




HEPIX SPRING 2018



NIKHEF SITE REPORT

Bart van der Wal

NIKHEF SITE REPORT

- Grid compute, storage and network
- Grid management to new hardware
- Saltstack

- Local cluster
- Dcache

- Computer Lab
- LDAP and Active Directory
- Active Directory service disruption



NEW DELL CLUSTER 'PEPERNOOT'



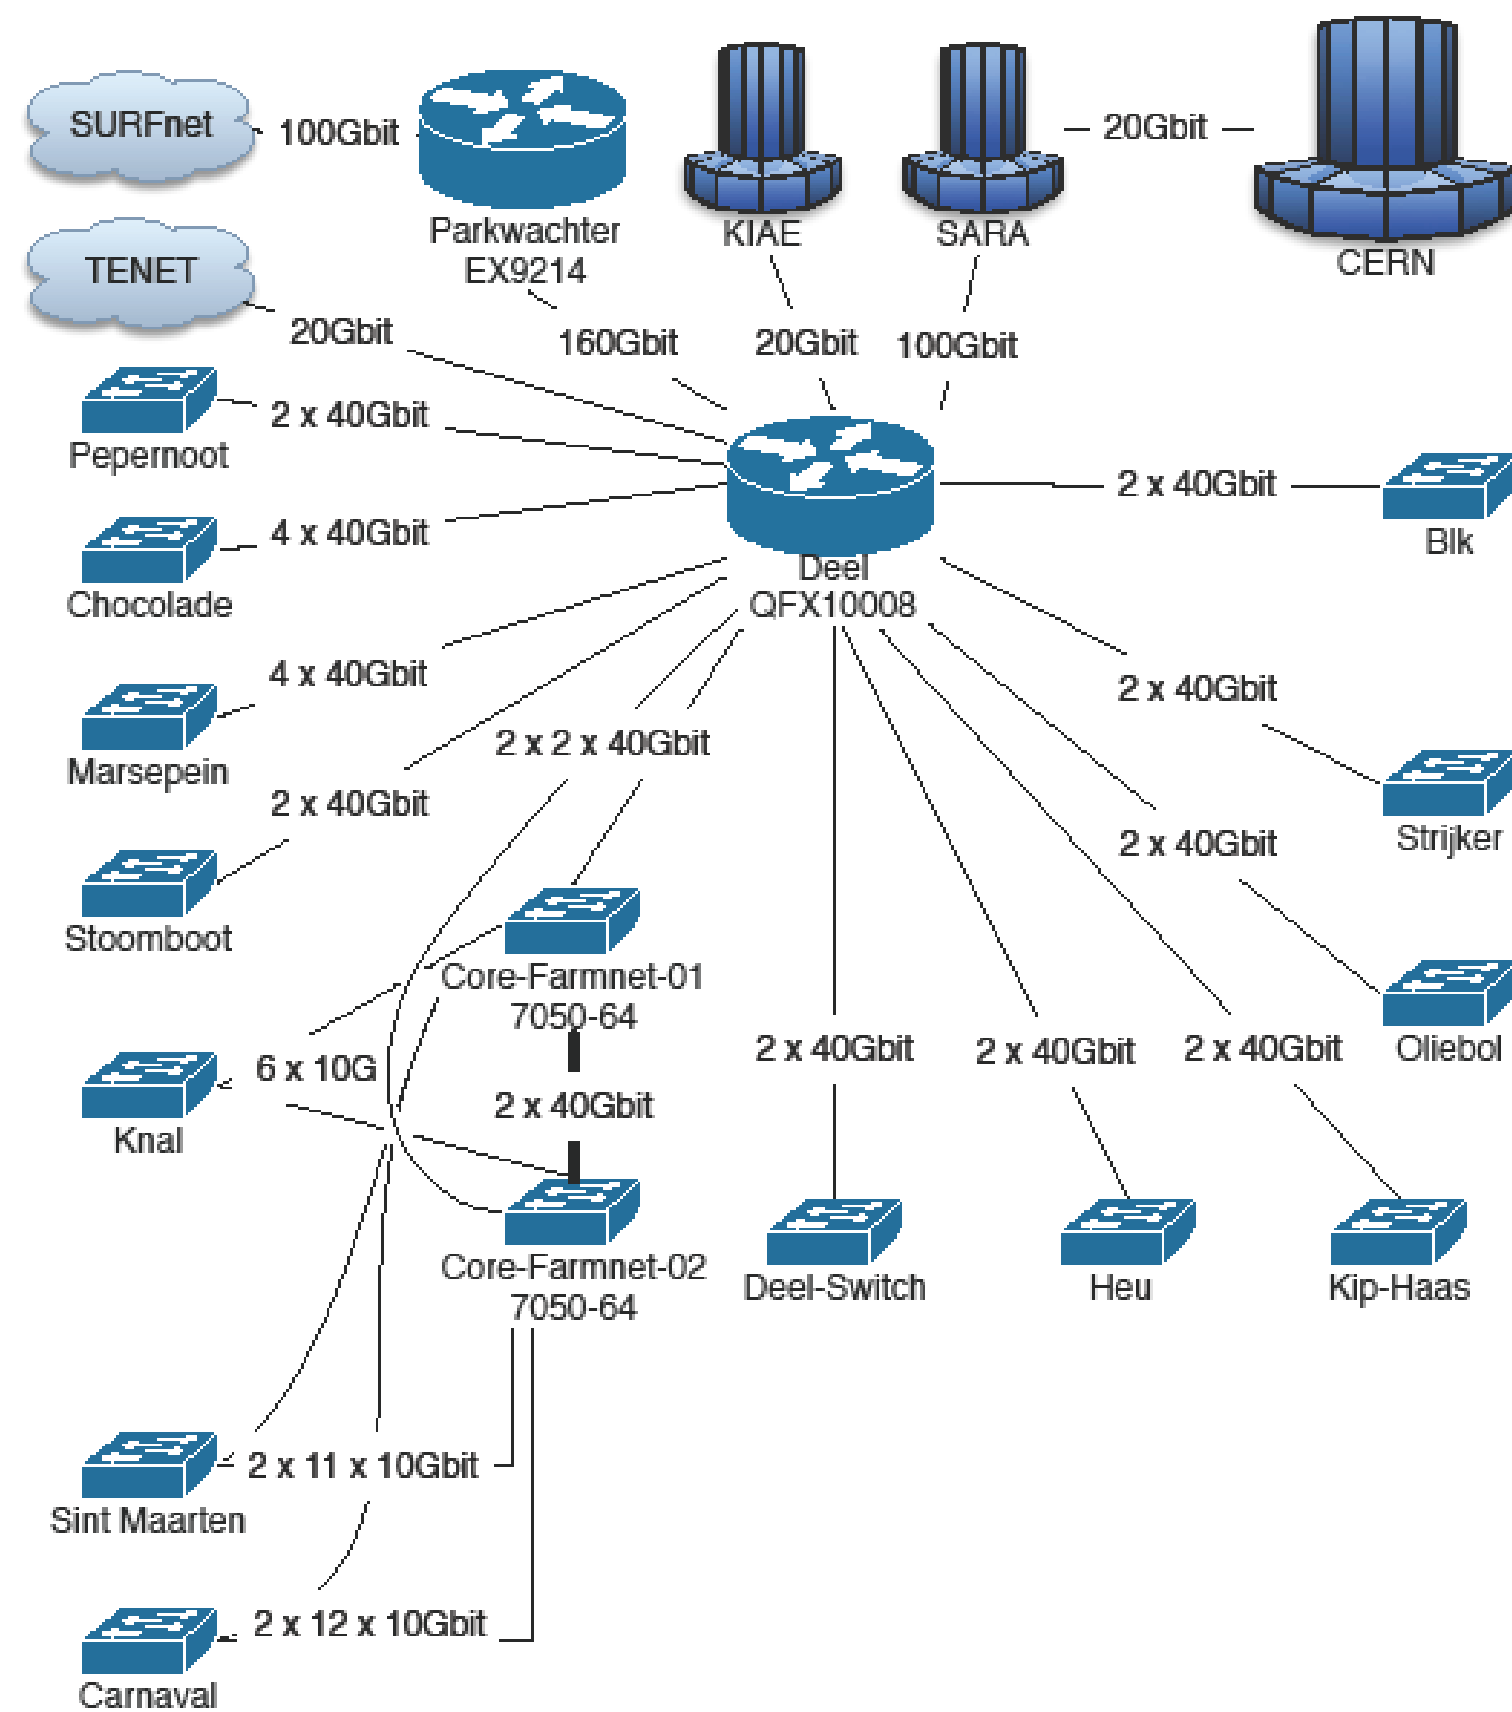
- 13 DELL PowerEdge R640
- dual socket Intel(R) Xeon(R) Gold 6148 CPU @ 2.40GHz
- 20 cores/cpu.
- Hyperthreading → 80 virtual cores
- 384 GB memory (9.6 GB/core)
- 3 TB NVMe (ridiculously fast) storage.
- 25Gb/s network per node

