode.
2. Perform search operation by using an LDAP client in SSL mode.
3. Restart the cryptographic hardware used.
4. Perform search operation by using an LDAP client in SSL mode.

**Reason**

You must not restart the cryptographic hardware if the instance uses PKCS#11 in key storage or accelerator mode. If the cryptographic hardware that is used by a server instance for cryptographic operations is reset, then the instance stops logging appropriate messages in trace file.

## Error with `idsldapdiff` tool query

When you query an entry of large size by using the **idsldapdiff** tool, an error might occur.

The IBM SDK Java Technology Edition implementation of the **idsldapdiff** tool has a limitation. Because of this limitation, it is unable to handle entries on the Directory Server that are more than 50 MB in size. As a result, the tool might throw an Out of Memory exception when it deals with entries with more than 50 MB in size.

## Operations error during null base search

Operations error is displayed when null base search is run against a Proxy Server.

Proxy Server does not support null base search and gives an operations error if null base search is fired against it.

## User account gets locked

When the **pwdLockout** attribute is set to `true`, the user account might get locked even if the number of invalid bind attempts is less than the **pwdMaxFailure** value.

A user account might get locked when all the invalid bind attempts are made within a specified time interval that is set in the **pwdFailureCountInterval** attribute. For example, consider the following attributes are set to:

```
ibm-pwdPolicyStartTime=20070217044605Z
pwdInHistory=0
pwdCheckSyntax=1
pwdGraceLoginLimit=0
pwdLockoutDuration=0
pwdMaxFailure=3
pwdFailureCountInterval=0
passwordMaxRepeatedChars=0
pwdMaxAge=99
pwdMinAge=0
pwdExpireWarning=0
pwdMinLength=5
passwordMinAlphaChars=0
passwordMinOtherChars=0
passwordMinDiffChars=0
ibm-pwdPolicy=true
pwdLockout=false
pwdAllowUserChange=true
pwdMustChange=false
pwdSafeModify=false
ibm-pwdGroupAndIndividualEnabled=true
```

With this setting, if a user makes three invalid bind attempts, the user can still continue with bind attempts because the **pwdLockout** attribute is set to `false`.

However, **pwdFailureTime** is registered even when **pwdLockout** is `false`. Therefore, if a user does three invalid bind attempts with `pwdLockout=false` or `pwdMaxFailure=0`, **pwdFailureTime** logs one recent time stamp out of the consecutive authentication failures.

Set the **pwdLockout** attribute to `true`:

```
# idsldapmodify -D cn=RDN_value -w password
-p port_number -h host_name
dn:cn=pwdpolicy,cn=ibmpolicies
pwdLockout:true
```

Now, when the **pwdLockout** attribute is set to `true`, another two invalid bind attempts causes lockout of the user account. The lockout occurs because the invalid bind attempts made when `pwdLockout=false` is also taken into account according to the number of values in the **pwdFailureTime** attribute that are younger than **pwdFailureCountInterval**.

## Possible memory leak with PKCS#11 support configured

When you configure a Directory Server over SSL to use PKCS#11 SYMMETRIC acceleration support, there are chances for memory leak.

Consider a scenario where a Directory Server over SSL is configured to use PKCS#11 SYMMETRIC acceleration support. This configuration is for cryptographic operations by using nFast cryptographic library. In this scenario, memory leak is observed during operations.

**Note:** nFast cryptographic library is a third-party library. It is used for PKCS#11 support that is provided by IBM Security Directory Suite.

## Multivalued attributes in a virtual list view search

Duplicate entries might be returned in a virtual list view search if the sort key is a multivalued attribute.

**Explanation**

In virtual list view searches, the search filter resolutions are done in the database. The entire result set is not read from the database at one time. However, in a normal search operation a list of EIDs is

maintained in the memory. It ensures that duplicate entries are not returned to clients, even if DB2 returns duplicate EIDs.

Because the entire result set (list of EIDs) is not read into the memory, the constraint of identifying and preventing duplicates exists. Suppose that a virtual list view search is done with a primary sort key attribute that has multiple values. Then, the entries that are returned might not be in sorted order. Additionally, duplicate entries might also be returned.

**Example**

Consider a Directory Server with the following data set:

*Table 7. Entries and multivalued attribute of the entries*

| EID | Values of the cn attribute |
|---|---|
| 1 | A, Y |
| 2 | C, J |
| 3 | E |

In a normal search with `cn` as the sort key, the entries are returned in the following order: 1, 2, 3. However, the search filter resolution for the DB2 query returns EIDs in the following order: 1, 2, 3, 2, 1, based on the values of `cn`. In this case, the duplication is prevented by maintaining the list of EIDs at the server end.

In a virtual list view search, the entire result set is not maintained in memory and therefore preventing duplication is not possible.

