Redpill Linpro

Monitoring at The Gathering

The Whatnow?

- The Gathering is a yearly digital festival, attracting 6000+ people a year during the easter holiday.
- Typical participant is between 15 and 20 years old
- Gaming, creative competitons, booths, socializing, etc.
- Bring your computer (or gaming console, or whatnot)
- Non-profit

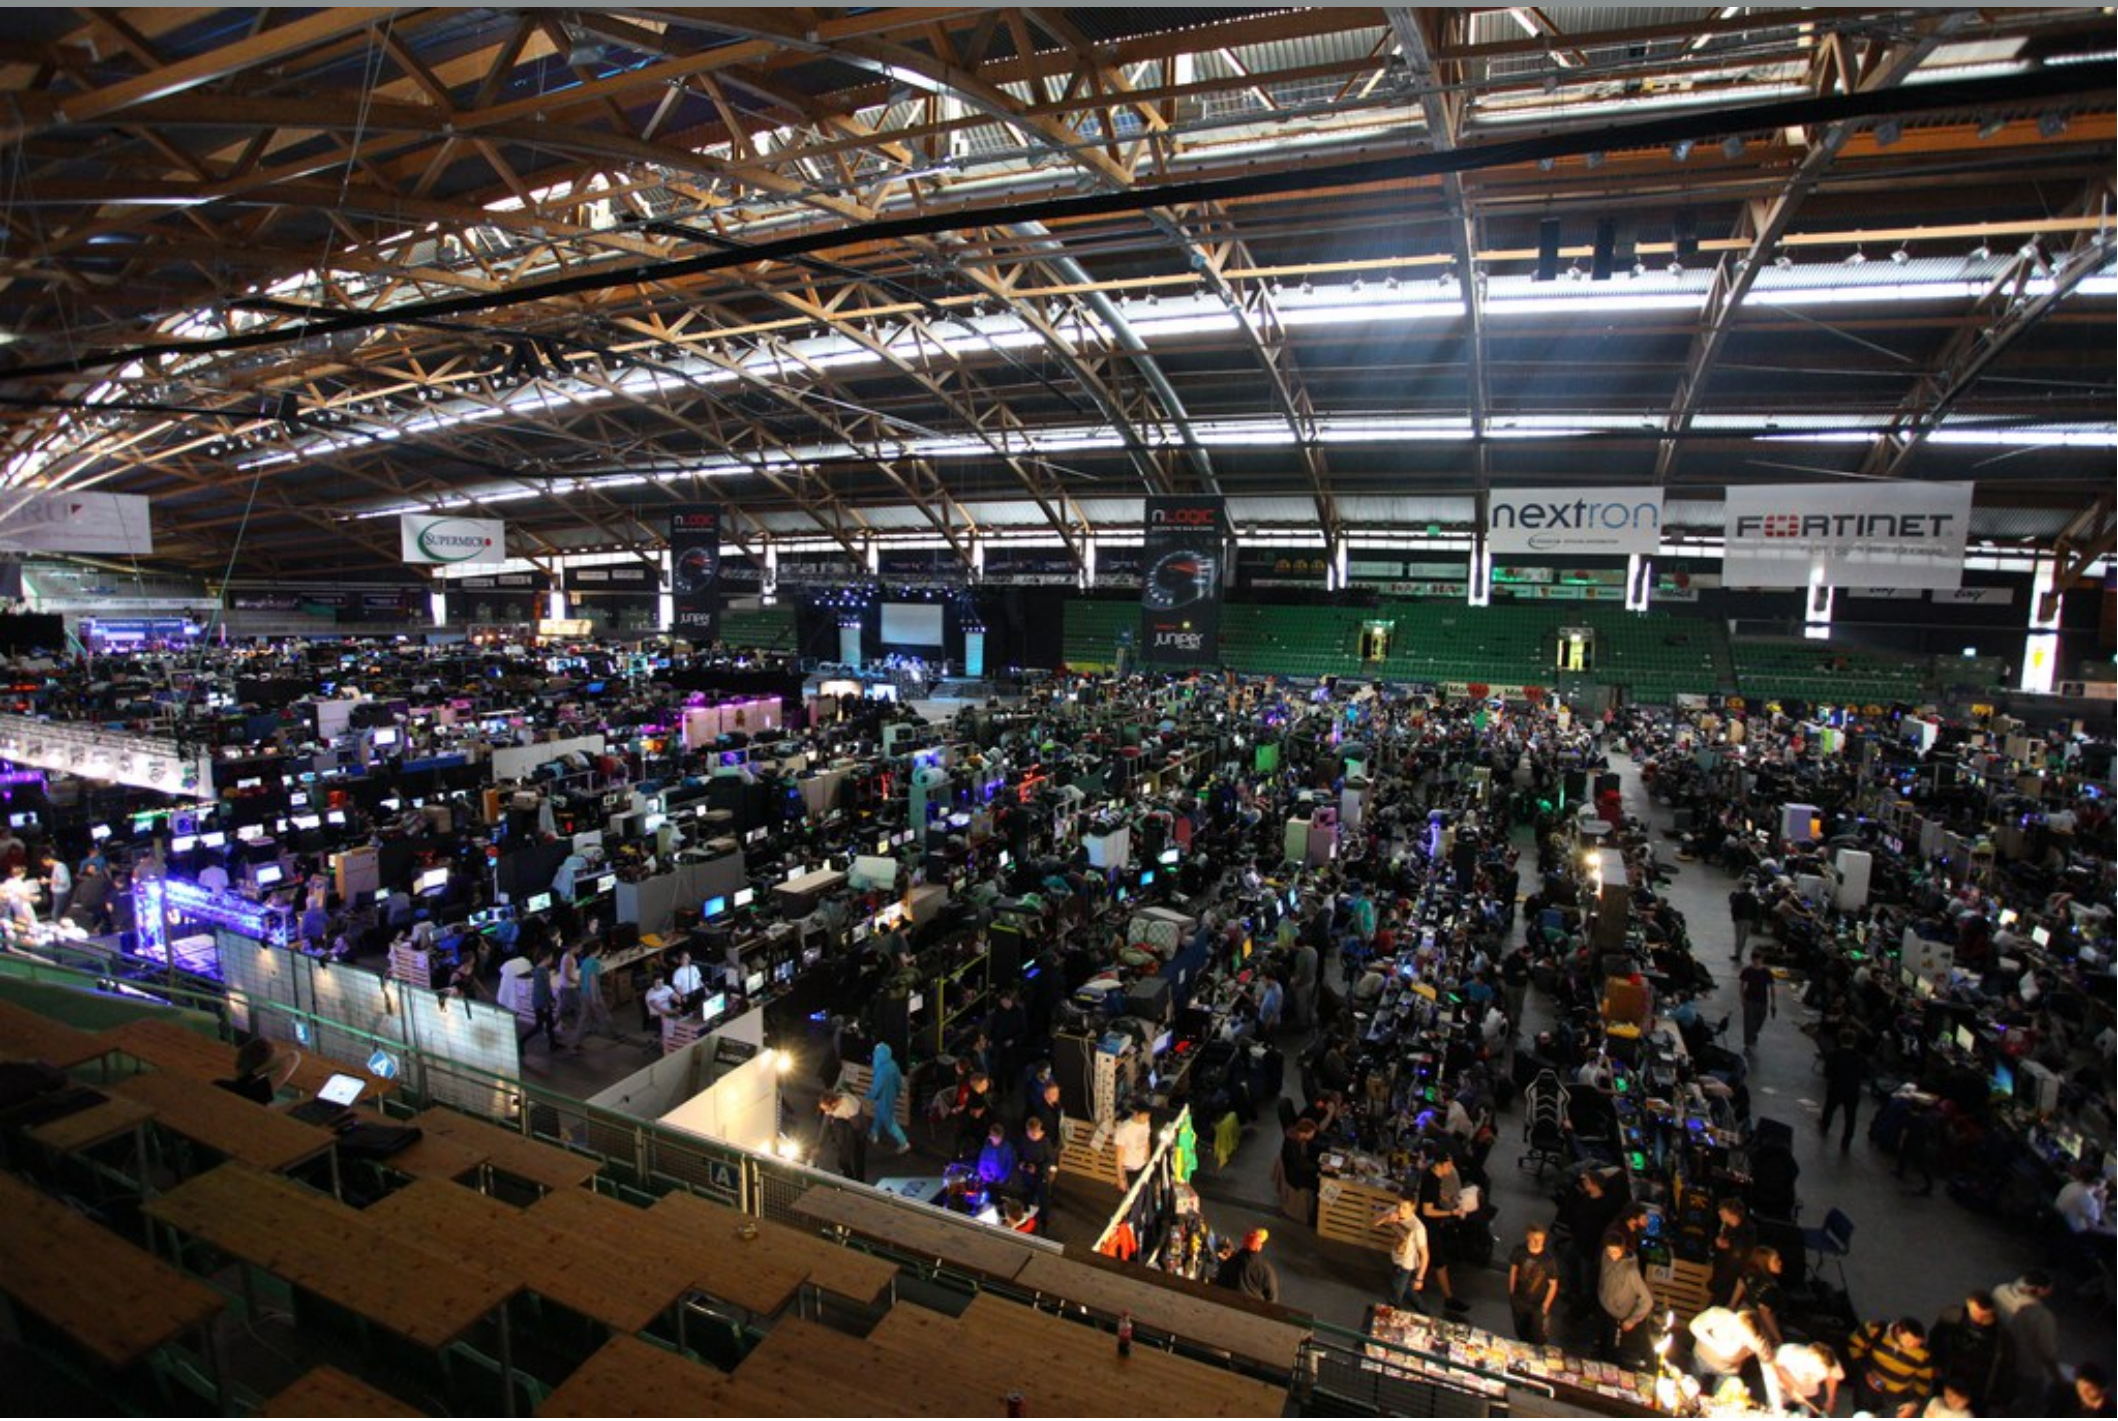

PRODUCTS • CONSULTING • APPLICATION MANAGEMENT • IT OPERATIONS • SUPPORT • TRAINING

Redpill Linpro

PRODUCTS • CONSULTING • APPLICATION MANAGEMENT • IT OPERATIONS • SUPPORT • TRAINING

Redpill Linpro

Technical breakdown

Redpill Linpro

200 or so access switches.
Juniper EX2200 48P

PRODUCTS • CONSULTING • APPLICATION MANAGEMENT • IT OPERATIONS • SUPPORT • TRAINING

Redpill Linpro

9 (or so) distribution switches.

3x Juniper EX3300 (in virtual chassis)

