Installing and maintaining clusters of FreeBSD servers using PXE and Rsync

Cor Bosman

XS4ALL

cor@xs4all.net

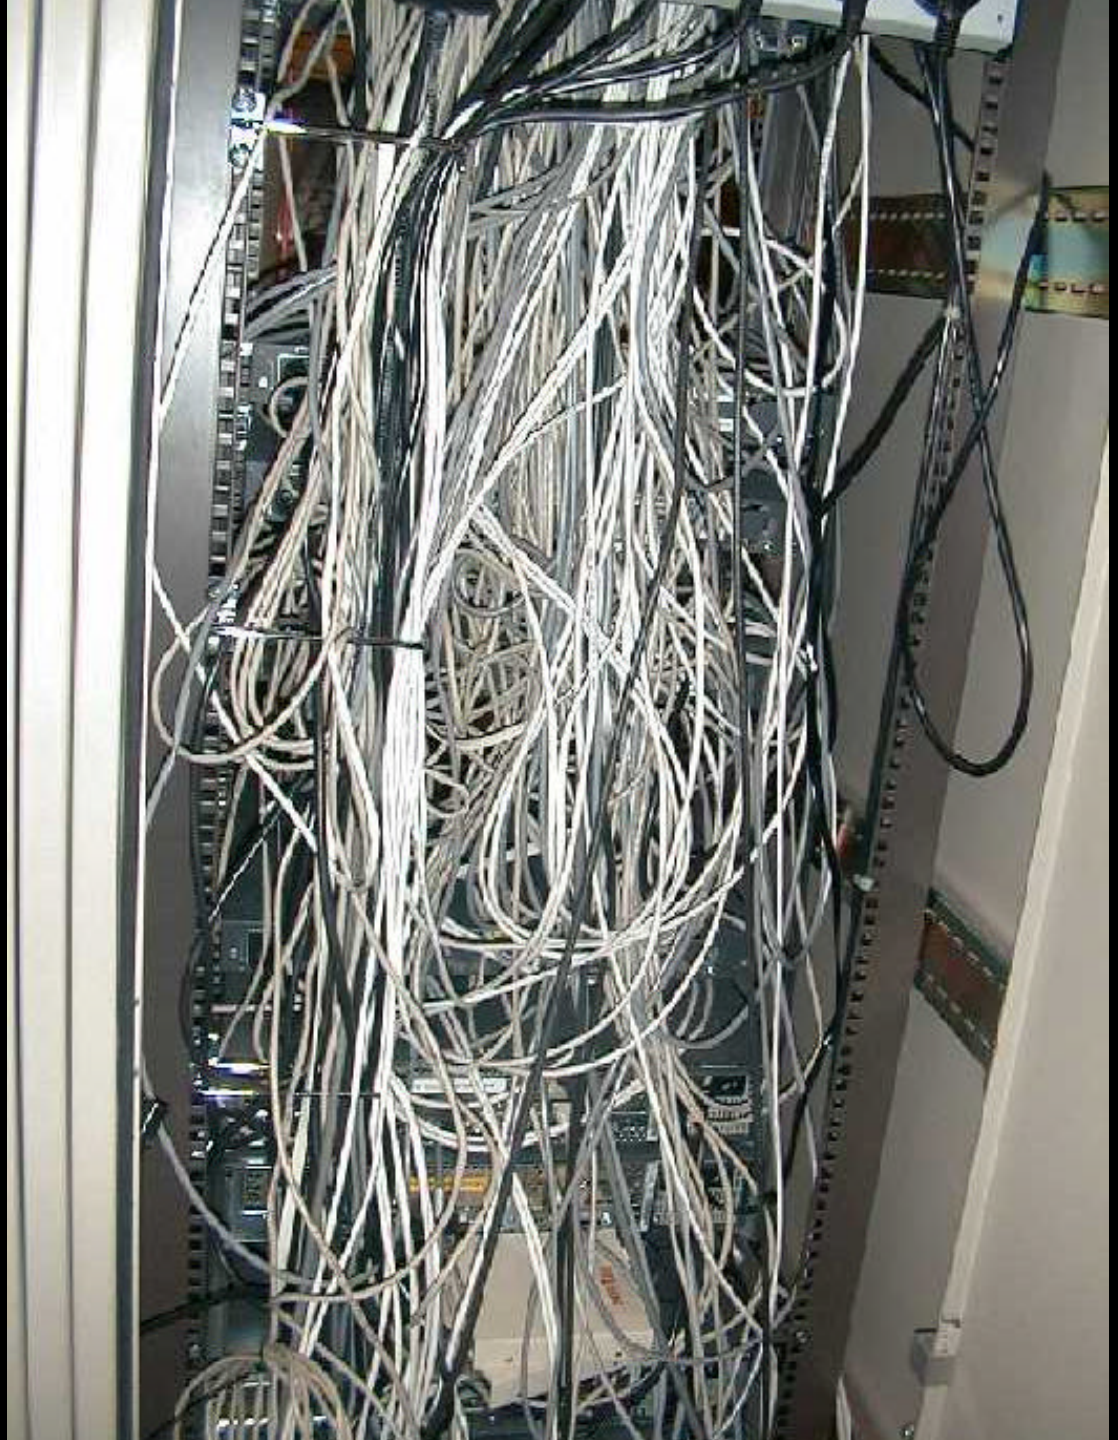

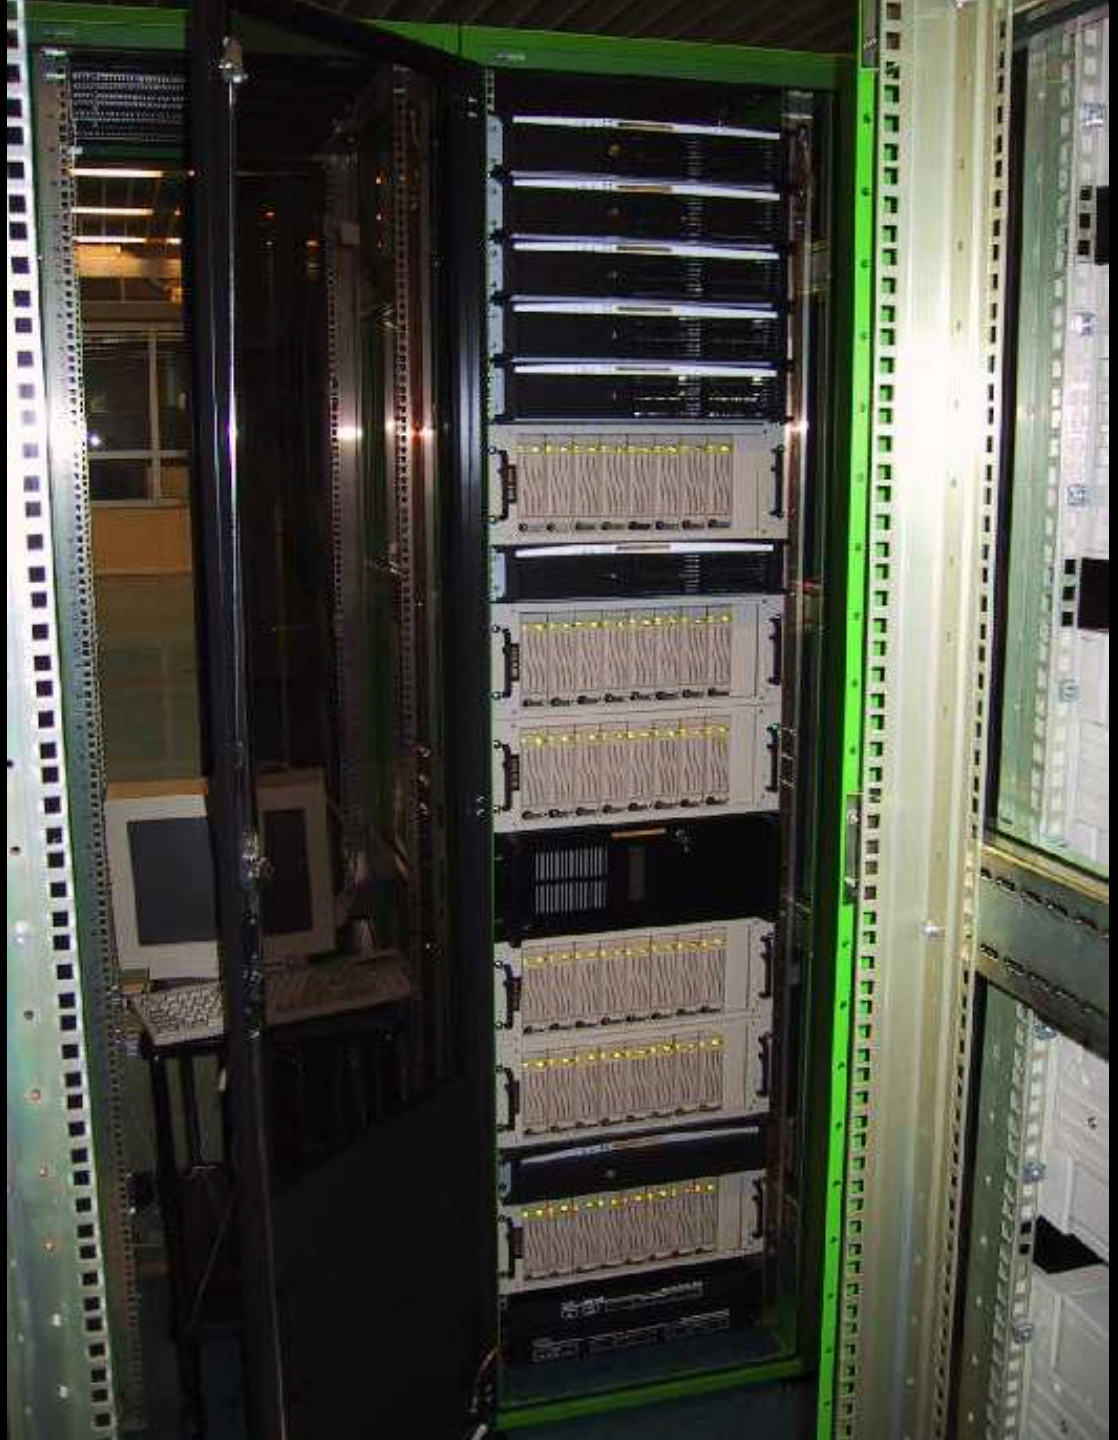

The problem

Operating dozens of servers individually

- Installation
- Maintenance
- Security
- Upgrades
- Number of servers increased rapidly. A solution is necessary

Fixing the problem

- Standardize hardware
- Create fast and easy installation
- Centralize maintenance to keep groups of servers identical

Standardize hardware

- Choose a few different server layouts. Low, medium, high performance. We prefer on-board devices so we can use 1U rack mounted servers
- Spare policy is simple. Keep one or more spare boxes. Swap complete box when hardware fails
- You always know what to expect. Hardware will be supported by your OS

installation

Different possibilities

- Copy an image using tools like dd. Slow, can’t easily be done in parallel. Problem with drive sizes
- Scripted install using floppy/CD
- PXE
  - Preboot Execution Environment
  - Standard created by Intel
  - Built on TCP/IP, DHCP, TFTP