Consider a virtual list view search that is sent with the following values: before count = 1, after count = 1, offset = 3, and content count = 0. If the virtual list view control is applied over the DB2 result set, the entries 2, 3, 2 are returned. Here, the entry with EID=2 is returned twice. The result shows that there is a possibility of returning duplicate entries in a virtual list view search if the sort key is a multivalued attribute.

## Distributed directory environment search scope

In a distributed directory environment, only base scope search with `ibm-allMembers` is supported.

If distributed group and dynamic distributed group are enabled in the configuration file, then only base scope search with `ibm-allMembers` is supported. If `onelevel` or `subtree` scope search is attempted with `ibm-allMembers`, then an appropriate error message is logged in the `ibmslapd.log` file and `LDAP_UNWILLING_TO_PERFORM` is returned.

However, if distributed group and dynamic distributed group are disabled in the configuration file, then a search for `ibm-allMembers` is forwarded to a single back-end server. In this case, the search returns group members for all search scopes.

## Instance fails to start if system date is modified

A Directory Server instance might fail to start if the system date is modified.

Suppose that the system date is set to a previous date. For example, it might be set to one month before the date when the Directory Server instance was configured on the system. If such a significant change is made to the system date, then the Directory Server instance might fail to start. The following error messages can be seen:

```
GLPRDB001E Error code -1 from function:" SQLTables ".
GLPRDB001E Error code -1 from function:" SQLFetch ".
GLPRDB001E Error code -1 from function:" SQLFetch ".
GLPRDB001E Error code -1 from function:" SQLFetch ".
GLPSRV064E Failed to initialize be_config.
```

In this scenario, the Directory Server instance is run with server trace set to ON and the debug level is set. The following error messages can be seen in the server trace:

```
188:22:35:24 T1 retrieving SQLGetDiagRec info
188:22:35:24 T1 Error - map_rc_fnc: henv=0,hdbc=0,hstmt=10001,native
retcode = -443; state = "38553"; message = "[IBM][CLI Driver][DB2/SUN64]
SQL0443N Routine "SYSIBM.SQLTABLES" (specific name "TABLES") has returned
an error SQLSTATE with diagnostic text "SYSIBM:CLI:-727". SQLSTATE=38553"
```

These error messages can also be seen in the `db2diag.log` file.

## Format of the DN gets changed

In the configuration file, the format of the DN gets changed when a composite DN is added as suffix.

Suppose that a composite DN is added as suffix. Then, the format of the DN that gets added to the configuration file is different from the DN value that was provided. For example, a composite DN, `o=sample+c=in` gets updated in configuration files as `c=i\20 + o=sample`.

## `idsdbmaint` tool error message

The **`idsdbmaint`** tool might give an error message, which states that it is unable to estimate the database size. This error is related to the privileges of the instance owner.

When the **`idsdbmaint`** tool is run with root or administrator privileges the tool inherits those privileges and therefore, the tool is able to access a directory even if it does not have write permissions or sufficient privileges for the directory instance owner. The **`idsdbmaint`** tool attempts to estimate the directory size with the privileges of the instance owner. In this case, if the instance owner does not have sufficient privileges to run the operation, the tool gives the following error.

```
GLPDBA054E Unable to estimate the database size.
```

## Error opening `filename.cat`

An error message which states that there is an error opening `filename.cat` gets displayed when you run a Directory Server. This error is related to the language pack or locale.

If a Directory Server is set to a locale that does not have corresponding message files for that locale, then an error message `Error opening filename.cat` is displayed along with an appropriate message in English locale.

The reason for this error can be the following conditions:

- An incorrect language pack is installed on the system.
- IBM Security Directory Suite does not support that particular locale.

## The values TRUE and FALSEare not translated

The Directory Server messages do not translate the values TRUEand FALSE to the corresponding locales of the translated version. You can see the issue in the translated versions of IBM Security Directory Suite, the graphical user interface (GUI) tools, such as the **Web Administration Tool**.

## Some schema-related keywords are not translated

The values of some schema-related keywords such as syntax and matching rules are not translated. You might see the issue in the **Web Administration Tool** for the translated versions of IBM Security Directory Suite.

## Date is not displayed properly for the Russian locale

You might see the following issue in the translated version of the **Web Administration Tool** in the Russian locale: Sometimes, the date format either gets displayed in wrong format or the last character of the month name gets truncated. This issue is a limitation with the tool.

## Date and time are displayed in English in translated versions

You can see the following issue on certain panels, such as **Manage backup and restore**, in the translated versions of the **Web Administration Tool**: The date and time values that are displayed on the panels are in the English locale instead of the locale of the translated version.

## Error logo is not displayed with error messages

If you access panels on the **Web Administration Tool** when the Directory Server is in the stopped state, an error panel is displayed with error messages. However, on this error panel, the error logo is not displayed. This issue is a limitation with the **Web Administration Tool**.