We own about 50+ EX3300s

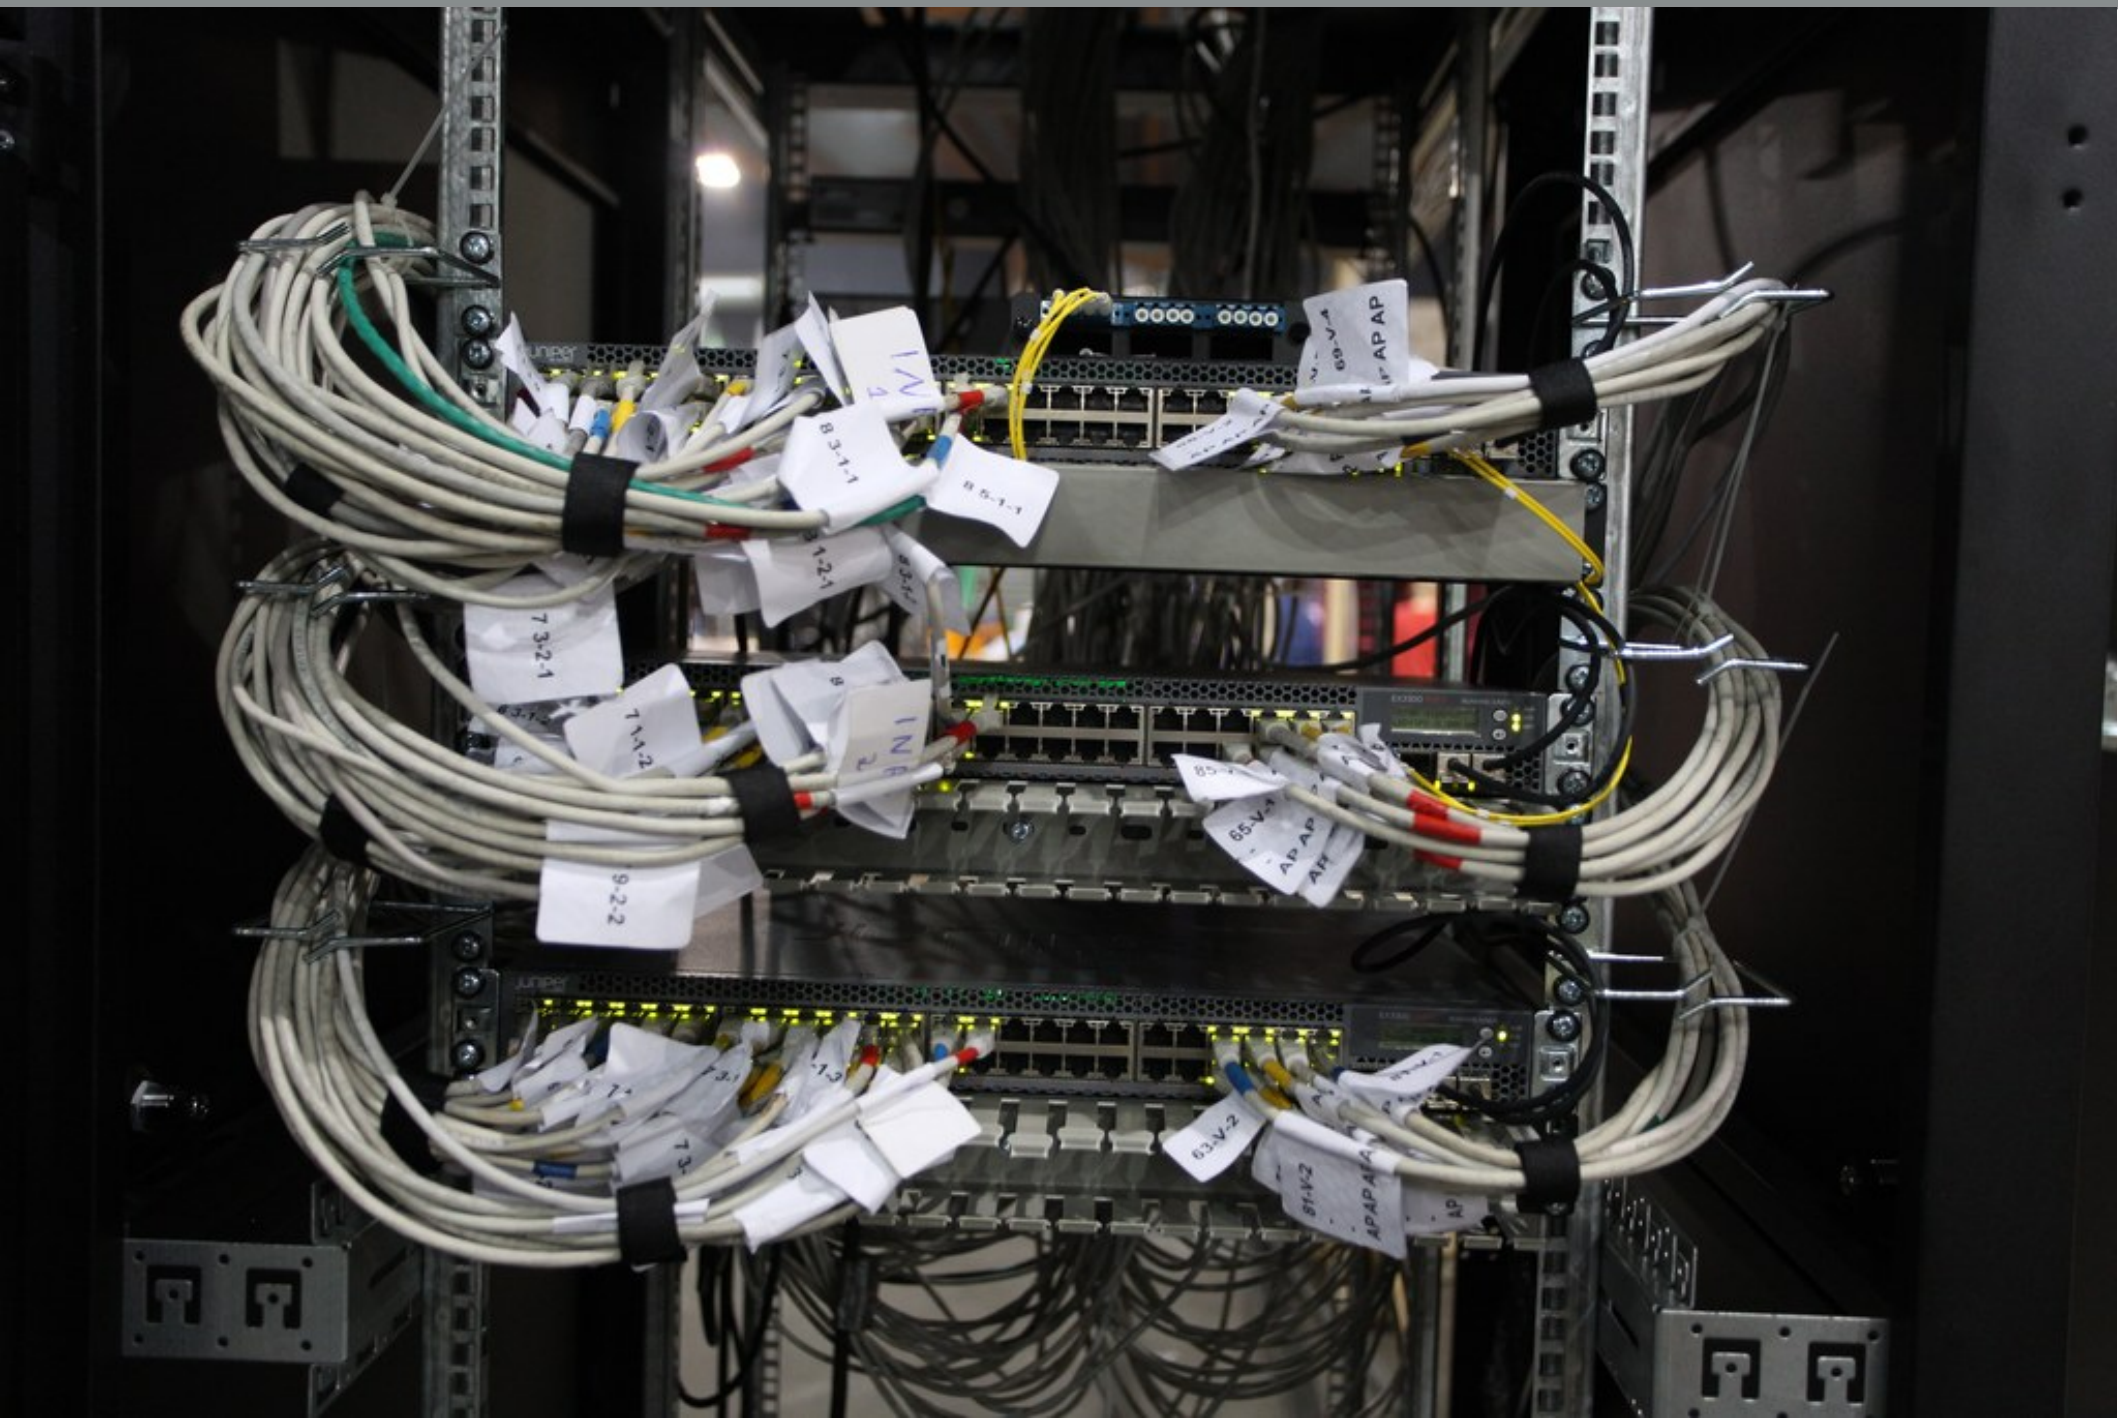

PRODUCTS • CONSULTING • APPLICATION MANAGEMENT • IT OPERATIONS • SUPPORT • TRAINING

Redpill Linpro

Redpill Linpro

Meru wireless. A few hundred Access points (I honestly don’t know the specs)

Redpill Linpro

Etc

- Peaks at about 7000 live DHCP leases with unique MAC addresses
- Total of about 10500 unique MAC addresses seen during The Gathering 2016.
- Core switch varies from year to year (in the roof – literally)
- Redundant ring network mainly consisting of EX3300’s
- A few EX4300’s here and there for gateways and whatnot

Redpill Linpro

Servers: Vastly overpowered Supermicro boxes on loan from Nextron

“We didn’t have a suitable 15TB storage server for you, so we sent you this 128TB one instead”

PRODUCTS • CONSULTING • APPLICATION MANAGEMENT • IT OPERATIONS • SUPPORT • TRAINING

PRODUCTS • CONSULTING • APPLICATION MANAGEMENT • IT OPERATIONS • SUPPORT • TRAINING

Redpill Linpro

Unused:
Two QFX10002-72Q Juniper “Switches” (24 100GBit ports or 72 40GBit ports)

Price tag: 2x “oh god, oh god, I hope I don’t drop this”

PRODUCTS • CONSULTING • APPLICATION MANAGEMENT • IT OPERATIONS • SUPPORT • TRAINING

Redpill Linpro

Core: Juniper qfx5100 (this year)
Sits about 40 meters above the main floor. Literally.