PXE

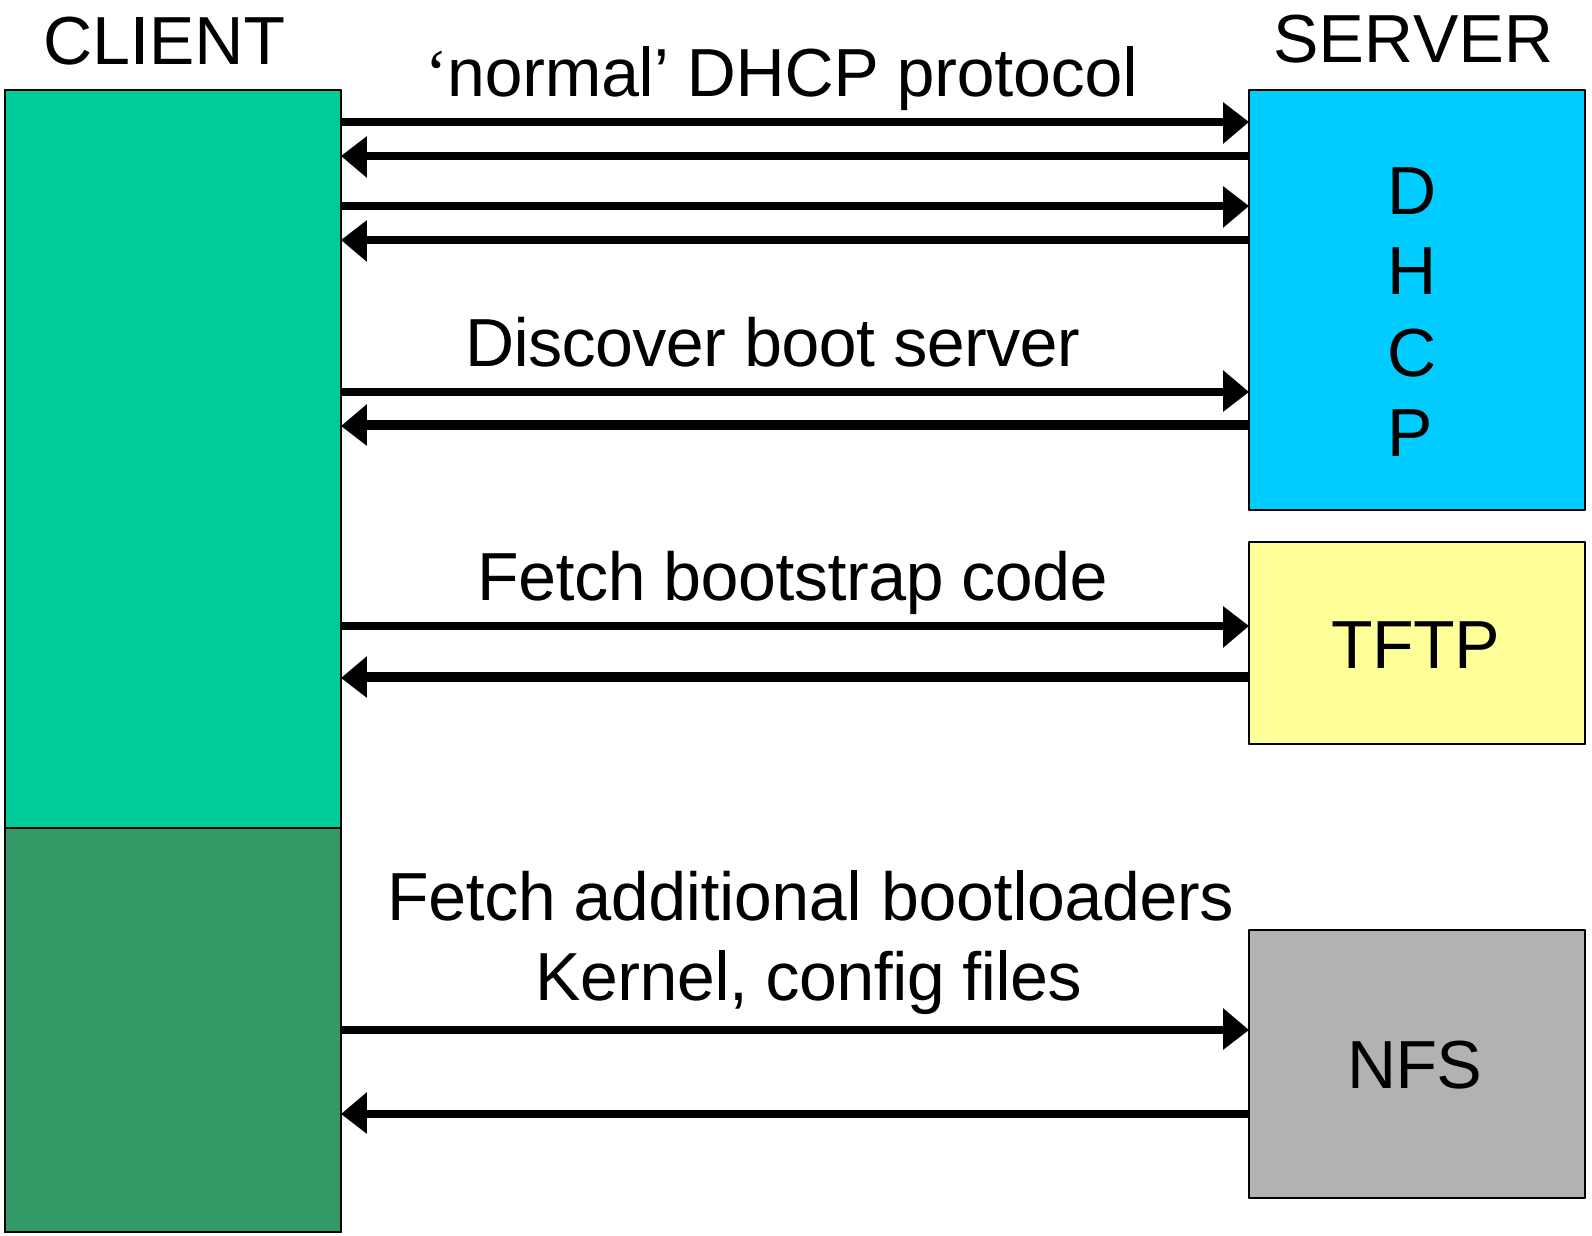
CLIENT
SERVER
‘normal’ DHCP protocol
D
H
C
P
Discover boot server
Fetch bootstrap code
TFTP
Fetch additional bootloaders
Kernel, config files
NFS