## Attribute encryption in RDN of an entry

The encrypted attribute of the RDN is displayed in clear text instead of being displayed in the encrypted format. This issue is a known limitation.

When you add an entry to a Directory Server instance with an RDN that has encrypted attribute in it, the following error message is displayed:

```
GLPWDM003E An error occurred while adding entry uid=5user,uid=5user,o=ibm,c=us :
uid=5user,uid=5user,o=ibm,c=us: [LDAP: error code 34 - GLPSRV156I
Encrypted attributes are not allowed in entry distinguished names.].
```

An attribute that is already present in the RDN of an entry can be encrypted without getting any error message. When the Directory Server instance is started, the encrypted attribute of the RDN must be displayed in the encrypted format. Instead, the entry displays the RDN in clear text.

This inconsistency is a limitation in the existing design.

## LDAP search filters that exceed 4K are not supported

If an LDAP search filter exceeds the 4K limit, then the server might throw an `ldap_search:bad search filter` error. To avoid this error, you must use search filters that do not exceed the 4K limit.

An error message might also be logged in the `db2cli.log` file, which indicates a syntax error in the query sent to DB2. For example, the error message that is logged in the `db2cli.log` can be of the following format:

```
12/04/07 10:38:24 native retcode = -104; state = "42601"; message =
"[IBM][CLI Driver][DB2/6000] SQL0104N An unexpected token
"END-OF-STATEMENT" was found following ".ORGANIZATIONDN WHER". Expected
tokens may include:")". SQLSTATE=42601
```

Also, the `ibmslapd.log` file might contain the following error:

```
12/04/07 10:37:44 AM GLPRDB001E Error code -1 from function:" SQLExecute " .
```

To avoid these errors, you must use search filters that do not exceed the 4K limit.

## Error during creation of a Directory Server instance from an existing instance

If the version of DB2 on the source and target server are different, the **idsideploy** tool displays an error when you create a Directory Server instance from an existing Directory Server instance.

During the creation of a Directory Server instance from an existing Directory Server instance, the **idsideploy** tool takes online backup of the source database, including the logs. At the target server, the database is restored with rolling forward of logs to bring the database to a consistent state. However, there is a limitation when the target database DB2 version is later than the source database DB2 version. The rolling forward of logs from a previous version to a later version of DB2 is not supported.

Therefore, when you use the **idsideploy** tool, you must use the same DB2 versions on the source and target server.

## The idsideploy tool fails to restore a database

The **idsideploy** tool might fail to restore a database if the backup location has backup images of the database.

When the **idsideploy** tool is run to create a copy of a Directory Server instance with data of an existing directory server instance, you must ensure that the directory path specified with the **-L** option does not already have backup image of the database, which the tool is attempting to restore. If a backup image is already present, then the restore operation of **idsideploy** fails.

## Creation of online backup image fails

The **idsdbback** command might fail to create an online backup image of a Directory Server instance that is created by the **idsideploy** tool.

When the **idsideploy** tool is used to create a copy of a Directory Server instance (along with database), the tool backs up the source database and restores it on the target server. During this process, all the internal database settings are also copied on to the target server as it is. The error messages that might get displayed are:

```
GLPDBB051E Failed to create path '/export/home/mybkup/back/INACTIVE_LOGS'
for logging inactive log files.
GLPDBB010E Failed to back up Directory Server instance 'inst2'.
```

One of the reasons for this error is that the target database uses the same settings as the source database. You must set the archive path for the target server instance before you do the online backup operation for the target server instance. Otherwise, the online backup might fail. To know more about the **idsdbback** and **idsideploy** commands, see *IBM Security Directory Suite Command Reference*.

## Inconsistent data when transaction updates are replicated

There is a possibility of inconsistent data on a directory server when transaction updates are replicated in an environment with failover setup.

When transactional updates are replicated by a supplier, the updates are not replicated in a transactional manner by the supplier to its consumers. In a replicated environment, the supplier replicates the transactional updates to its consumers only when the transaction is complete (committed or rolled back state). If a supplier goes offline during replication of the transactional updates, it is possible that only a part of the update is replicated to its consumers. In this case, when the supplier is brought online the remaining updates that are in its replication queue is replicated automatically.

However, in a replication environment with failover, if the primary master fails during replication of updates, the Proxy Server fails over to the peer server. The data might not be entirely consistent because it is possible for the peer server to not get all the updates made to master.

### Directory Server instance creation fails

IBM Security Directory Suite might fail to create a Directory Server instance.

On AIX, Solaris, and Linux systems, IBM Security Directory Suite might fail to create a Directory Server instance because of one of the following reasons:

- Not enough disk space in the `/home` or `/export/home` directory
- The root user might not have write permission on the `/home` or `/export/home` directory

### Unable to log on to a system

When migrated users use the LDAP operating system authentication mechanism, they might not be able to log on to the system. Follow the steps to work around this limitation.

