

Univention Corporate Server - Performance guide

Release 5.0

Sep 22, 2023

Contents

1	OpenLDAP and listener/notifier domain replication	2
1.1	Indexes	2
1.2	Configuration of the database back end	3
1.3	OpenLDAP ACLs	4
1.4	Univention Directory Listener	4
2	Name Service Cache Daemon (NSCD)	4
3	Performance issues during the join process	4
3.1	Samba	5
4	Local group cache	5
5	UCS management system	6
5.1	Disabling automatic search	6
5.2	Imposing a size limit for searches	6
5.3	Adjusting the limit on open file descriptors	6
5.4	Vertical performance scaling	6
6	Further services and components	6
6.1	Apache	6
6.2	SAML	7
6.3	Squid	7
6.4	BIND	7
6.5	Kernel	7
6.6	Samba	8
6.7	System statistics	8
6.8	Dovecot high-performance mode	8
6.9	UDM REST API performance scaling	8
7	Bibliography	9
	References	9

By default UCS is suitable for environments with up to 5,000 users. This document describes configuration modifications which can increase performance in larger environments.

1 OpenLDAP and listener/notifier domain replication

As a core element in the operation and administration of a UCS domain, the performance of the LDAP server plays a central role in the overall performance.

1.1 Indexes

Comparable with other database systems, OpenLDAP uses indexes about commonly requested attributes. For indexed attributes a search isn't performed through the full database contents, but over an optimized subsection.

With recent UCS versions, the indexes are occasionally expanded and automatically activated. You can deactivate the automatic activation with the UCR variable `ldap/index/autorebuild`. In this case, you must set the indexes manually to ensure that there is no loss of performance as a result.

Optimizing an LDAP index requires some thoughts before hand. It isn't recommended to add all wanted attributes to the LDAP index. The index requires maintenance and updating an index costs performance on write and sometimes also on read operations. Evaluate your frequent searches and add the frequent used search attributes to the LDAP index.

See also:

Can more indexes improve performance?¹ in the *OpenLDAP Faq-O-Matic*
for more information about indexes influence the performance in OpenLDAP

To add LDAP attributes to an index, use the following steps:

1. Stop the OpenLDAP server.
2. Add or remove the LDAP attributes in the indexes by changing the respective UCR variable.
3. Run the command `slapindex` to re-index the entries in the LDAP database.
4. Start the OpenLDAP server.

The UCR variables `ldap/index/eq` (page 3), `ldap/index/pres` (page 3), `ldap/index/sub` (page 3), and `ldap/index/approx` (page 2) control the LDAP index configuration for the OpenLDAP server. After changing one of those variables, UCR rewrites the OpenLDAP server configuration file.

Consider the following hints for your index:

- Negations don't use an index and therefore negation searches suffer performance.
- Range comparisons like `>=` and `<=` only work for non-string syntaxes like `integer` and `generalizedTime`. They use the `eq` (page 3) index for equality tests.

The following UCR variables control the LDAP index:

`ldap/index/approx`

This index tests for approximate matches. For example adding the `uid` attribute to the index, it corresponds to the LDAP filter `uid~=value` and finds all approximate objects.

See also:

Ldapwiki: ApproxMatch^{Page 3, 2}

Overview about approximate match in OpenLDAP search filters defined in **RFC 4511**³.

¹ <https://www.openldap.org/faq/data/cache/42.html>

ldap/index/eq

This index tests for equality. For example adding the `uid` attribute to the index, it corresponds to the LDAP filter `uid=value` and finds all objects with exact that `value`. LDAP uses this index as fallback for a missing presence index in `ldap/index/pres` (page 3).

ldap/index/pres

This index tests for the presence. For example adding the `uid` attribute to the index, it corresponds to the LDAP filter `uid=*` and finds all objects that have something within the `uid` attribute.

ldap/index/sub

This index runs a sub string search. For example adding the `uid` attribute to the index, it corresponds to the LDAP filter `uid=value*` and finds all objects that match with the filter including the wildcard.