Installing FreeBSD using PXE

- PC with PXE capable network card
- Ethernet connection
- DHCP/Boot server
- TFTP server
  - Boot image
- NFS server
  - Boot loader files
  - Loader config file
  - Kernel
  - Memory file system
    - Install.cfg
- FreeBSD distribution

```
server-name "DHCPserver";
default-lease-time 86400;
option subnet-mask 255.255.255.0;
option broadcast-address 192.168.1.255;
option domain-name "xs4all.nl";
option domain-nameservers 194.109.6.66,194.109.9.99;
option routers 192.168.1.2;
subnet 192.168.1.0 netmask 255.255.255.0 {
        range dynamic-bootp 192.168.1.10 192.168.1.254;
        filename "pxeboot";
        next-server 192.168.1.3;
        option root-path "/usr/local/export/pxe";
```

Installing FreeBSD using PXE

- PC with PXE capable network card
- Ethernet connection
- DHCP/Boot server
- TFTP server
  - Boot image
- NFS server
  - Boot loader files
  - Loader config file
  - Kernel
  - Memory file system
    - Install.cfg
- FreeBSD distribution

oot/loader.rc

```
cho Loading Kernel...
oad /kernel
et choice=default
cho
cho Please select one of the following installs within 15 seconds
cho
cho default
cho scsi
cho dh
cho
ead -t 15 -p "Type in the exact word of your selection: " choice
cho
nclude /boot/loader.rc.$choice
cho booting...
et vfs.root.mountfrom="ufs:/dev/md0c"
oot
```

oot/loader.rc.default

```
oad -t mfs_root /mfsroot-default
```

Tips

- Use an install server. Insert empty HD
- Keep a stack of installed harddisks ready
- You can install multiple servers at the same time

Centralized maintenance

- Scaling choices have created clusters of servers
- Every cluster has a parent (“golden master”)
  - Doesn’t do any production work. Merely a repository
  - Low end hardware
  - Parent is called the ‘zero’ server. Production servers are named after their task. smtp0, smtp1, smtp2….
- OS has a parent