The Directory Server does not support the password encryption mechanism that UNIX supports. Hence, the migrated users might not be able to log on to the system by using LDAP operating system authentication mechanism.

The Directory Server supports CRYPT and MD5 encryption schemes. However, the UNIX system uses a mix of MD5 and CRYPT password encryption scheme, which Directory Server does not support.

You can use one of the following workaround for this problem:

- LDAP administrator can reset the user password for the migrated users on the LDAP system.
- Create new users on the LDAP system for LDAP - operating system authentication.

### Accessibility tool is unable to read messages in the Configuration Tool

The Accessibility tool, JAWS, is not able to read the message that is displayed on two dialog boxes of the **Configuration Tool**, which is a limitation.

The JAWS tool is unable to read the message that is displayed on two dialog boxes because of the limitation in the design that implements messages. The following messages that are displayed on the dialog boxes are not read by the JAWS tool:

- Configuration of the instance was changed, causing this task to become invalid. Would you like to dispose this task?
- Are you sure you want to close this window?

## General troubleshooting

Use the workaround and solutions to resolve general issues in IBM Security Directory Suite.

### IBM Installation Manager generates an error when the GSKit repository contains multiple installable

To provide secure communication mechanism, IBM Security Directory Suite uses GSKit. IBM Security Directory Suite requires both GSKit SSL package and GSKit crypt package to be installed on the computer.

For the installation of GSKit with IBM Installation Manager, you must provide a path that contains GSKit installable. If the path contains installable for multiple GSKit versions, then IBM Installation Manager generates an error when installing GSKit.

The reason for the error is that GSKit versions of the same base version, for example, version 8, are installed at the same location by default. IBM Installation Manager might not be able to sequence the order of versions and install an appropriate version.

To avoid such problems, you must store SSL and crypt packages for a GSKit version in a directory. You can use the same directory to store both 32-bit and 64-bit GSKit packages of a version if they contain unique file names.

## Instance owner unable to access core file

The instance owner is sometimes unable to access the core file for a core file that is produced during server initialization. Follow the steps to work around this issue.

If the root user starts the server, a core file might be produced early during initialization of the server. The core file might not be accessible to the instance owner user. Instead, the root user has access to the core file.

If this error occurs, the root user can manually set the core file's ownership to the instance owner user if required.

This problem occurs only on AIX, Linux, and Solaris operating systems.

## Key labels do not match

If the key labels in the `.kdb` file and `ibmslapd.conf` file do not match, follow the steps to resolve this error.

If the key label in the SSL key database certificate does not match the key label in the Directory Server configuration file (`ibmslapd.conf`), the following error occurs:

```
The default SSL key database certificate is incorrect in file
c:/keytabs/pd_ldapkey.kdb.
```

Check the key label in the configuration file and the SSL key database certificate. If they do not match, create a self-signed SSL key database certificate that matches the key label in the configuration file. For more information about how to create a self-signed key database certificate, see the *Administering* section in the IBM Security Directory Suite documentation.

## GSKit certificate error

If the GSKit fails with an error when you try to import a signer or personal certificate, follow the steps to troubleshoot and resolve this error.

When you import a signer or personal certificate from an external certificate authority (CA) such as Entrust, the GSKit might fail with the following error:

```
An error occurred while receiving the certificate from the given file.
```

The problem might occur because certificate returned from Entrust is a chain certificate, not a root certificate. You must have a root certificate to start a certificate chain. A chain certificate cannot start a certificate chain.

If you do not already have a root certificate, the following method is one way to obtain the root certificate.

An example of a root certificate is GTE Cybertrust, which is included in Internet Explorer (IE). However, it is not included by default in the GSKit kdb database. To obtain this certificate, you must do the following steps:

1. Export one of the GTE Cybertrust certificates (there are 3) from Internet Explorer as Base64 encoded.
2. Add the certificate as a trusted root certificate.

   **Note:** To use the GSKit option to set a certificate as a trusted root, the certificate must be self-signed.
3. Add the chain CA certificate from Entrust.
4. Receive the SSL certificate from Entrust.

## Server instance fails to start because of incorrect file permissions

The server instance might fail to start because of incorrect file permissions. You must ensure that the file permissions are readable by the user ID, **idsldap**.

On AIX, Linux, and Solaris systems, file permissions are frequently altered inadvertently by the actions of copying or editing a key database file. Because these actions are generally done as the user ID `root`, file permissions are set for the user `root`. For the Directory Server instance to use this file, you must change the file permissions so that it is readable by the user ID **idsldap**. Otherwise, the Directory Server instance fails to start.

```
chown idsldap:idsldap mykeyring.*
```

## Server instance fails to start because host name is incorrect

The server instance fails to start because the localhost host name is not set correctly. You must ensure that the host name meets the specified requirements.