To determine whether OpenLDAP uses not-indexed variables, you can activate OpenLDAP debug level `-1` and search for the string `not indexed` in the log file `/var/log/syslog`. For example:

```
$ ucr set ldap/debug/level=-1
$ systemctl restart slapd
$ grep 'not indexed' /var/log/syslog
```

1.2 Configuration of the database back end

The memory mapped database (MDB) has been used for new installations since UCS 4.0. If BDB is still in use, a migration to MDB should be performed for *amd64* systems. The database back end can be controlled via the UCR variable `ldap/database/type`. A migration can be performed as follows:

```
$ systemctl stop slapd
$ slapcat -f /etc/ldap/slapd.conf -l ldif
$ mkdir /var/lib/univention-ldap/ldap.BACKUP
$ mv /var/lib/univention-ldap/ldap/* /var/lib/univention-ldap/ldap.BACKUP
$ ucr set ldap/database/type=mdb
$ slapadd -f /etc/ldap/slapd.conf -l ldif
$ systemctl start slapd
```

By default the memory mapped database needs more I/O operations than the BDB back end. With the Univention Configuration Registry Variable `ldap/database/mdb/envflags` this behavior can be configured. The following flags can be set (multiple values are separated by spaces):

nosync

Specify that on-disk database contents should not be immediately synchronized with in memory changes. Enabling this option may improve performance at the expense of data security. In particular, if the operating system crashes before changes are flushed, some number of transactions may be lost. By default, a full data flush/sync is performed when each transaction is committed.

nometasync

Flush the data on a commit, but skip the sync of the meta page. This mode is slightly faster than doing a full sync, but can potentially lose the last committed transaction if the operating system crashes. If both `nometasync` and `nosync` are set, the `nosync` flag takes precedence.

writemap

Use a writable memory map instead of just read-only. This speeds up write operations but makes the database vulnerable to corruption in case any bugs in `slapd` cause stray writes into the memory mapped region.

mapasync

When using a writable memory map and performing flushes on each commit, use an asynchronous flush instead of a synchronous flush (the default). This option has no effect if `writemap` has not been set. It also has no effect if `nosync` is set.

² <https://ldapwiki.com/wiki/ApproxMatch>

³ <https://datatracker.ietf.org/doc/html/rfc4511.html>

nordahead

Turn off file read-ahead. Usually the OS performs read-ahead on every read request. This usually boosts read performance but can be harmful to random access read performance if the system's memory is full and the DB is larger than RAM.

1.3 OpenLDAP ACLs

Access to the information contained in the LDAP directory is controlled by access control lists (ACLs) on the server side. General information on the configuration of ACLs in UCS can be found in [Access control for the LDAP directory](#)⁴ in *Univention Corporate Server - Manual for users and administrators* [1].

Nested groups are also supported. The Univention Configuration Registry Variable `ldap/acl/nestedgroups`⁵ can be used to deactivate the nested groups function for LDAP ACLs, which will result in a speed increase for directory requests.

1.4 Univention Directory Listener

The Univention Directory Listener can perform safety checks to prevent a user name being added into a group twice. These checks add some overhead to replication and can be deactivated by setting the Univention Configuration Registry variables `listener/memberuid/skip` and `listener/uniquemember/skip` to `no`. Starting with UCS 3.1 the variables are not set and the checks are not activated any longer by default.

2 Name Service Cache Daemon (NSCD)

Name resolutions can be cached by the *Name Service Cache Daemon* (NSCD) in order to speed up frequently recurring requests for unchanged data. Thus, if a repeated request occurs, instead of querying the LDAP server, the data are simply drawn directly from the cache.

The size of the cache held by the NSCD is preconfigured for an environment with 5,000 users. If more users or hosts are created, the cache should be enlarged as otherwise it will not be possible to cache enough entries.

The following Univention Configuration Registry variables can be set:

- `nscd/hosts/size` should be at least the same as the number of all the computers entered in the DNS.
- `nscd/passwd/size` should be at least the same as the number of users.

To allow an efficient cache allocation, the value selected should always be a prime number, in case of doubt the next highest prime number should be selected.

A script can be downloaded from <https://updates.software-univention.de/download/scripts/nscdCachesize.sh> which suggests corresponding values based on the objects currently included in the system.

3 Performance issues during the join process

The size of the UCS domain can have an impact on the duration of the join process. Here is some information how to deal with such problems.

⁴ <https://docs.software-univention.de/manual/5.0/en/domain-ldap/ldap-directory.html#domain-ldap-acls>

⁵ <https://docs.software-univention.de/manual/5.0/en/appendix/variables.html#envvar-ldap-acl-nestedgroups>

3.1 Samba

One of the join scripts for samba requires that the samba connector has synchronized all domain objects into samba. This script has a timeout of 3h (from UCS 4.4-7 on). This is sufficient for normal sized environments. But in large environments this script may hit the timeout and abort the join process. To increase the timeout the Univention Configuration Registry Variable `create/spn/account/timeout` can be set prior to the join process.