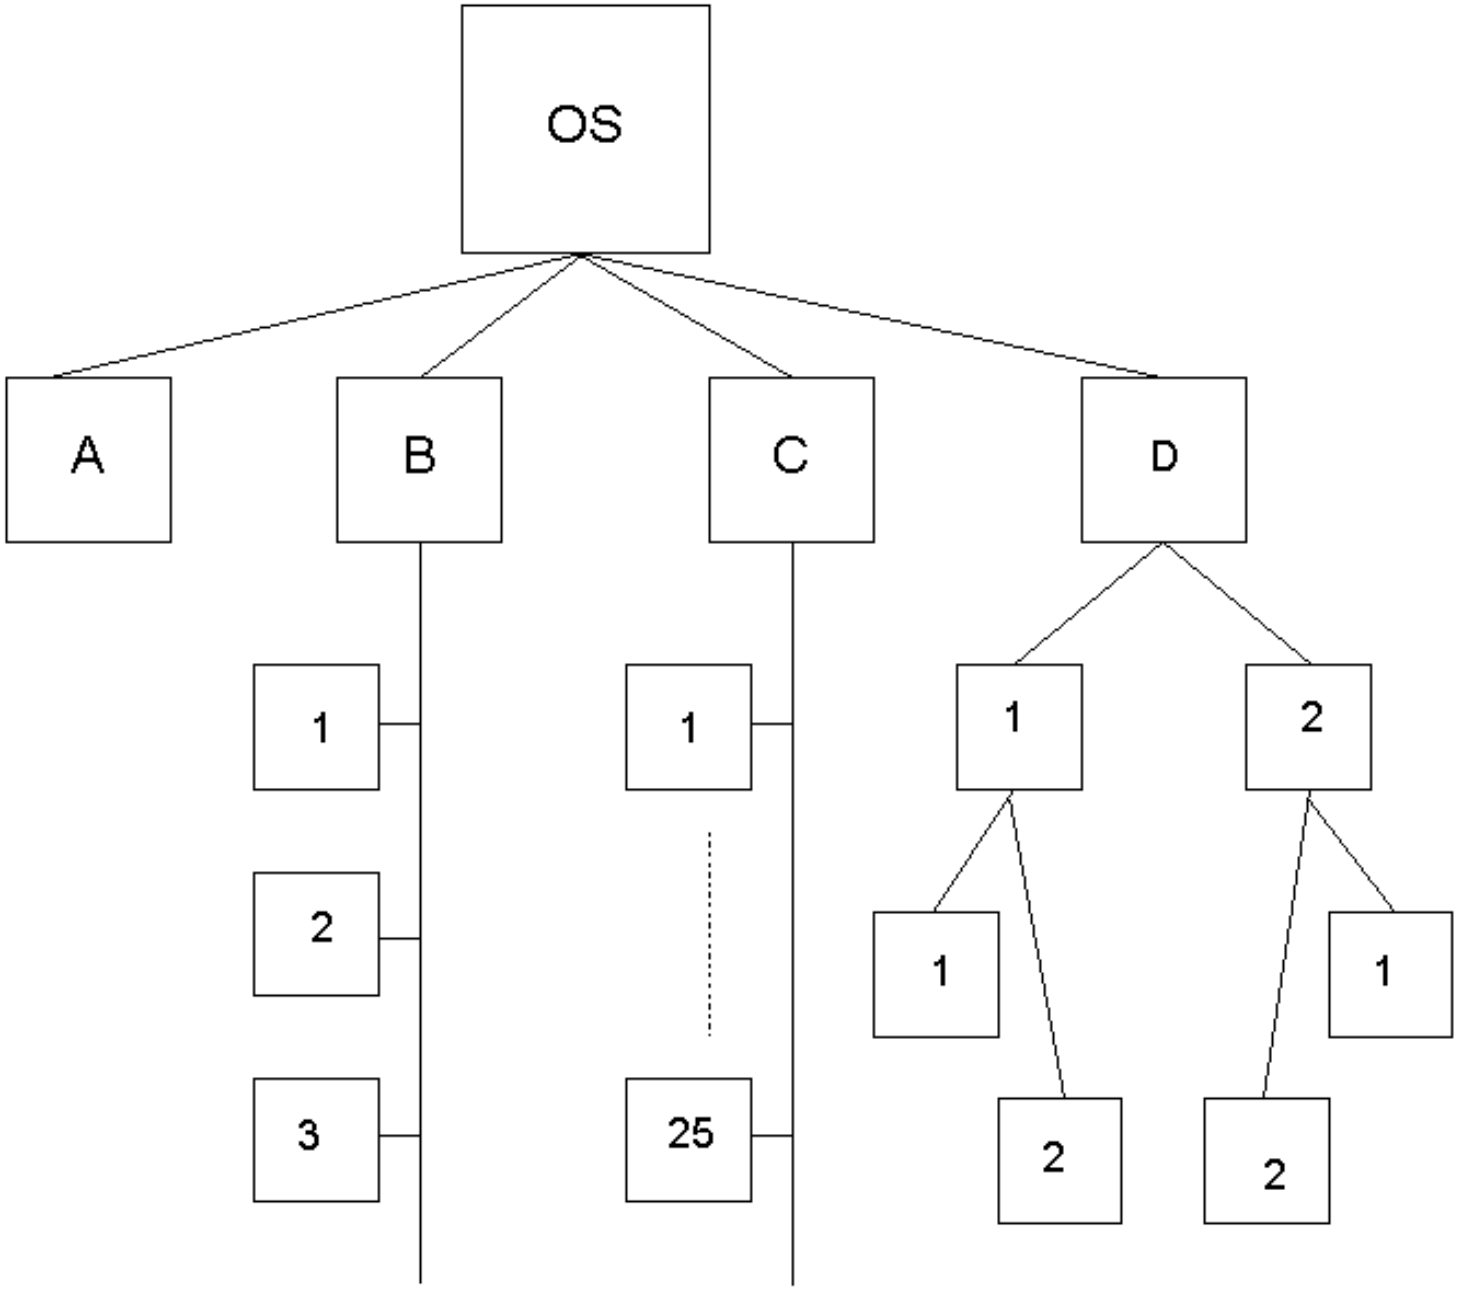
OS
A
B
C
D
1
2
3
1
25
1
2
1
2
1
2

Testing

- Testing on production server is a bad idea
- Testing on parent is also a bad idea
- Use specific test and accept servers

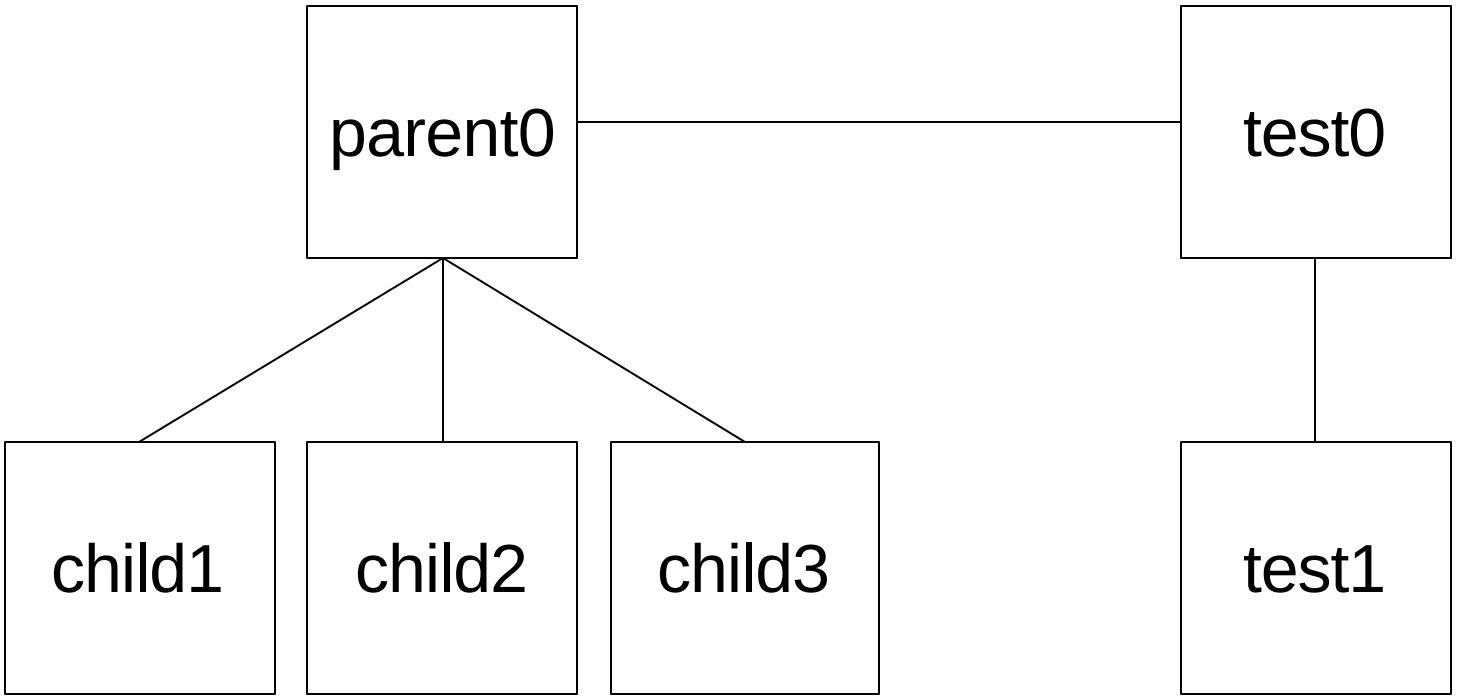

Synchronizing servers

- Mirror parent over the network
  - Exclude files
    - Machine specific config files
    - Temporary files/directories
    - Customer data
  - Preview changes