The localhost host name must correspond to the local loopback address of 127.0.0.1. If the localhost is renamed or the TCP/IP address changes, the Directory Server instance does not start.

## Server instance cannot be started except by instance owner

If a user other than the instance owner cannot start a server instance, you must verify the group and access rights of the user.

On AIX, Linux, and Solaris systems, a user other than the Directory Server instance owner might not be able to start the Directory Server instance. The user must meet the following requirements to start the Directory Server instance.

- The user who is attempting to start the Directory Server instance is a member of the primary group of the Directory Server instance owner.
- The Directory Server instance owner's primary group has Write access to the location where the database was created.

For more information about users and groups, see the *Installing* section in the IBM Security Directory Suite documentation.

## DSML file client produces error

When a user tries to connect to an LDAP server that does not use SSL, the DSML file client might give an error. This error is not serious.

The DSML file client produces the following error when the DSML file client is set up to communicate by using SSL. The error occurs when a user tries to connect to an LDAP server that does not use SSL:

```
SSL IS ON
javax.naming.CommunicationException: 9.182.21.228:389. Root exception is javax.
net.ssl.SSLProtocolException: end of file
at com.ibm.jsse.bd.a(Unknown Source)
at com.ibm.jsse.b.a(Unknown Source)
at com.ibm.jsse.b.write(Unknown Source)
at com.sun.jndi.ldap.Connection.<init>(Connection.java:226)
at com.sun.jndi.ldap.LdapClient.<init>(LdapClient.java:127)
at com.sun.jndi.ldap.LdapCtx.connect(LdapCtx.java:2398)
at com.sun.jndi.ldap.LdapCtx.<init>(LdapCtx.java:258)
at com.sun.jndi.ldap.LdapCtxFactory.getInitialContext(LdapCtxFactory.java:91)
at javax.naming.spi.NamingManager.getInitialContext(NamingManager.java:674)
at javax.naming.InitialContext.getDefaultInitCtx(InitialContext.java:255)
at javax.naming.InitialContext.init(InitialContext.java:231)
at javax.naming.InitialContext.<init>(InitialContext.java:207)
at javax.naming.directory.InitialDirContext.<init>(InitialDirContext.java:92)
at com.ibm.ldap.dsml.DsmlRequest.processRequests(DsmlRequest.java:767)
at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:253)
at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:402)
at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:373)
```

```
at com.ibm.ldap.dsml.DsmlServer.processDsmlRequest(DsmlServer.java:296)
at com.ibm.ldap.dsmlClient.DsmlFileClient.main(DsmlFileClient.java:203)
```

The error is not serious and the output XML file is generated.

## Non-default log files need valid path

If you want to store your log files in a nondefault path, you must ensure that the path exists and is valid.

You must create the directory before you can configure the log files.

## Null searches retrieve entries of deleted suffixes

If you deleted a suffix without first removing the entries of the suffix from the database, those entries are returned by the null search. The entries are returned even though the suffix no longer exists.

A null search (`ldapsearch -s sub –b "" objectclass=*`) returns all the entries that are found in the database.

## Error occurs with the `idsldapsearch` command

The **`idsldapsearch`** command with `-h` option gives an error with the DIGEST-MD5 mechanism. Follow the steps to resolve this error.

The DIGEST-MD5 SASL bind mechanism requires the client to be able to resolve the fully qualified host name of the server. If the client cannot resolve the server's fully qualified host name, the bind fails with an LDAP_PROTOCOL_ERROR. To correctly resolve the host name, you might be required to make system changes or make DNS configuration changes, such as enabling reverse DNS mapping.

For example, AIX, Linux, and Solaris systems have lines in the `/etc/hosts` file with the syntax:

```
IP address fully qualified distinguished name alias
```

This syntax is used to define the local host name to the IP address mappings.

If the syntax is something like:

```
127.0.0.1 localhost
```

When localhost is resolved, it is seen as the fully qualified distinguished name of the system, which causes DIGEST-MD5 to fail.

For the DIGEST-MD5 mechanism to work correctly, the syntax must be similar to the following syntax:

```
127.0.0.1 ldap.myserver.mycompany.com localhost
```

The syntax of the line is now such that `ldap.myserver.mycompany.com` is a valid fully qualified distinguished name for the localhost system.

## Server behavior when language tags are disabled

To avoid potential problems and unexpected behavior after you enable language tags, you must not disable the language tags.

After you enable the language tag feature, if you associate language tags with the attributes of an entry, the server returns the entry with language tags. This behavior occurs even if you later disable the language tag feature. The behavior of the server might not be what the application is expecting. Hence, to avoid potential problems, do not disable the language tag feature after it is enabled.

## Key database certificate

You must create the key database certificate before you set up SSL.

Before you set up SSL communications on your server, you must use the GSKit utility, **ikeyman**, to create the necessary certificates. For more information about **ikeyman**, see the *Administering* section in the IBM Security Directory Suite documentation.

