Solving Common Problems

LDAP is Not Replicating

Symptom

Many OpenLDAP installations have multiple data centers, for example DC-1 and DC-2. A search of the users in DC-1 does not match a list of users in DC-2.

Error Messages

No errors appear, the search simply does not show the list of users that should have been replicated across all OpenLDAP servers.

Possible Causes

Typically the cause of this issue is a misconfigured OpenLDAP replication configuration, not the installation itself. Also, replication may break if the network between the OpenLDAP servers does not allow traffic on port 389 or 636 (or whatever custom port slapd is configured to use).

Diagnosis

Use the following steps to diagnose the problem:

- Check if ldapsearch returns data from each OpenLDAP server:

```
ldapsearch -W -D "cn=replicator,ou=application accounts,dc=example,dc=com" -b
"dc=example,dc=com" -LLL -H ldap://<host-ip>:389
```

- Check if you can connect to each OpenLDAP node from the other OpenLDAP nodes on port 10389. If telnet is installed, use the following command:

```
telnet <OpenLDAP_Peer_IP> 389
```

- If telnet is not available, use netcat to check the connectivity as follows:

```
nc -vz <OpenLDAP_Peer_IP> 389
```

- Check the replication configuration in the following file:

  /opt/symas/etc/openldap/slapd.d/cn=config/olcDatabase={1}mdb.ldif

  The file should contain configuration similar to this:

```
olcSyncRepl: rid=001
  provider=ldap://__OTHER_LDAP_SERVER__/
  binddn="cn=replicator,ou=application accounts,dc=example,dc=com"
  bindmethod=simple
  credentials=__LDAP_PASSWORD__
  searchbase="dc=example,dc=com"
  attrs="*,+"
  type=refreshAndPersist
  retry="60 1 300 12 7200 +"
  timeout=1
```

- Also for Multi-Master Replication check the same file for the value of the olcMirrorMode attribute. It should be set to the value TRUE:

```
grep olcMirrorMode
/opt/symas/etc/openldap/slapd.d/cn=config/olcDatabase={1}mdb.ldif
```

- Check for iptables and tcp wrapper rules. Please remove any rules that do not allow the peer OpenLDAP servers to communicate with each other. Work with your network administrator to set the rules appropriately.
- Make sure the OpenLDAP system password is the same on each OpenLDAP node.
- Check for hidden characters in the ldif configuration files that are being used to configure N-Way OpenLDAP replication by running dos2unix against the ldif files that have been created to update the configuration. Typically an ldif file that has bad characters would cause the ldapmodify

command to fail to run and so replication may not be set up. Remove any bad characters and save the configuration files.

- If the problem persists, contact Symas Support for assistance with the replication issue.

Unable to Start OpenLDAP

Symptom

OpenLDAP does not start.

Error Messages

SLAPD process dead but PID file exists

Possible Causes

This issue is typically caused by a lock file that is left behind on the file system and needs to be removed.

Diagnosis

Use the following steps to diagnose this issue:

- Verify slapd is actually stopped

```
ps -aux |grep slapd
```

- If a Process ID (PID) is returned, issue a kill command to terminate the process

```
kill -9 <PID>
```

- Either way, check for an OpenLDAP slapd PID file in the following location:

```
ls /var/symas/run/slapd.pid
```

- If it exists and the process has been killed, delete the PID file and try to restart openldap.

```
rm –f ls /var/symas/run/slapd.pid
```

If the OpenLDAP slapd process starts, then skip the below steps.

- If the OpenLDAP slapd process does not start, try running slapd in debug mode and look for any errors:

```
slapd -H ldap://:389/ -F /opt/symas/etc/openldap/slapd.d -d -1
```

Errors may point to resource issues. Check memory and CPU utilization on the system.

- Check the version of OpenLDAP and upgrade if it is old. Check for the supported versions of OpenLDAP in our Supported Software document.

```
slapd -V
```

- Use strace to troubleshoot slapd process, and to provide strace output to Symas Support:

```
strace -tt -T -f -F -i -v -e read=all -s 8192 -e write=all -o /tmp/strace.out
-p <PID>
```

Server Unwilling to Perform Operation

Symptom

In Producer/Consumer Replication the Consumer server will be unwilling to replicate from its respective Producer.