  Possible tools: rdist, rsync, …
- Manual syncing prone to human error

```
 sync

sage: sync -t -s <host> -f <conf file>

f      alternate file. Default is sync.conf
t      test, don't actually do it. Recommended before ALL syncs
s      hostname to sync. Use "all" for all sites.
c      allow recursive updates
```

onfig file

```
 cat sync.conf
 hostname            exclude_file                recursive?

mtp0.xs4all.nl       exclude.smtp                   yes
ebmail0.xs4all.nl    exclude.webmail                yes
h0.xs4all.nl         exclude.dh                     yes
ackup.xs4all.nl      exclude.backup                  no
```

Security

- Security and ease of use often go hand in hand
  - Define what is important to you
  - We try to avoid tasks that need interaction per individual server.
- Protect your parent servers
  - Harden servers
  - Firewall
  - IDS

Security

- Connecting to children
  - Use ssh key agent for authentication
- Use sync to check for changes
- Quick reaction possible

Enhanced PXE installation

- Installation is adding a server to a group
- Choices possible through loader.rc
- DHCP server drives the PXE installation

VLANS

- VLAN allows subnetting of physical network
- Divide switch into separate networks
- DHCP server can reply differently for each network

```
server-name "DHCPserver";
default-lease-time 86400;
option subnet-mask 255.255.255.0;
option broadcast-address 192.168.1.255;
option domain-name "xs4all.nl";
shared-network "VLAN0" {
        subnet 192.168.1.0 netmask 255.255.255.240 {
                range dynamic-bootp 192.168.1.4 192.168.1.14;
                filename "pxeboot";
                next-server 192.168.1.1;
                option root-path "/usr/local/export/pxe-freebsd-45-scsi";
                option domain-name-servers 192.168.1.1;
                option routers 192.168.1.1;
         }
 }
 shared-network "VLAN1" {
        subnet 192.168.1.16 netmask 255.255.255.240 {
                range dynamic-bootp 192.168.1.20 192.168.1.30;
                filename "pxeboot";
                next-server 192.168.1.17;
                option root-path "/usr/local/export/pxe-freebsd-45-ide";
                option domain-name-servers 192.168.1.17;
                option routers 192.168.1.17;
         }
}
```

Immediate Synchronization

- Create FreeBSD package for each type of server
- Install package from install.cfg
- “post” script runs sync with correct parent
- End result is fully installed and configured server, ready for production.
- Total install time less than 5 minutes