## idsbulkload hangs during parsing phase

If the **idsbulkload** seems like it is hanging during the parsing phase, you can resolve this issue by changing some variable values.

The **idsbulkload** utility has special code to handle nested groups, and the extra processing takes time.

For example, if an LDIF file contains 50,000 nested groups with 100 member groups in each of the nested groups, **idsbulkload** might need about 1 to 2 seconds to process each one of the nested groups during the parsing phase.

In this case, **idsbulkload** seems like it is hanging before it shows any progress.

An environment variable, *BULKLOAD_REPORT_CHUNK*, can be used to increase the frequency of progress reporting.

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

- On AIX, Linux, and Solaris systems: `export BULKLOAD_REPORT_CHUNK=100`
- On Windows systems: `set BULKLOAD_REPORT_CHUNK=100`

**idsbulkload** then reports parsing progress at 100 entry interval. For example:

```
...
GLPBLK061I Parsing entries ...
GPBLK004I 100 entries parsed successfully out of 100 attempts.
LPBLK004I 200 entries parsed successfully out of 200 attempts.
..
```

## Size of log file exceeds the system file size limit

A Directory Server might fail if the size of any log file exceeds the system file size limit. This failure typically occurs when tracing is enabled on the server.

## Directory Server fails to start after running bulkload

When you run ldap operations after you run **bulkload**, the Directory Server fails to start or shows an error. Follow the steps to troubleshoot and resolve this issue.

After performing **bulkload**, if the directory server fails to start or displays error when performing LDAP operations, it could be because of one of the following reasons:

- Check the log file, db2diag.log, if there is an error that states `ACCESS TABLE WHEN IN RESTRICTED STATE`. This means that loading data or **bulkload** was not complete or was unsuccessful.
- The table is in the “Load Pending” or “Locked” state. A previous LOAD attempt on the table might have resulted in failure. Accessing the table is not allowed until the LOAD operation is restarted or terminated.

Consider the following options to rectify the problem:

- Stop or restart the failed LOAD operation on the table by issuing LOAD with the TERMINATE or RESTART option.
- Check if the `bulkload_status` file is present. This file is created in the home directory of the instance. If this file is present, it means that **bulkload** was unsuccessful. Check the file for errors and rectify it, and try running the bulk load utility again.

## Unable to open a new connection for an LDAP client

You might be unable to open a new connection for an LDAP client to connect to a Directory Server instance. This issue is specific to the Directory Server instance that is running on a Linux or Solaris operating system. You can work around this restriction by increasing the limit on open file descriptors.

On Linux and Solaris operating system, there is a limit on the maximum number of file descriptors that can be opened by a process. The default value of the maximum number for open file descriptors is 1024 for Linux operating systems and 256 for Solaris operating systems.

A Directory Server instance uses 15 file descriptors for logging messages. So on Linux, an instance stops accepting new connections after 1009, that is, 1024 – 15 concurrent client connects. Whereas on Solaris, an instance stops accepting new connections after 241, that is, 256 – 15 concurrent client connects. If an error is encountered when new connections are opened, an appropriate message is logged. This error does not affect any existing connections; only new LDAP clients fail to connect to the Directory Server.

To increase the maximum open file descriptors, user must issue the following command and restart the server from the same command prompt.

```
#ulimit –Hn number of connections
```

**Note:** The performance with a high number of concurrent client connections depends on the hardware and the operations that are run. With thousands of concurrent client connections that are sending operations simultaneously, the performance of the Directory Server might decrease.

## Error occurs when you deploy with `idsideploy` tool

When you deploy a replica or a peer in a replication environment by using the **idsideploy** tool, an error might occur. You can resolve this issue by ensuring that there is only one replication subentry.

When you deploy a replica or a peer in a replication environment by using the **idsideploy** tool, if the tool detects more than one replication subentry that contains the same **serverID** value for the attribute **ibm-replicaServerId** with the attribute **ibm-replicationServerIsMaster** set to `true`, the tool gives an error.

For any replication context, multiple replication subentries are not required, only one replication subentry is required. For example, if the entries are made as shown in the following example, **idsideploy** fails.

```
dn: ibm-replicaServerId=Peer1,ibm-replicaGroup=default, ou=ounit1, o=sample
objectclass: top
objectclass: ibm-replicaSubentry
ibm-replicaServerId: Peer1
ibm-replicationServerIsMaster: true
cn: Peer1
description: Peer1

dn: cn=Peer1_entry,ibm-replicaGroup=default, ou=ounit1, o=sample
objectclass: top
objectclass: ibm-replicaSubentry
ibm-replicaServerId: Peer1
ibm-replicationServerIsMaster: true
cn: Peer1_entry
description: Peer1
```

To rectify the problem in this example, you must create only one entry.