NEW GRID STORAGE 'HOOIBRAND'

- 4 Blocks of 350 TiB
- Fujitsu RX2530M4
- 40Gb/s network
- Seagate Assured SAN trays
- 56 X 8TB NLSAS
- 4 times raid 6
- 1.4 PiB in 20U
- Details in talk of Tristan Suerink



GRID NETWORK



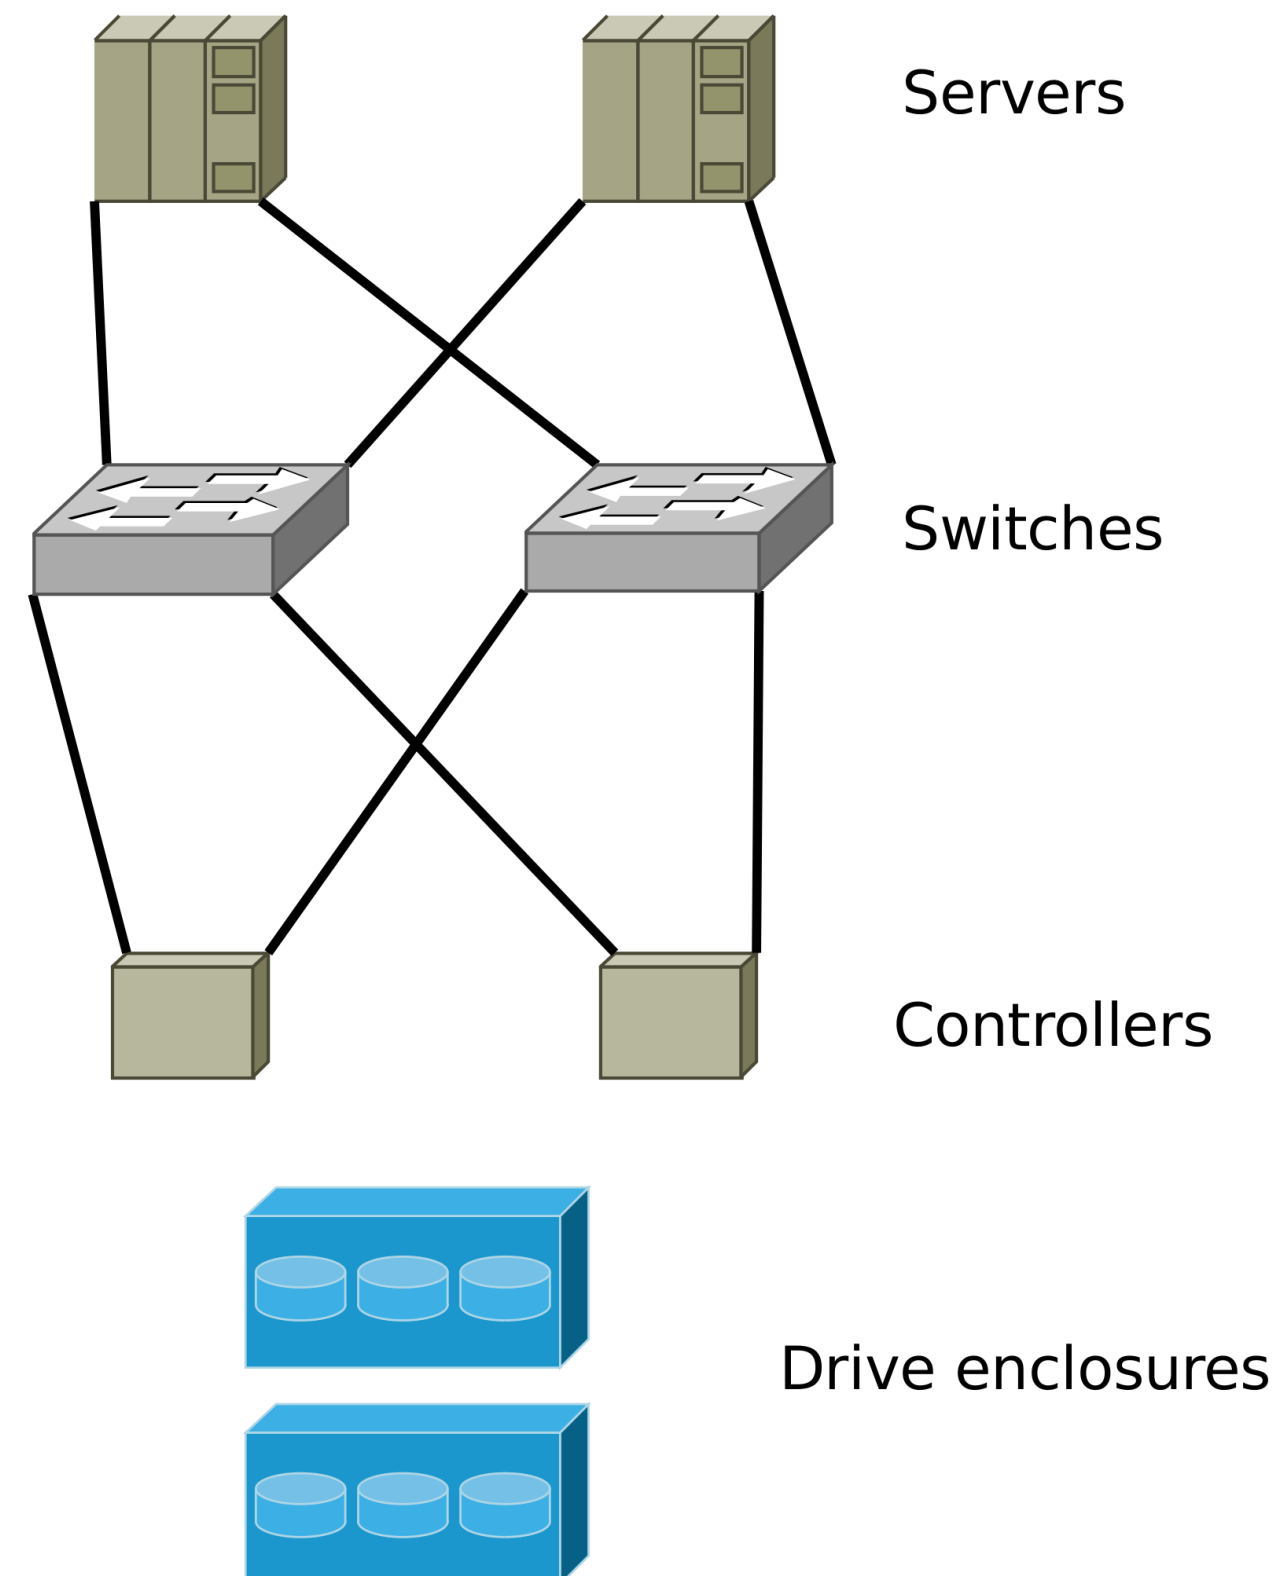
- New Juniper QFX10008 'Deel'
- 60 times 40Gb/s ports in use
- 6 times 100Gb/s ports in use
- Able to grow to ~90Tb/s
- 400Gb/s ready, waiting for Juniper
- Multi VRF for traffic separation
- Gateway for HTC cloud

DELL COMPELLENT → FUJITSU ETERNUS

- 5 y/o Fibre Channel Compellent
- Backbone of VM platform for GRID
- New Fujitsu ETERNUS DX200S4
- 17 times 2TB SSD
- 26 times 6TB HDD
- No Fibre Channel just iSCSI
- Auto-tiering (working fine)
- Auto-QoS bad performance
- Auto-QoS turned off.



GRID MANAGEMENT TO NEW HARDWARE



- Platform for Grid services
- Citrix Xenserver 7.0
- 4 PRIMERGY RX2530 M4 servers
- Double Xeon 6148 CPU
- 4x10GB iSCSI to ETERNUS

GRID MANAGEMENT TO NEW HARDWARE

- (Mostly) live migration
- ~160 servers ~20GB per server
- More than one week
- Older kernels failed at the end
- Probably hardware compatibility
- New OS and data migration
- No upgrade because of licencing
- Near future Xen Cloud Platform?



SALT STACK



- Implementing Saltstack
- Talk by Dennis van Dok

LOCAL CLUSTER 'STOOMBOOT'



- Running on 10 year old hardware
- New hardware planned this year
- Torque updated to 4.2.10
- Head node updated to CentOS 7
- New interactive nodes (pepernoot)
- Under Salt control

DCACHE

- NFS mount on local cluster
- Added 300TB repurposed hardware
- 844TB total
- Management to CentOS 7
- Upgraded to DCache 3.1
- High load new bugs found
- Mostly NFS client issues
- Solved with Dcache developers
- Now stable!