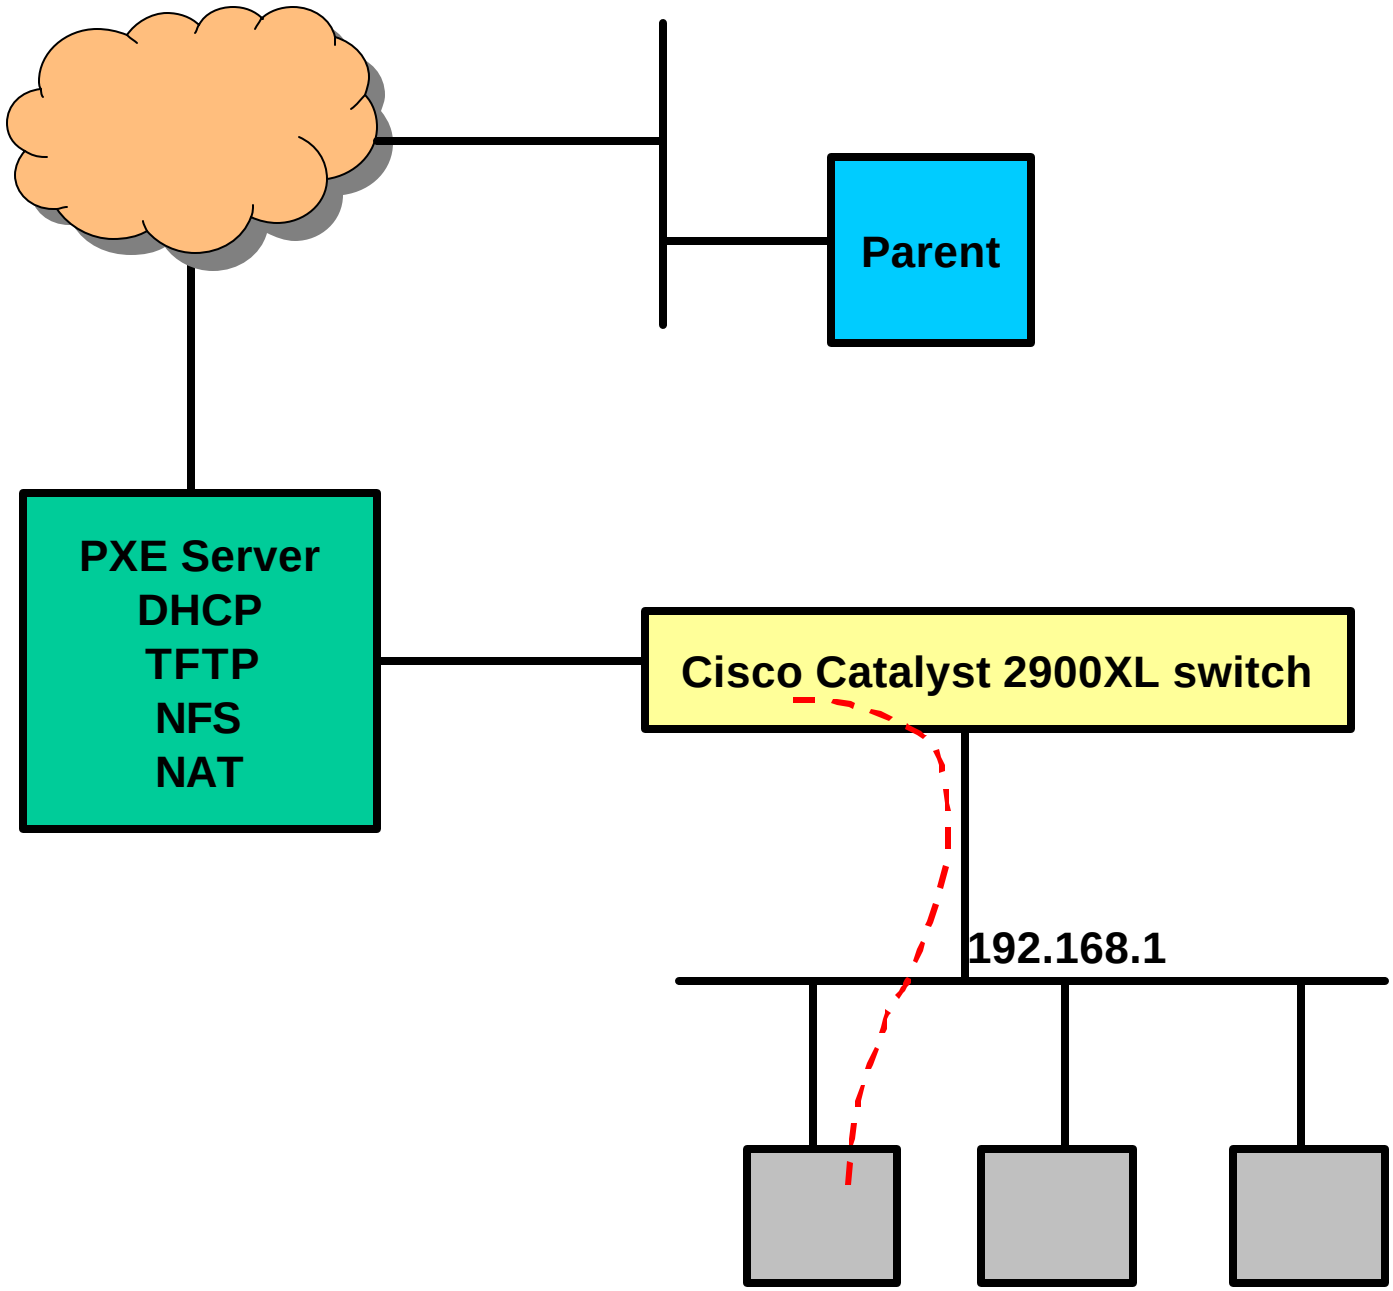
Parent
PXE Server
DHCP
TFTP
NFS
NAT
Cisco Catalyst 2900XL switch
192.168.1

