



linux for your business

Peter H. Ganten

ganten@univention.de

Testing Samba for Bigger Environments -

Samba / Linux / OpenLDAP at the

german federal parliament

Agenda

1. Advertisement for the speakers company (very short, I promise)
2. History and general conditions of the MigOS project at the german federal parliament
3. The crash
4. The race towards a functional setup
5. Behavior and performance of the final setup
6. Conclusion and questions



Univention GmbH

- └ Founded in 2002
- └ Manufacturer of Univention Corporate Server (UCS) and Univention Groupware Server (UGS)
- └ UCS: a complete infrastructure solution based on Debian GNU/Linux, with an integrated infrastructure and identity management system (iims) build on top of OpenLDAP, Heimdal Kerberos, Samba and other components
- └ UGS: a complete Groupware solution based on Kolab 2 and UCS



Some References



[...]

The MigOS project of the german federal parliament

In 2002 the german federal parliament (Deutscher Bundestag) decided to migrate:

- └ its Windows NT 4.0 based clients to Windows XP
- └ its Windows NT 4.0 based servers to Linux / Samba / OpenLDAP

Reasons:

- └ Microsoft announced to discontinue support for Windows NT
- └ Windows based clients promised better compatibility to extern partners and existing applications
- └ independence, higher flexibility and higher transparency were the main reasons to select OSS for the servers

The challenge

- └ migrate an environment with approx. 5.000 clients, 120 servers and 8.000 user accounts from a single Windows NT 4.0 Domain to a Samba / OpenLDAP domain
- └ keep effects on users as low as possible
- └ keep availability during migration as high as possible
- └ Integrate Samba with an OpenLDAP directory, which is used to manage DHCP, DNS, Mail and other services, too
- └ develop tools for an integrated administration

The order for the migration was won by Computacenter

Tests and Preparation

- └ a complete, separate test environment was established
- └ migration of the account database and data was tested
- └ for all applications, compatibility tests were done
- └ performance tests were done by means of smbclient scripts, which tested among others:
 - └ authentication against the samba servers
 - └ opening connections to specific shares (netlogon and home share)
 - └ copying files from the logon-server (the content of the netlogon share)
 - └ copying the users profile data from a file server
 - └ using NT file servers with Samba logon servers

The breakdown

- └ the productive environment was migrated to Samba on the weekend 9th and 10th of october 2004.
- └ in the morning of monday the 11th the server reached critical load and were unable to process logon requests
- └ at 2:00 p.m. the fall back plan was executed and the NT servers were put in place, again
- └ bad headlines for Linux, OpenLDAP and Samba :-)



The screenshot shows a news article from a website. The header includes the word 'news' in yellow, the date '12.10.2004 12:09', and navigation links '<< Vorige | 1'. The article title is 'Linux-Migration im Bundestag: Kommando zurück beim Anmeldedienst'. The text of the article begins with 'Nach schweren Problemen bei der Verbindung der Windows-Clients mit den laufenden Open-Source-Servern im Bundestag haben die Verwaltung und der verantwortliche IT-Dienstleister CC CompuNet zunächst die alte Windows-L reaktiviert. Die mit der Migration beauftragte Firma untersucht derzeit die ITs:'. On the left side of the screenshot, there is a vertical menu with links: 'News', 'rchiv', 'terwegs', 'ter', 'abinden', 'nfe', 'nfe', and 'örungen'.