The join scripts `97univention-s4-connector` and `98univention-samba4-dns` wait for the replication of the DNS record of the joining system to verify that the local Samba backed DNS server can answer requests for Active Directory related requests. By default the scripts wait for 600 seconds, but in case there are a lot of objects that need to be replicated (e.g. DNS zones) then this default may be too short. In that case the timeout can be adjusted by setting the Univention Configuration Registry Variable `join/samba/dns/replication/timeout` to a bigger value before joining.

Samba traditionally uses TDB as back end database storage, that has an internal 32 bit address space limitation. UCS supports provisioning Samba using LMDB instead, which doesn't have this strict limitation. For more information, see *Lightning Memory-Mapped Database Manager (LMDB) documentation* [2]. [KB 18014](#)⁶ describes how to migrate a productive UCS domain.

To use LMDB instead of TDB, set the corresponding Univention Configuration Registry Variable `samba/database/backend/store` to `mdb` before you install the app **Active Directory-compatible Domain Controller** in your UCS domain.

The Univention Configuration Registry Variable `samba/database/backend/store/size` defines the current maximal size of the individual back end database store files and has the default value of 8GB. Since there is one back end storage file per Active Directory naming context, this amounts to a total of 40GiB. Take care that the storage can accommodate this amount of space.

If required, you can increase the value monotonically. After changing the Univention Configuration Registry Variable, you need to restart Samba, so that the value can take effect. You can check the number of used storage pages, 4KiB each, by running the command `mdb_stat -nef` on the individual files and calculating `Number of pages used minus Free pages`. The value at `Max pages` shows the current effective limit. The back end storage files locate in `/var/lib/samba/private/sam.ldb.d/` and have the file extension `.ldb`.

4 Local group cache

By default the group cache is regenerated every time changes are made to a group. This avoids cache effects whereby group memberships only become visible for a service after the next scheduled group cache rewrite (by default once a day and after 15 seconds of inactivity in the Univention Directory Listener). In larger environments with a lot of group changes, this function should be deactivated by setting the Univention Configuration Registry Variable `nss/group/cachefile/invalidate_on_changes`⁷ to `false`. This setting takes effect immediately and does not require a restart of the Univention Directory Listener.

When the group cache file is being generated, the script verifies whether the group members are still present in the LDAP directory. If only the Univention Management Console is used for the management of the LDAP directory, this additional check is not necessary and can be disabled by setting the Univention Configuration Registry Variable `nss/group/cachefile/check_member`⁸ to `false`.

⁶ <https://help.univention.com/t/18014>

⁷ https://docs.software-univention.de/manual/5.0/en/appendix/variables.html#envvar-nss-group-cachefile-invalidate_on_changes

⁸ https://docs.software-univention.de/manual/5.0/en/appendix/variables.html#envvar-nss-group-cachefile-check_member

5 UCS management system

5.1 Disabling automatic search

By default all objects are automatically searched for in the domain management modules of the Univention Management Console. This behavior can be disabled by setting the Univention Configuration Registry Variable `directory/manager/web/modules/autosearch` to 0.

5.2 Imposing a size limit for searches

The Univention Configuration Registry Variable `directory/manager/web/sizelimit` is used to impose an upper limit for search results. If, e.g., this variable is set to 2000 (as is the default), searching for more than 2000 users would not be performed and instead the user is asked to refine the search.

5.3 Adjusting the limit on open file descriptors

The Univention Configuration Registry Variable `umc/http/max-open-file-descriptors` is used to impose an upper limit on open file descriptors of the **univention-management-console-web-server**. The default is 65535.

5.4 Vertical performance scaling

A single Univention Management Console instance does not use multiple CPU cores by design, therefore it can be beneficial to start multiple instances. Set the following Univention Configuration Registry Variables `umc/server/processes` and `umc/http/processes` and restart the Univention Management Console:

```
$ systemctl restart apache2 \
univention-management-console-web-server \
univention-management-console-server
```

The number of instances to configure depends on the workload and the server system. As a general rule of thumb these should not be higher than the machines CPU cores. Good throughput values had resulted in tests with the following combinations:

- 6 CPU cores: `umc/http/processes=3` and `umc/server/processes=3`
- 16 CPU cores: `umc/http/processes=15` and `umc/server/processes=15`
- 32 CPU cores: `umc/http/processes=25` and `umc/server/processes=25`

Note that the number of Apache processes may also need to be increased for the customization to take effect.

6 Further services and components

6.1 Apache

In environments with many simultaneous accesses to the web server or Univention Portal and Univention Management Console, it may be advisable to increase the number of possible Apache processes or reserve processes. This can be achieved via the UCR variables `apache2/server-limit`, `apache2/start-servers`, `apache2/min-spare-servers` and `apache2/max-spare-servers`. After setting, the Apache process must be restarted via the command **systemctl restart apache2**.