Error Messages

```
SEARCH RESULT tag=101 err=53 duration=1.992ms nentries=0 text=consumer state is newer
than provider!
       do_syncrep2: rid=002 (53) Server is unwilling to perform
        slap_client_connect: URI=ldap://<master server>/
DN="cn=replicator,dc=example,dc=com" ldap_sasl_bind_s failed (-1)
```

Possible Causes

This occurs when the Consumer server has newer data than the Producer does. This is easy to do with delta-syncrepl because starting up delta-syncrepl on an empty accesslog database will create the cn=accesslog container object with a brand new timestamp. This can be avoided by starting the Producer first and then the Consumer.

Diagnosis

Use the following steps to diagnose this issue:

- First, verify the clocks on both servers are in sync.

```
date
```

- Next check the contextCSN attribute values on both servers.

```
ldapsearch -x  -H ldap://<consumer FQDN> -D dc=example,dc=com -w secret -b
dc=example,dc=com -s base contextCSN

ldapsearch -x  -H ldap://<producer FQDN> -D dc=example,dc=com -w secret -b
dc=example,dc=com -s base contextCSN
```

The values should match

- If the Consumer's contextCSN value is newer than the Producer's, then issue an modification to the Producer's primary MDB database which will create a new and updated contextCSN attribute value.
- If this does not resolve the problem, replace the Consumer's accesslog database with a copy of the Producer's accesslog database.

On the Consumer:

```
service solserver stop

rm -rf /var/symas/openldap-data/accesslog/*.mdb
```

On the Producer:

```
service solserver stop

slapcat -n 2 -l /opt/symas/etc/openldap/access.ldif
```

Copy the ldif file from the Producer to the Consumer:

```
scp /opt/symas/etc/openldap/access.ldif root@<consumer
FQDN>:/opt/symas/etc/openldap/

service solserver start
```

Import the LDIF on the Consumer and start slapd:

```
slapadd -n 2 -l /opt/symas/etc/openldap/access.ldif

service solserver start
```

Check logs to verify replication is now completing successfully.

Slow Operation Processing

Symptom

An operation is issued against an OpenLDAP Server, but the process takes a long time to complete.

Error Messages

```
Nov 19 10:23:03 core01 slapd[24753]: warning: accept(9) failed errno=24: Too many
open files
```

Possible Causes

This can occur when the number of files opened by the slapd process meet or exceed existing nofile (Maximum number of files that can be opened by a single process) limits.

Processor thread availability can also be a cause. Essential the operation must wait for a thread to free up before it can complete.

Diagnosis

Use the following steps to diagnose this issue:

- View current file descriptors in use by slapd

```
pgrep slapd
```

This will give you the Process ID for slapd

```
ls /proc/<slapd PID>/fd |wc -l
```

This will give you the number of file descriptors in use. Compare that number to the NOFILE limits

- For Non-SystemD operating systems, the limits are set in /etc/security/limits.conf

```
less /etc/security/limits.conf |grep nofile
```

The soft limit is the current limit for users or processes. It can be overridden up to, but not more than, the hard limit. The hard limit is the maximum limit set by root for any/all users or processes.

- For SystemD systems use systemctl show to see what limits are applied to the slapd process.

```
systemctl show slapd | grep NOFILE
```

These values are the soft and hard limits of how many files can be open for the slapd process.

- If the value of file descriptors for the slapd process is equal to or greater than the maximum NOFILE limit for slapd, the limit needs to be increased.
- If the value of file descriptors for the slapd process is significantly less than the NOFILE limits, check for process thread usage. First, examine the number of threads available on the server.

```
lscpu |egrep 'Thread|CPU\(s\):'
```

Multiple the Threads per core value by the number of CPU(s) for the number of available threads. Remember, the default thread setting for slapd is 16 unless otherwise set. It is possible that value is too high for the threads available.

- To reduce the thread count for slapd, modify slapd.conf or slapd.d changing the threads/olcThreads value to a number reflective the servers available thread count. Remember, for slapd.conf a solserver restart is required.
- If the threads setting in the slapd configuration is adequate for the server, check how many threads are in use by the slapd process.

```
pgrep slapd

cat /proc/<slapd PID>/status |grep Threads
```

- If the number of threads used by slapd is equal to the threads/olcThreads setting in the slapd configuration, that value may need to be increased to handle the additional workload, or additional CPUs may need to be added to the server, thus making more threads available allowing for a higher threads/olcThreads setting in slapd.conf/slapd.d.
- Finally, adding some sort of replication may be needed to distribute the load over multiple servers. Be aware than replication adds network connections which may use up available threads.