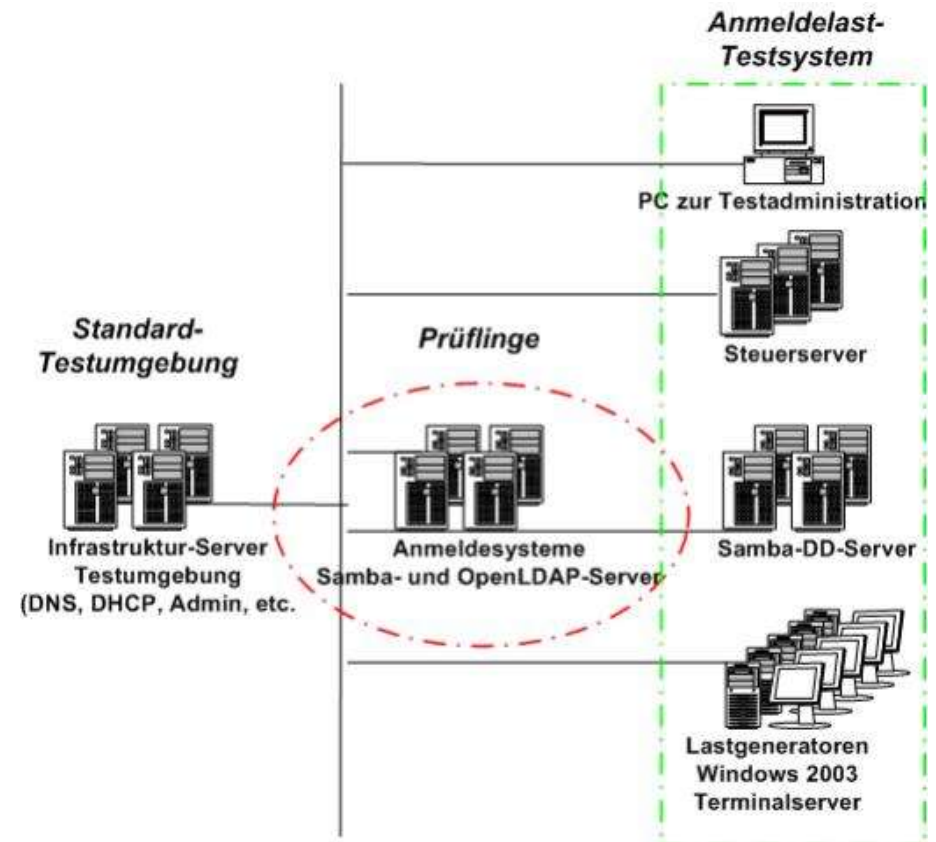
Clearing up

- └ After the crash the parliament charged a group of four companies (Gonicus, PSI, It-ec and Univention) with:
 - └ planing and building a test environment to:
 - reproduce the crash
 - test improved setups
 - └ find possible reasons for the crash
 - └ as well as to propose an optimized setup

The test environment

consisted of three groups of systems:

- └ test systems
- └ load generators
- └ control systems
- └ the test system were the systems, which previously failed in the productive environment
- └ we had about 30 load generators, among them 8 strong servers, on all Windows 2003 terminal services was installed (up to 6.000 sessions possible)
- └ 3 servers were used as control systems