Detailed information about useful values for the UCR variables can be found at [ServerLimit Directive⁹](#) and [Start-Servers Directive¹⁰](#) in *Apache HTTP Server Version 2.4 Documentation* [3].

6.2 SAML

By default, SAML assertions are valid for 300 seconds and must be renewed by clients no later than then to continue using them. In scenarios where refreshing SAML assertions at such short intervals is too expensive (for clients or servers), the lifetime of SAML assertions can be increased via the UCR variable `umc/saml/assertion-lifetime`. This can be achieved on each UCS system with the role Primary Directory Node or Backup Directory Node by executing the following commands:

```
$ ucr set umc/saml/assertion-lifetime=3600
$ cd /usr/share/univention-management-console/saml/
$ ./update_metadata --binddn $USERDN --bindpwfile $FILENAME
```

`$USERDN` has to be replaced with a valid DN of a user, that is member of the group `Domain Admins` and the file specified by `$FILENAME` has to contain the corresponding password of that user.

It should be noted that increasing the lifetime has security implications that should be carefully considered.

6.3 Squid

If the Squid proxy service is used with NTLM authentication, up to five running NTLM requests can be processed in parallel. If many proxy requests are received in parallel, the Squid user may occasionally receive an authentication error. The number of parallel NTLM authentication processes can be configured with the Univention Configuration Registry Variable `squid/ntlmauth/children`.

6.4 BIND

BIND can use two different back ends for its configuration: OpenLDAP or the internal LDB database of Samba/AD. The back end is configured via the Univention Configuration Registry Variable `dns/backend11`. On UCS Directory Nodes running Samba/AD, the back end **must not** be changed to OpenLDAP.

When using the Samba back end, a search is performed in the LDAP for every DNS request. With the OpenLDAP back end, a search is only performed in the directory service if the DNS data has changed. For this reason, using the OpenLDAP back end can reduce the load on a Samba/AD domain controller.

6.5 Kernel

In medium and larger environments the maximum number of open files allowed by the Linux kernel may be set too low by default. As each instance requires some unswappable memory in the Linux kernel, too many objects may lead to a resource depletion and denial-of-service problems in multi-user environments. Because of that the number of allowed file objects is limited by default.

The maximum number of open files can be configured on a per-user or per-group basis. The default for all users can be set through the following Univention Configuration Registry Variables:

security/limits/user/default/hard/nofile

The hard limit defines the upper limit a user can assign to a process. The default is 32768.

security/limits/user/default/soft/nofile

The soft limit defines the default settings for the processes of the user. The default is 32768.

A similar problem exists with the Inotify sub-system of the kernel, which can be used by all users and applications to monitor changes in file systems.

⁹ https://httpd.apache.org/docs/2.4/en/mod/mpm_common.html#serverlimit

¹⁰ https://httpd.apache.org/docs/2.4/en/mod/mpm_common.html#startservers

¹¹ <https://docs.softwire-univention.de/manual/5.0/en/appendix/variables.html#envvar-dns-backend>

kernel/fs/inotify/max_user_instances

The upper limit of inotify services per user ID. The default is 511.

kernel/fs/inotify/max_user_watches

The upper limit of files per user which can be watched by the inotify service. The default is 32767.

kernel/fs/inotify/max_queued_events

The upper limit of queued events per inotify instance. The default is 16384.

6.6 Samba

Samba uses its own mechanism to specify the maximum number of open files. This can be configured through the Univention Configuration Registry Variable `samba/max_open_files`. The default is 32808.

If the log file `/var/log/samba/log.smbd` contains errors like `Failed to init inotify - Too many open files`, the kernel and Samba limits should be increased and the services should be restarted.

6.7 System statistics

The log file `/var/log/univention/system-stats.log` can be checked for further performance analyses. The system status is logged every *30 minutes*. If more regular logging is required, it can be controlled via the UCR variable `system/stats/cron`¹².

6.8 Dovecot high-performance mode

Univention Corporate Server configures Dovecot to run in *High-security mode* by default. Each connection is served by a separate login process. This security has a price: for each connection at least two processes must run.

Thus installations with 10.000s of users hit operating system boundaries. For this case Dovecot offers the *High-performance mode*. To activate it, login processes are allowed to serve more than one connection. To configure this run

```
$ ucr mail/dovecot/limits/imap-login/service_count=0
```

If `client_limit=1000` and `process_limit=100` are set, only 100 login processes are started, but each serves up to 1000 connections — a total of 100.000 connections.