## Error occurs because environment variable values contain spaces

The **idssnmp** and **idslogmgmt** tools might give errors if the environment variable values contain spaces. Ensure that you do not use spaces in the value of the environment variables.

If you installed IBM Security Directory Integrator in a different location other that the default location, set the following environment variable:

- For the Log management (**idslogmgmt**) tool and SNMP (**idssnmp**) tools function correctly, you must explicitly set the *IDS_LDAP_TDI_HOME* environment variable to point to the directory where you installed IBM Security Directory Integrator.

The value that you set for the environment variable *IDS_LDAP_TDI_HOME* must not have space or double quotation marks, otherwise the tools do not work properly. On Windows, the tools work properly when tilde, "~" (that is, short path or path with no spaces) is used.

## The `idscfgdb` command fails with error code GLPCTL028E

The **idscfgdb** command might fail with error code GLPCTL028E while it is creating a database. You might be required to tune the kernel parameters to resolve this issue.

On AIX, Linux, and Solaris systems, the **idscfgdb** command might fail while it is creating a database.

An example of the `db2cli.log` file with the information logged:

```
retcode = 1478; state = "01626"; message = "SQL1478W
The defined buffer pools could not be started.
Instead, one small buffer pool for each page size supported by DB2 has been started.
SQLSTATE=01626
```

An example of the `db2diag.log` file with the information logged:

```
MESSAGE : ZRC=0x850F0005=-2062614523=SQLO_NOSEG
          "No Storage Available for allocation"
          DIA8305C Memory allocation failure occurred.
DATA #1 :
Unable to attach 3 segments totalling 2478440448 bytes starting at address
0x0000000000000000. One possible cause may be an improper setting for the
shmmax Linux kernel tuneable.
```

## Problem with monitoring server instances on a Solaris system

On a Solaris system, you might face a problem with monitoring Directory Server instances with an SNMP agent. The problem might occur with an SNMP agent that tries to log on with SSH from IBM Security Directory Integrator.

You must start an `rsh` session on the Solaris system and then try logging with `rsh` on to the Solaris system. After you log on to the Solaris system, you can monitor directory server instances by using an SNMP agent.

## The `idsdbrestore` utility displays error messages

The **idsdbrestore** utility displays error messages if the `ldapdb.properties` file is modified. You must replace the `ldapdb.properties` file to resolve this issue.

The **idsdbrestore** utility refers to the `etc/ldapdb.properties` file in the IBM Security Directory Suite installation location and not the instance-specific `ldapdb.properties` file in the *install-home*`/idsslapd-`*instance-name*`/etc` directory during a restore operation.

If a user updated or modified the **currentDB2InstallPath** parameter in the `ldapdb.properties` file to a different DB2 installation path or to a different DB2 major version after the Directory Server instance creation, error messages are displayed when you run a restore operation.

To resolve this problem, user can temporarily copy the *install-home*`/idsslapd-`*instance-name*`/etc/ldapdb.properties` file to the `etc` subdirectory in the IBM Security Directory Suite installation location before you run a restore request with the **idsdbrestore** utility. After **idsdbrestore** completes the request, restore the original `ldapdb.properties` file.

## File path causes backup and restore to fail

The backup and restore operations with the **Configuration Tool** might not function because the file path is not valid.

When you enter paths on the graphical user interface (GUI) tools, ensure that the path specified can be represented on the system. The file path string must be representable in the system's local code page as the GUI translates the Unicode input to local code page. For example, if the Unicode input for the path contains Chinese characters on a system with French locale, the translated file path is not valid.

## The warning message GLPSRV147W is displayed

In an IBM Security Directory Suite environment, a warning message with the message code GLPSRV147W might be displayed. This message might be because of the default value of write timeout that is set to 10 seconds.

If you see this error frequently for your IBM Security Directory Suite environment, you must consider increasing the write timeout value by modifying the **ibm-slapdWriteTimeout** attribute under the entry DN cn=Connection Management, cn=Front End, cn=Configuration.

You can either use the **Web Administration Tool** or the **ldapmodify** command to change the value of **ibm-slapdWriteTimeout**. To change the value, issue the **ldapmodify** command of the following format:

```
#idsldapmodify -D adminDN -w password -i filename
```

where *filename* contains:

```
dn: cn=Connection Management,cn=Front End, cn=Configuration
changetype: modify
replace: ibm-slapdWriteTimeout
ibm-slapdWriteTimeout: 120
```

# Appendix A. Support information

To obtain support for IBM products, you can use one or more of the several options such as, knowledge bases and the IBM Knowledge Center.

## 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.

### IBM Knowlege Center

IBM provides extensive online documentation in the IBM Knowledge Center. You can use the search function of the IBM Knowledge 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 IBM Knowledge Center, search the Internet for the latest and most complete information.