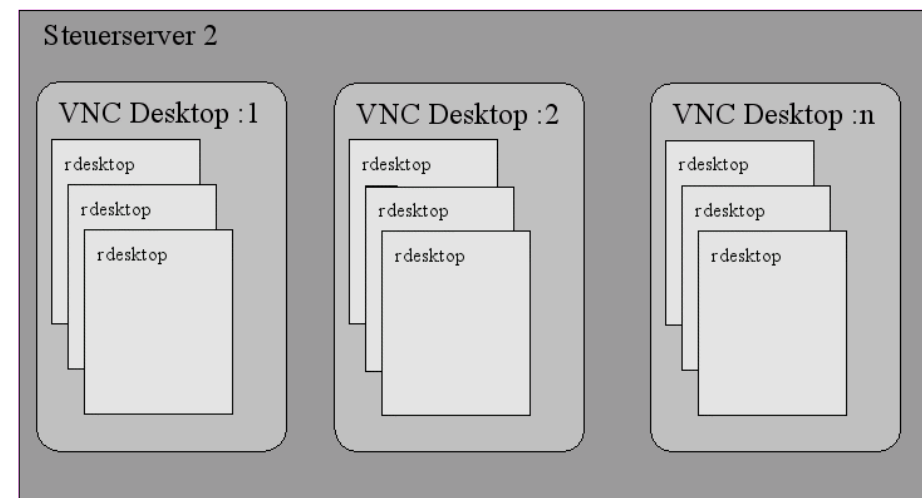
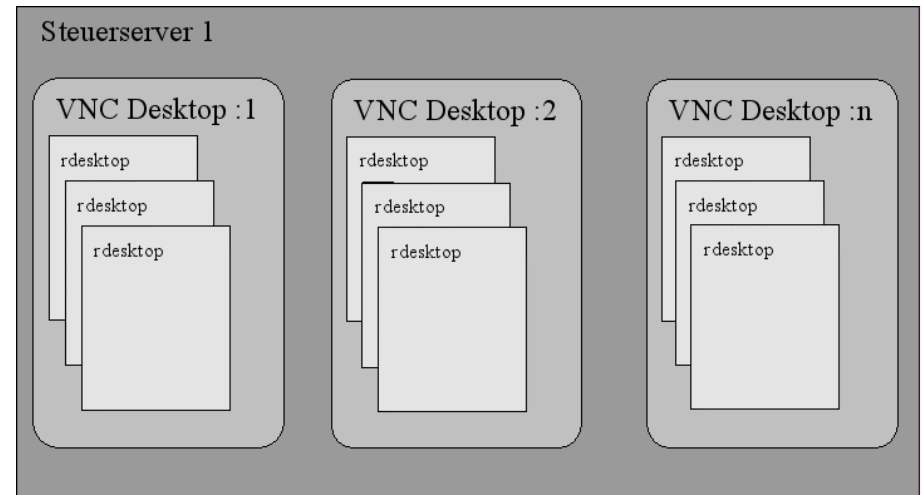


Why Windows Terminal Server sessions?

- └ obviously a terminal server can do more logons than a XP client at the same time :-)
- └ using Windows 2003 terminal server sessions has a high „face validity“ when simulating Windows XP sessions
(makes the client to believe in the test results)
- └ Windows 2003 terminal server sessions behave very similar to Windows XP sessions (very similar sequence of requests, same load on the samba server, same load on the OpenLDAP servers)


The control systems

- └ linux based
- └ were running several VNC sessions
- └ several rdesktop sessions inside each VNC session (because only a limited number of windows were possible inside one X session)
- └ about 3.000 sessions were possible from a single server with 4 Xeon CPUs and 16 GB RAM
- └ the test scripts were started and stopped from a standard PC via SSH




Planing the test

- └ a log file of a Windows Server was available to reconstruct the number of logon request during the morning of the crash day (approx 6.500 logon attampts, usually approx. 4.000 logons)




Usual logon requests

30.11. - 06.12.2004




Results of the crash reproduction (1)

Lastentwicklung



Samba-Prozesse



Results of the crash reproduction (2)

- └─ logon requests took up to 30 minutes (!)
- └─ out of approx. 6.500 logon request, only approx. 2.100 requests succeeded at all
- └─ in the log files:

```
smbd: can't contact ldap server
```


```
slapd: too many open files
```


Some reasons and possible improvements


- └ OpenLDAP was used with the ldbm database backend
- └ The OpenLDAP ACLs were not optimized
- └ The handling of group memberships was very inefficient (The code tested for group membership in every mapped group and the domain had several thousand groups); this was fixed by Andrew Bartlett one day before (!) the first migration attempt was done.
- └ NSCD was not configured to efficiently cache user and group data
- └ While we found and evaluated these improvements, Sernet developed further improvements in the Samba code (ldapsam:trusted=yes, getpwnam-Cache and more)

Results of the improvements (1)

Lastentwicklung



Samba-Prozesse



Results of the improvements (2)

- └ the domain was successfully migrated in summer 2005
- └ Better press for Linux, OpenLDAP and Samba :-)



Conclusion

- └ In a Samba 3.0 / OpenLDAP based domain both components should be seen (and tested) as one Unit
- └ Tests should be as realistic as possible (smbclient could not be used to simulate Windows XP logon requests; logon scripts can cause high loads)
- └ Samba ($\geq 3.0.7$) / OpenLDAP (with the bdb backend) and Linux proved themselves as a efficient and performant replacement of Windows NT

Thank you!

└ Questions?