The cost of this is that if a login process is compromised, an attacker might read the login credentials and emails of all users this login process is serving.

To distribute the load of the login processes evenly between CPU cores, `mail/dovecot/limits/imap-login/process_min_avail` should be set to the number of CPU cores in the system.

6.9 UDM REST API performance scaling

A single Univention Directory Manager REST API instance does not use multiple CPU cores by design, therefore it can be beneficial to start multiple instances. By setting the Univention Configuration Registry Variable `directory/manager/rest/processes` the number of processes can be increased. Afterwards the Univention Directory Manager REST API needs to be restarted:

```
$ systemctl restart univention-directory-manager-rest
```

The number of instances to configure depends on the workload and the server system. As a general rule of thumb these should not be higher than the machines CPU cores. With `directory/manager/rest/processes=0` all available CPU cores are used.

¹² <https://docs.software-univention.de/manual/5.0/en/appendix/variables.html#envvar-system-stats-cron>

7 Bibliography

References

- [1] *Univention Corporate Server - Manual for users and administrators*. Univention GmbH, 2021. URL: <https://docs.software-univention.de/manual/5.0/en/>.
- [2] *Lightning Memory-Mapped Database Manager (LMDB) documentation*. Symas Corporation, 2015. URL: <http://www.lmdb.tech/doc/>.
- [3] *Apache HTTP Server Version 2.4 Documentation*. The Apache Software Foundation. URL: <https://httpd.apache.org/docs/2.4/en/>.

Index

A

apache2/max-spare-servers, 6
apache2/min-spare-servers, 6
apache2/server-limit, 6
apache2/start-servers, 6

C

create/spn/account/timeout, 5

D

directory/manager/rest/processes, 8
directory/manager/web/modules/autosearch, 6
directory/manager/web/sizelimit, 6
dns/backend, 7

E

environment variable
 apache2/max-spare-servers, 6
 apache2/min-spare-servers, 6
 apache2/server-limit, 6
 apache2/start-servers, 6
 create/spn/account/timeout, 5
 directory/manager/rest/processes, 8
 directory/manager/web/modules/autosearch, 6
 directory/manager/web/sizelimit, 6
 dns/backend, 7
 join/samba/dns/replication/timeout, 5
 kernel/fs/inotify/max_queued_events, 8
 kernel/fs/inotify/max_user_instances, 8
 kernel/fs/inotify/max_user_watches, 8
 ldap/acl/nestedgroups, 4
 ldap/database/mdb/envflags, 3
 ldap/database/type, 3
 ldap/index/approx, 2
 ldap/index/autorebuild, 2
 ldap/index/eq, 2, 3
 ldap/index/pres, 2, 3
 ldap/index/sub, 2, 3
 listener/memberuid/skip, 4
 listener/uniquemember/skip, 4
 mail/dovecot/limits/imap-login/process_min_avail, 8
 nscd/hosts/size, 4
 nscd/passwd/size, 4
 nss/group/cachefile/check_member, 5
 nss/group/cachefile/invalidate_on_changes, 5

samba/database/backend/store, 5
samba/database/backend/store/size, 5
samba/max_open_files, 8
squid/ntlmauth/children, 7
system/stats/cron, 8
umc/http/max-open-file-descriptors, 6
umc/http/processes, 6
umc/saml/assertion-lifetime, 7
umc/server/processes, 6

J

join/samba/dns/replication/timeout, 5

K

kernel/fs/inotify/max_queued_events, 8
kernel/fs/inotify/max_user_instances, 8
kernel/fs/inotify/max_user_watches, 8
Knowledge Base
 KB 18014, 5

L

ldap/acl/nestedgroups, 4
ldap/database/mdb/envflags, 3
ldap/database/type, 3
ldap/index/approx, 2
ldap/index/autorebuild, 2
ldap/index/eq, 2
ldap/index/pres, 2, 3
ldap/index/sub, 2
listener/memberuid/skip, 4
listener/uniquemember/skip, 4

M

mail/dovecot/limits/imap-login/process_min_avail, 8

N

nscd/hosts/size, 4
nscd/passwd/size, 4
nss/group/cachefile/check_member, 5
nss/group/cachefile/invalidate_on_changes, 5

R

RFC
 RFC 4511, 2

S

samba/database/backend/store, 5
samba/database/backend/store/size, 5

samba/max_open_files, 8
squid/ntlmauth/children, 7
system/stats/cron, 8

U

umc/http/max-open-file-descriptors, 6
umc/http/processes, 6
umc/saml/assertion-lifetime, 7
umc/server/processes, 6