

Univention Corporate Server



Performance guide

Table of Contents

1. Introduction	3
2. OpenLDAP and listener/notifier domain replication	4
2.1. Indexes	4
2.2. Configuration of the database backend	4
2.3. OpenLDAP ACLs	5
2.4. Univention Directory Listener	5
3. Name Service Cache Daemon (NSCD)	6
4. Local group cache	7
5. UCS management system	8
5.1. Disabling automatic search	8
5.2. Imposing a size limit for searches	8
6. Further services and components	9
6.1. Squid	9
6.2. BIND	9
6.3. Kernel	9
6.4. Samba	9
6.5. System statistics	10
6.6. Dovecot high-performance mode	10
Bibliography	11

Chapter 1. Introduction


In its default configuration, UCS is suitable for environments with up to 5,000 users.

This document describes configuration modifications which can increase performance in larger environments.

Chapter 2. 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.

2.1. Indexes

Feedback 


Comparable with other database systems, OpenLDAP runs indexes about commonly requested attributes. For indexed attributes, a search is not performed via the full database contents, but over an optimized subsection.

With newer UCS versions, the indexes are occasionally expanded and automatically activated. The automatic activation can be deactivated using the UCR variable `ldap/index/autorebuild`. In this case, the indexes should be set manually to ensure that there is no loss of performance as a result. The indexes are controlled by the UCR variables `ldap/index/eq`, `ldap/index/pres`, `ldap/index/sub` and `ldap/index/approx`. Once the variables have been changed, the OpenLDAP server must be stopped and the `slapindex` command run.

To determine whether not-indexed variables are used, 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
invoke-rc.d slapd restart
grep 'not indexed' /var/log/syslog
```

2.2. Configuration of the database backend

Feedback 

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 (not i386). The database backend can be controlled via the UCR variable `ldap/database/type`. A migration can be performed as follows:


```
/etc/init.d/slapd stop
slapcat -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 -l ldif
/etc/init.d/slapd start
```

In the default configuration the memory mapped database needs more I/O operations than the BDB backend. 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.
- `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.

2.3. OpenLDAP ACLs


Feedback 

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 the LDAP chapter of the [ucs-manual].

The default setting of the LDAP server does not allow anonymous access to the LDAP directory. If the Univention Configuration Registry variable `ldap/acl/read/anonymous` is set to `yes`, these access restrictions do not apply and the LDAP performance is increased. Access to sensitive attributes such as user passwords remains protected via separate ACLs.

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

2.4. Univention Directory Listener

Feedback 

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/uniqemember/skip` to `no`. Starting with UCS 3.1 the variables are not set and the checks are not activated any longer by default.

Chapter 3. 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.


Chapter 4. Local group cache

In the default setting, 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 (in the default setting after 15 minutes 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/invalidate_cache_on_changes` 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` to `false`.


Chapter 5. UCS management system

5.1. Disabling automatic search

Feedback 

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

Feedback 

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.


Chapter 6. Further services and components

6.1. Squid

Feedback 

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`.


6.2. BIND

Feedback 

BIND can use two different backends for its configuration: OpenLDAP or the internal LDB database of Samba 4. The backend is configured via the Univention Configuration Registry variable `dns/backend`.

When using the Samba backend, a search is performed in the LDAP for every DNS request. With the OpenLDAP backend, a search is only performed in the directory service if the DNS data has changed. For this reason, using the OpenLDAP backend can reduce the load on a Samba 4 domain controller.

6.3. Kernel

Feedback 

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.

`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.


6.4. Samba

Feedback 

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.

6.5. System statistics

Feedback 

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`.

6.6. Dovecot high-performance mode

Feedback 

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.

Bibliography

[ucs-manual] Univention GmbH. 2018. *Univention Corporate Server - Manual for users and administrators*. <https://docs.software-univention.de/manual-4.3.html>.