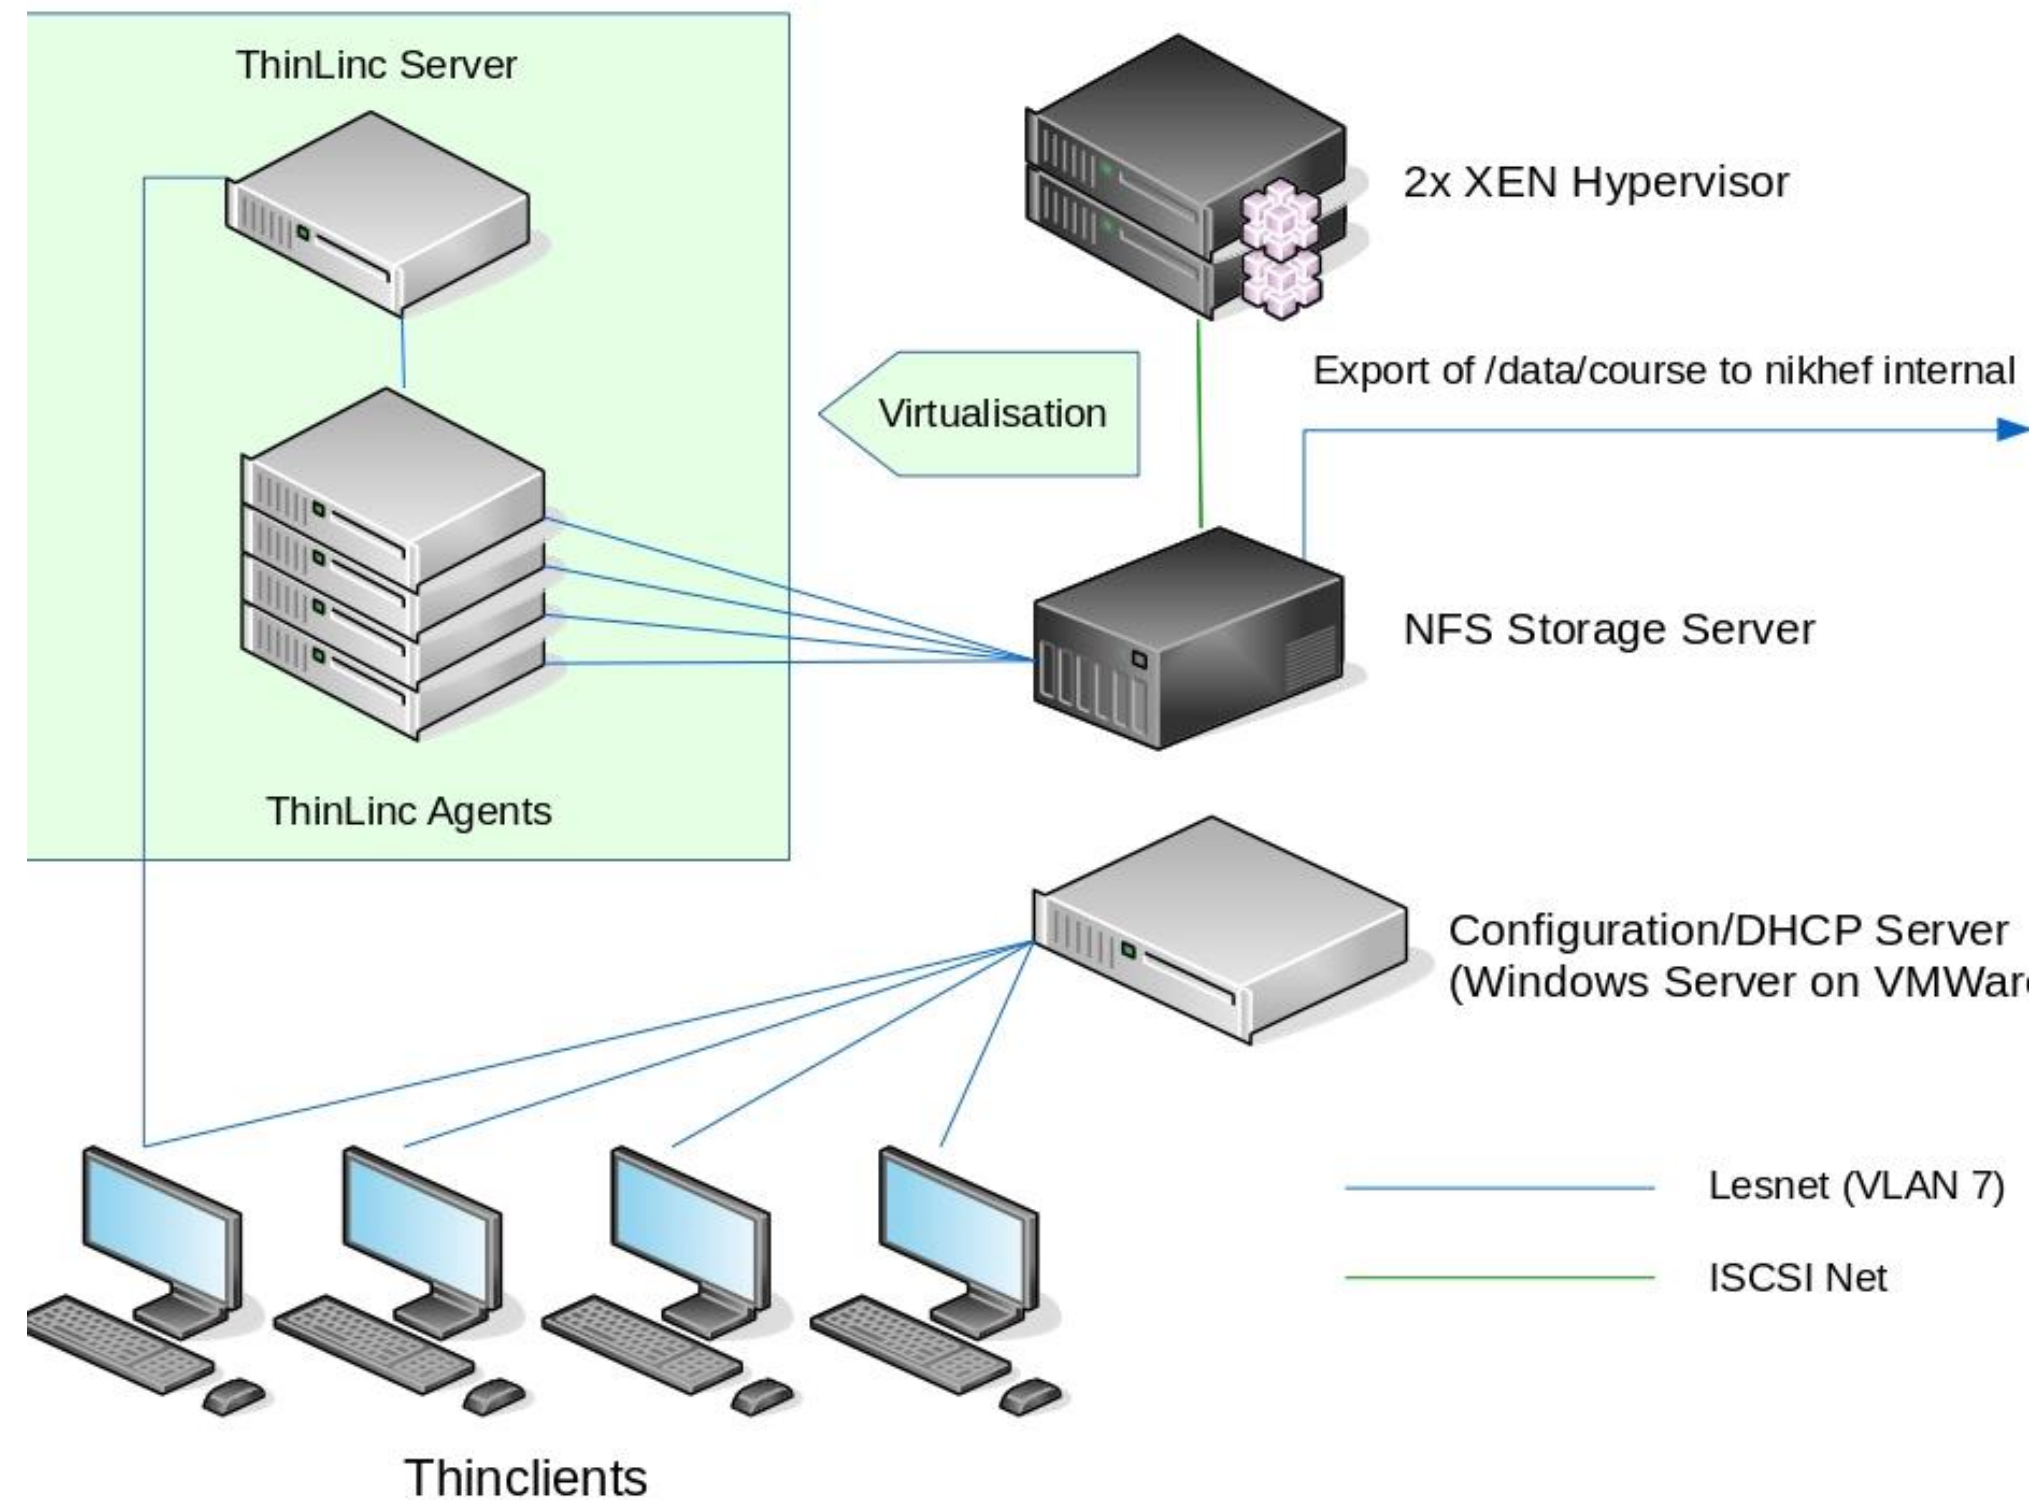
COMPUTER LAB

- Physics/programming courses
- Students and postdocs
- Different universities

- PC's with standard SLC
- Students share accounts
- On main network
- No external access