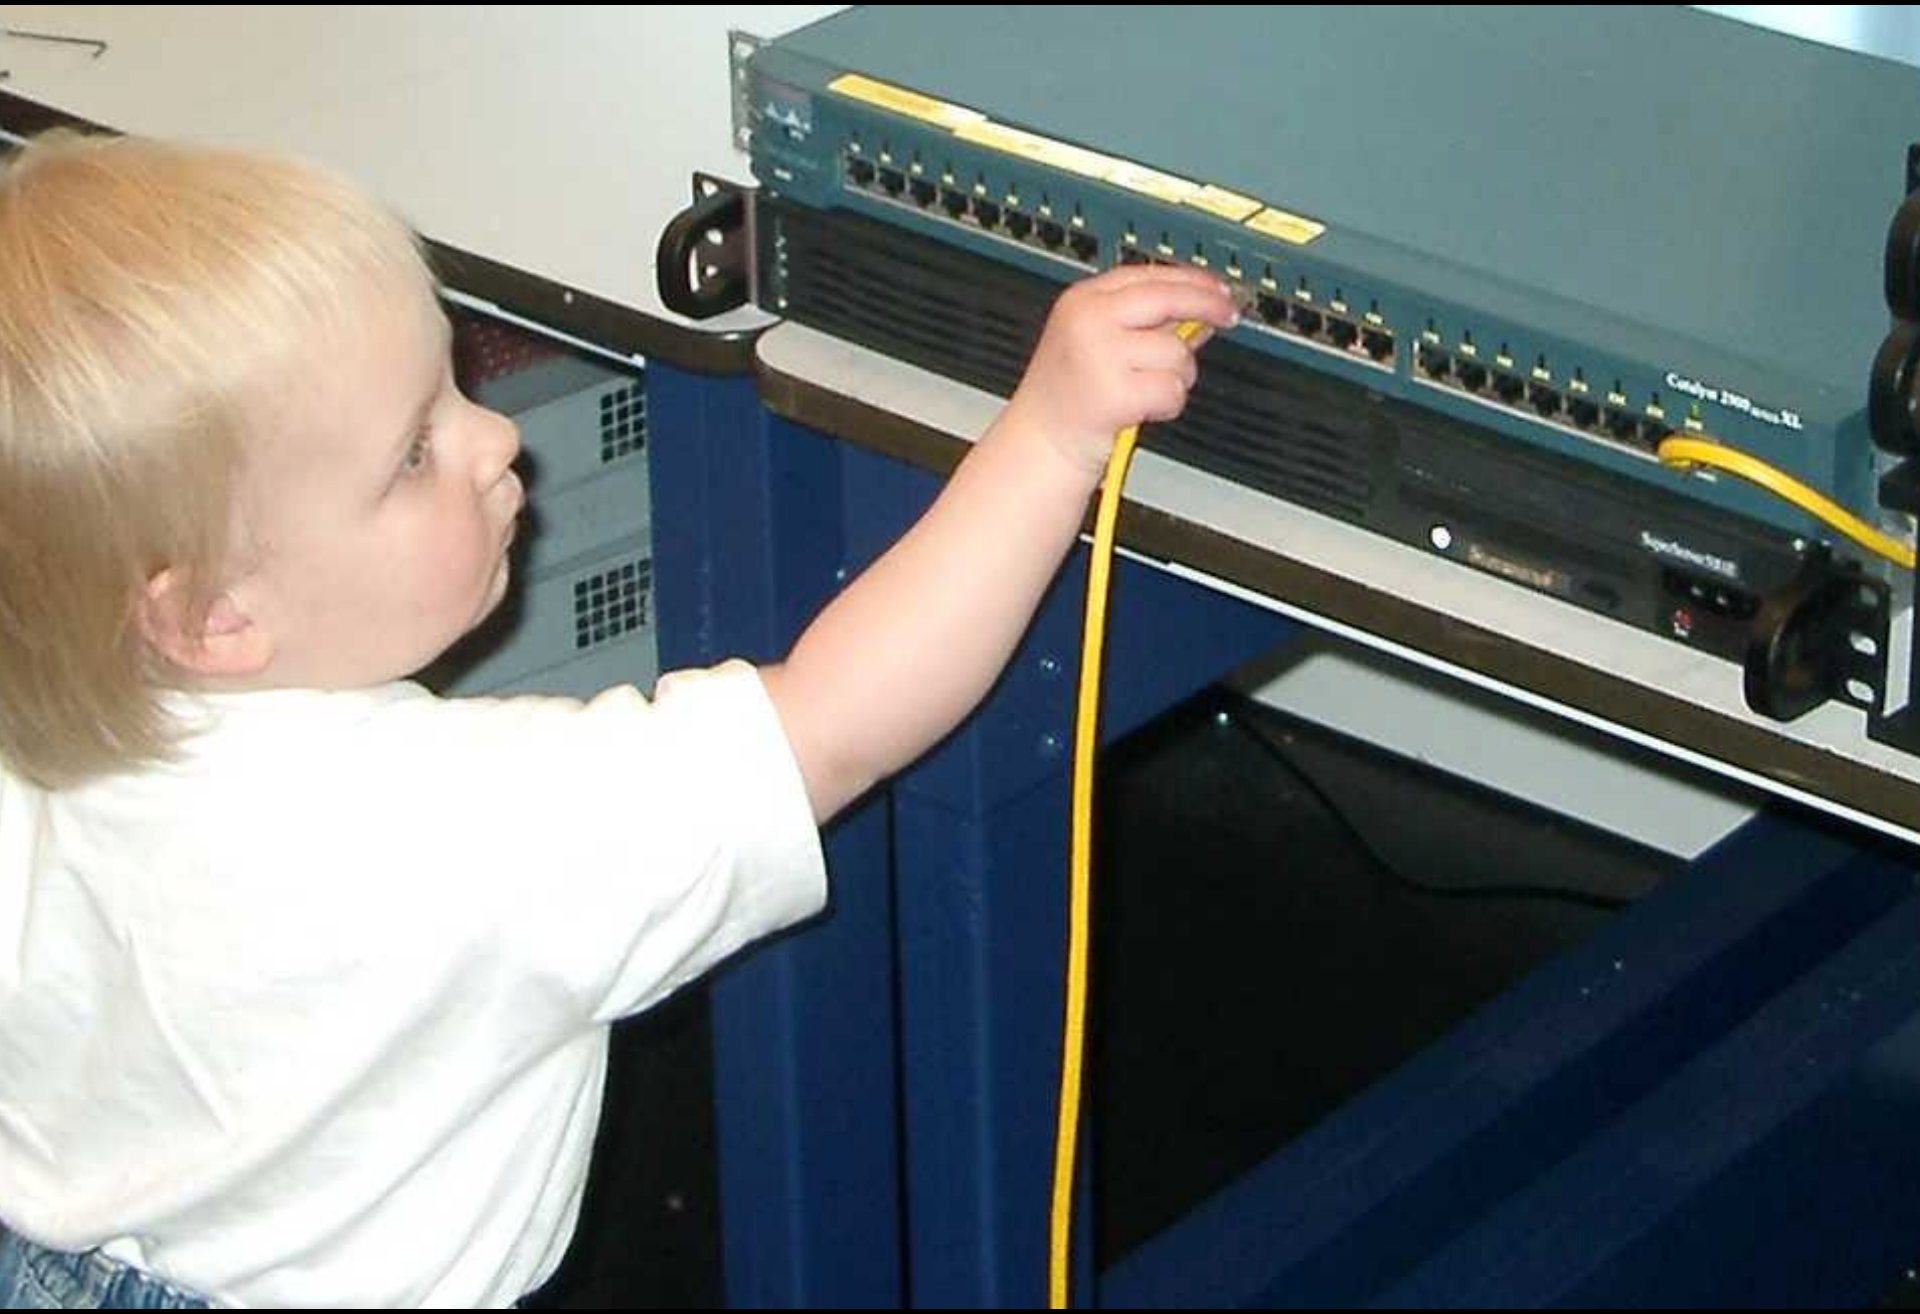

Questions ?

Slides, paper, scripts and example files:
http://www.xs4all.nl/~scorpio/sane2002

Email me at cor@xs4all.net