To search multiple Internet resources for your product, expand the product folder in the navigation frame to the left and select **Web search**. You can search various resources that include:

- IBM technotes
- IBM downloads
- IBM Redbooks®
- IBM developerWorks
- Forums and newsgroups
- Google

## Product fixes

A product fix might be available to resolve your problem. You can determine what fixes are available for your IBM software product by checking the product support website.

The following instructions can help you identify the product fix that might help you resolve your problem:

1. Go to the IBM Software Support website (http://www.ibm.com/software/support).
2. Under **Products A - Z**, select your product name. A product-specific support site is opened.
3. Under **Self help**, follow the link to **All Updates**, where you can find a list of fixes, fix packs, and other service updates for your product. For tips on refining your search, click **Search tips**.
4. Click the name of a fix to read the description and optionally download the fix.

To receive weekly email notifications about fixes and other news about IBM products, follow these steps:

1. From the support page for any IBM product, click **My support** in the upper-right corner of the page.
2. If you are already registered, skip to the next step. If you are not already registered, click register in the upper-right corner of the support page to establish your user ID and password.
3. Sign in to **My support**.
4. On the My support page, click **Edit profiles** in the left navigation pane, and scroll to **Select Mail Preferences**. Select a product family and check the appropriate boxes for the type of information you want.
5. Click **Submit**.

6. For email notification for other products, repeat Steps 4 and 5.

For more information about types of fixes, see the *Software Support Handbook* (http://techsupport.services.ibm.com/guides/handbook.html).

# Contact IBM Software Support

IBM Software Support assists with product defects. Before your contact IBM Software Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM.

The type of software maintenance contract that you need depends on the type of product you have:

- For IBM distributed software products, enroll in Passport Advantage®. These products include, but are not limited to, Security Systems, DB2, WebSphere, Lotus®, and Rational® products that run on Windows, AIX, Linux, and Solaris operating systems. You can enroll in one of the following ways:

  **Online**
  Go to the Passport Advantage web page (http://www-01.ibm.com/software/passportadvantage/) and click **How to Enroll** .

  **By phone**
  For the phone number to call in your country, go to the IBM Software Support website (http://techsupport.services.ibm.com/guides/contacts.html) and click the name of your geographic region.

- For IBM eServer™ software products, you can purchase a software maintenance agreement by working directly with an IBM marketing representative or an IBM Business Partner. These products include, but are not limited to, DB2 and WebSphere products that run in System z®, System p, and System i® environments).

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. From other countries, go to the contacts page of the IBM Software Support Handbook on the web (http://techsupport.services.ibm.com/guides/contacts.html). Click the name of your geographic region for phone numbers of people who provide support for your location.

Follow the steps to contact IBM Software Support.

## Determine the business impact of your problem

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you must understand and assess the business impact of the problem you are reporting.

Use the following criteria to determine the business impact of your problem:

*Table 8. Severity level and their description*

| **Severity** | **Business impact** |
|---|---|
| Severity 1 | Critical business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution. |
| Severity 2 | Significant business impact: The program is usable but is severely limited. |
| Severity 3 | Some business impact: The program is usable with less significant features (not critical to operations) unavailable. |
| Severity 4 | Minimal business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem is implemented. |

## Describe your problem and gather background information

When you explain a problem to IBM, you must be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? (For example, hardware, operating system, and networking software.)
- Are you currently using a workaround for this problem? If so, be prepared to explain it when you report the problem.

## Submit your problem to IBM Software Support

You can submit your problem to IBM Software Support either online or by phone.

You can submit your problem in one of two ways:

**Online**

Go to the "Submit and track problems" page on the IBM Software Support site (http://www.ibm.com/software/support/probsub.html). Enter your information into the appropriate problem submission tool.

**By phone**

For the phone number to call in your country, go to the contacts page of the IBM Software Support Handbook on the web (`techsupport.services.ibm.com/guides/contacts.html`) and click the name of your geographic region.

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround for you to implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the IBM product support web pages daily. Other users who experience the same problem can also benefit from the same resolutions.

# Index

# Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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

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

*IBM Director of Licensing*
*IBM Corporation*
*North Castle Drive, MD-NC119*
*Armonk, NY 10504-1785*
*US*

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

*Intellectual Property Licensing*
*Legal and Intellectual Property Law*
*IBM Japan Ltd.*
*19-21, Nihonbashi-Hakozakicho, Chuo-ku*
*Tokyo 103-8510, Japan*

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 jurisdictions 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 publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

IBM may use or distribute any of the information you provide 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 Director of Licensing*
*IBM Corporation*
*North Castle Drive, MD-NC119*
*Armonk, NY 10504-1785*
*US*

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.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

The client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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-IBMproducts. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:
© (your company name) (year).
Portions of this code are derived from IBM Corp. Sample Programs.
© Copyright IBM Corp. _enter the year or years_.

## Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

## Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

## Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

## Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

## Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

## Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.