COMPUTER LAB

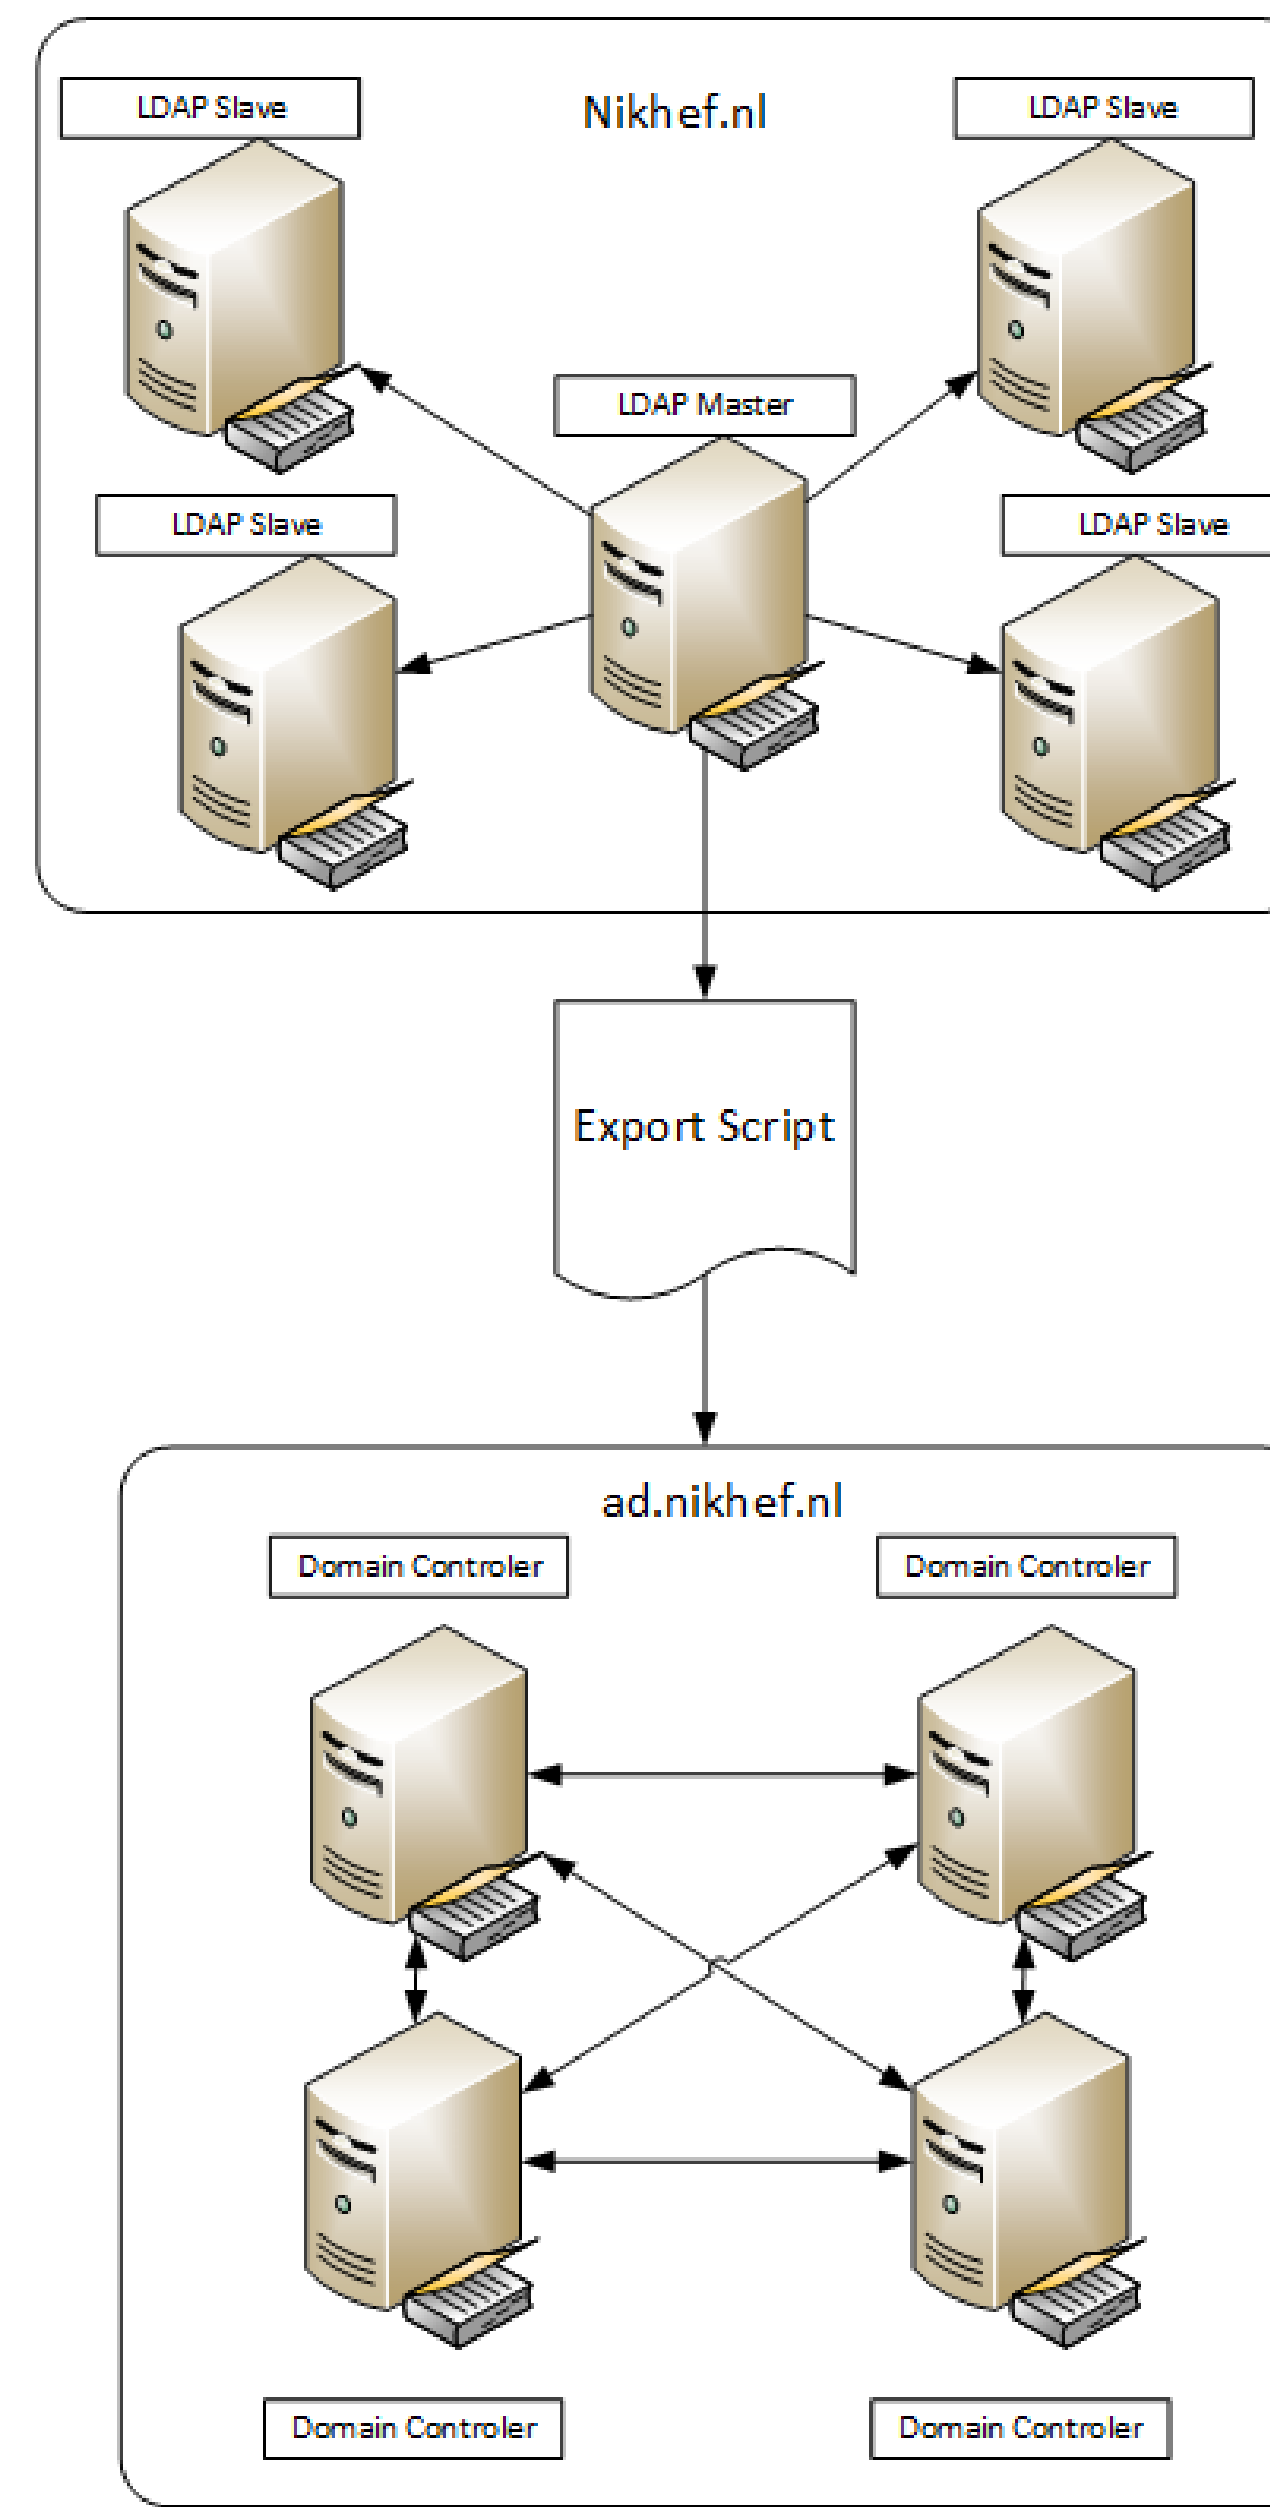


- Cendio Thinlinc remote desktop
- Virtual server and desktop
- All in one Thin Clients (Fujitsu)
- Custom image per course
- Personal student accounts
(created by course organizer)
- External access via https
- Seperate secure network

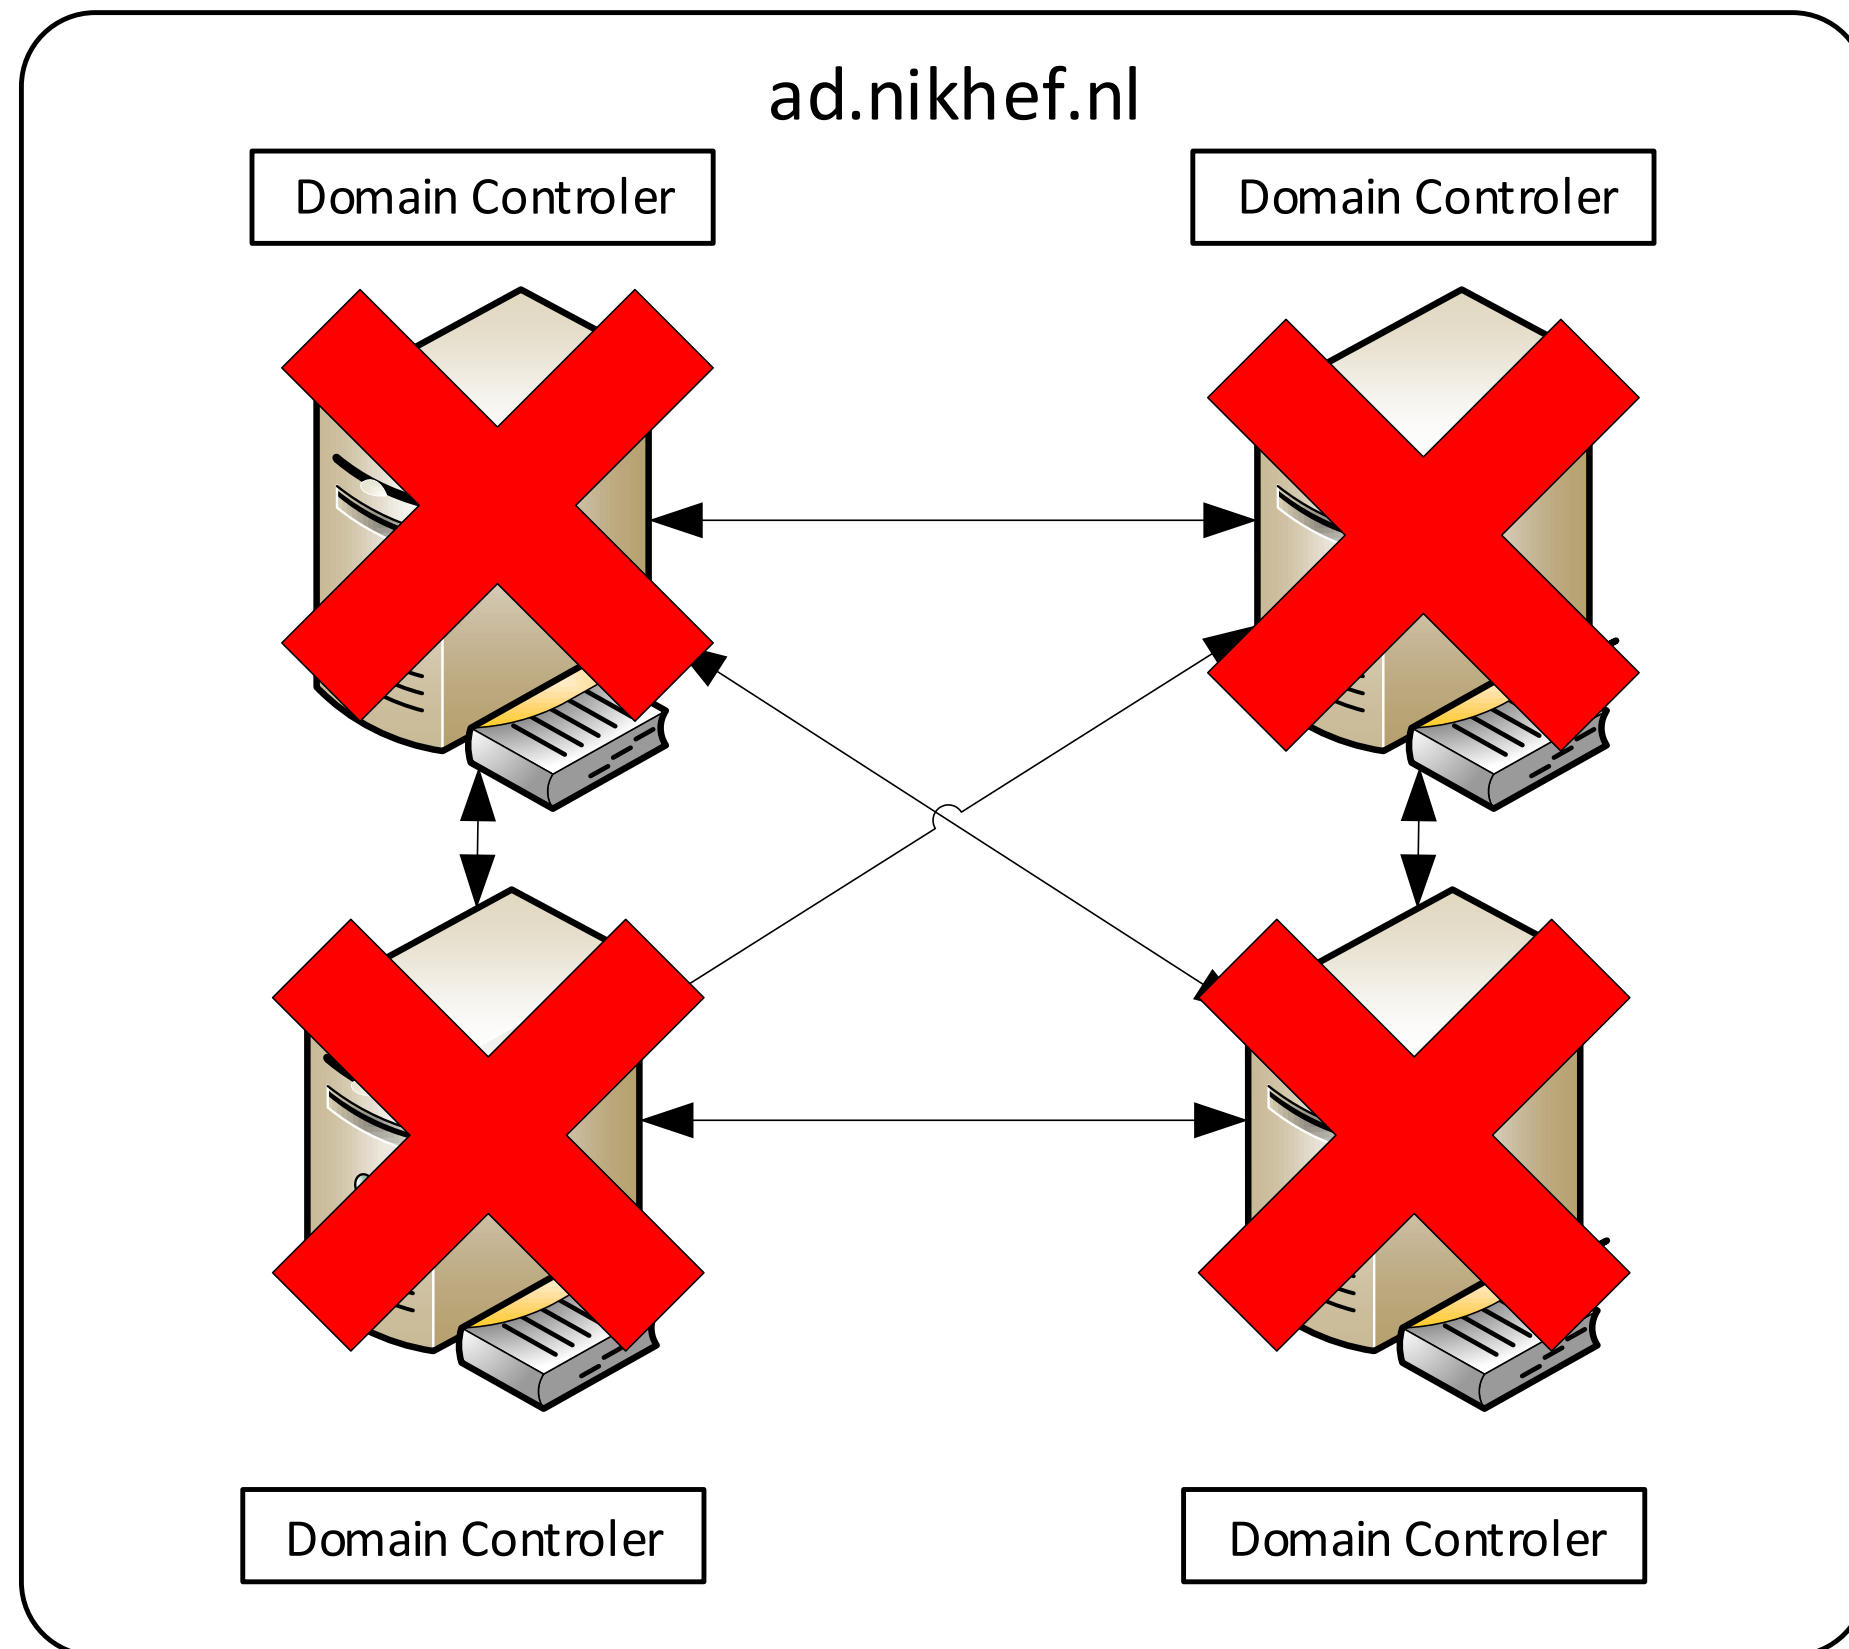
LDAP & ACTIVE DIRECTORY

- Separate LDAP and AD
- Two accounts same username
- Only passwords synced (mostly..)
- Groups were a mess

- Synchronization complex
- Instead LDAP → AD
- Custom python script
- Most administration now in LDAP
- Nice cleanup of groups



ACTIVE DIRECTORY SERVICE DISRUPTION



- Friday one Domain Controller
- Restore taking several hours
- Monday all! Domain Controllers
- Restore of first DC 20 minutes
- Restore of the rest took hours.
- Update KB4088883 (preview!)
- WSUS no difference preview patch
- Script to decline preview patches
- New backup system
- Staggered updates for DC's

My thanks to:

Andrew Pickford, Daniel Kolmer, Dennis van Dok, Paul Kuipers and Tristan Suerink

Questions?