Redpill Linpro

I just do server/development stuff. Computer things. Beep boop.

Monitoring: Gondul

- Open Source
- Small, simple components
- Home brew
- Used at other parties too (Digitality X in Stavanger, Sørlanet in Kristiansand, probably ByLAN in Bergen)

Redpill Linpro

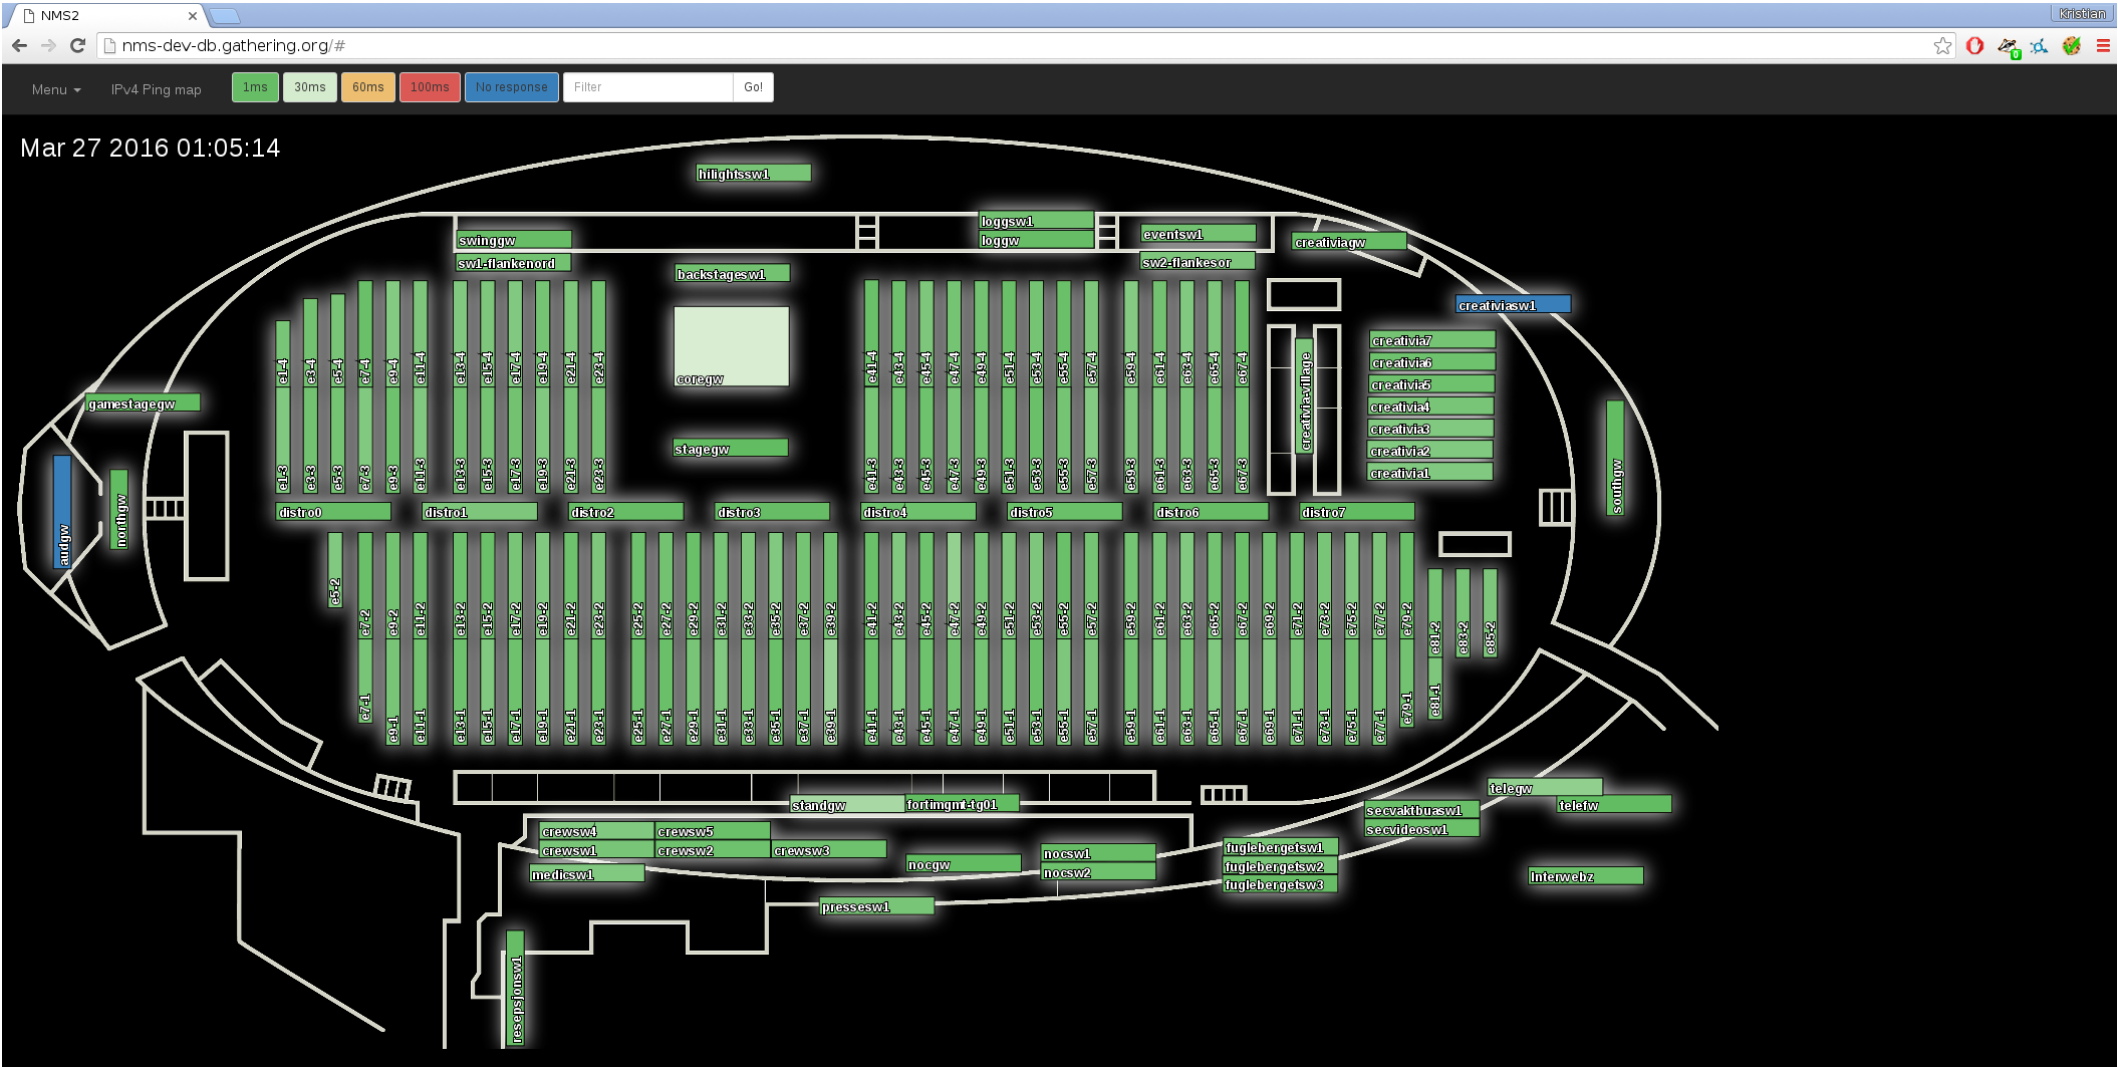

Two phases: 1. Pre-production

- Less than a week on-site to set everything up
- No point in graphs or alerts. But need the overview.
- Pre-production is also when we get to test new features for the first time and tweak them if need be.
- Other crews need network too (logistics, event, info…). Some is covered with “Just use 4G”, some is covered by us.

Gondul helps with...

- Tracking progress
- Catching regressions
- Configuring access switches
- Catching sillyness (...)

Two phases: 2. Production

- Less than a week of production
- We open at 9AM. At 9PM we have 7000 units on the network.
- 20 minutes of downtime is a big deal.
- This year: At 8AM we were short 1 distro switch (so roughly 700 participants)
- Alerting is largely redundant. (*)
- Trending and metrics are interesting for “bling effect”
- Typically 2-4 feature requests per 8-hour shift.

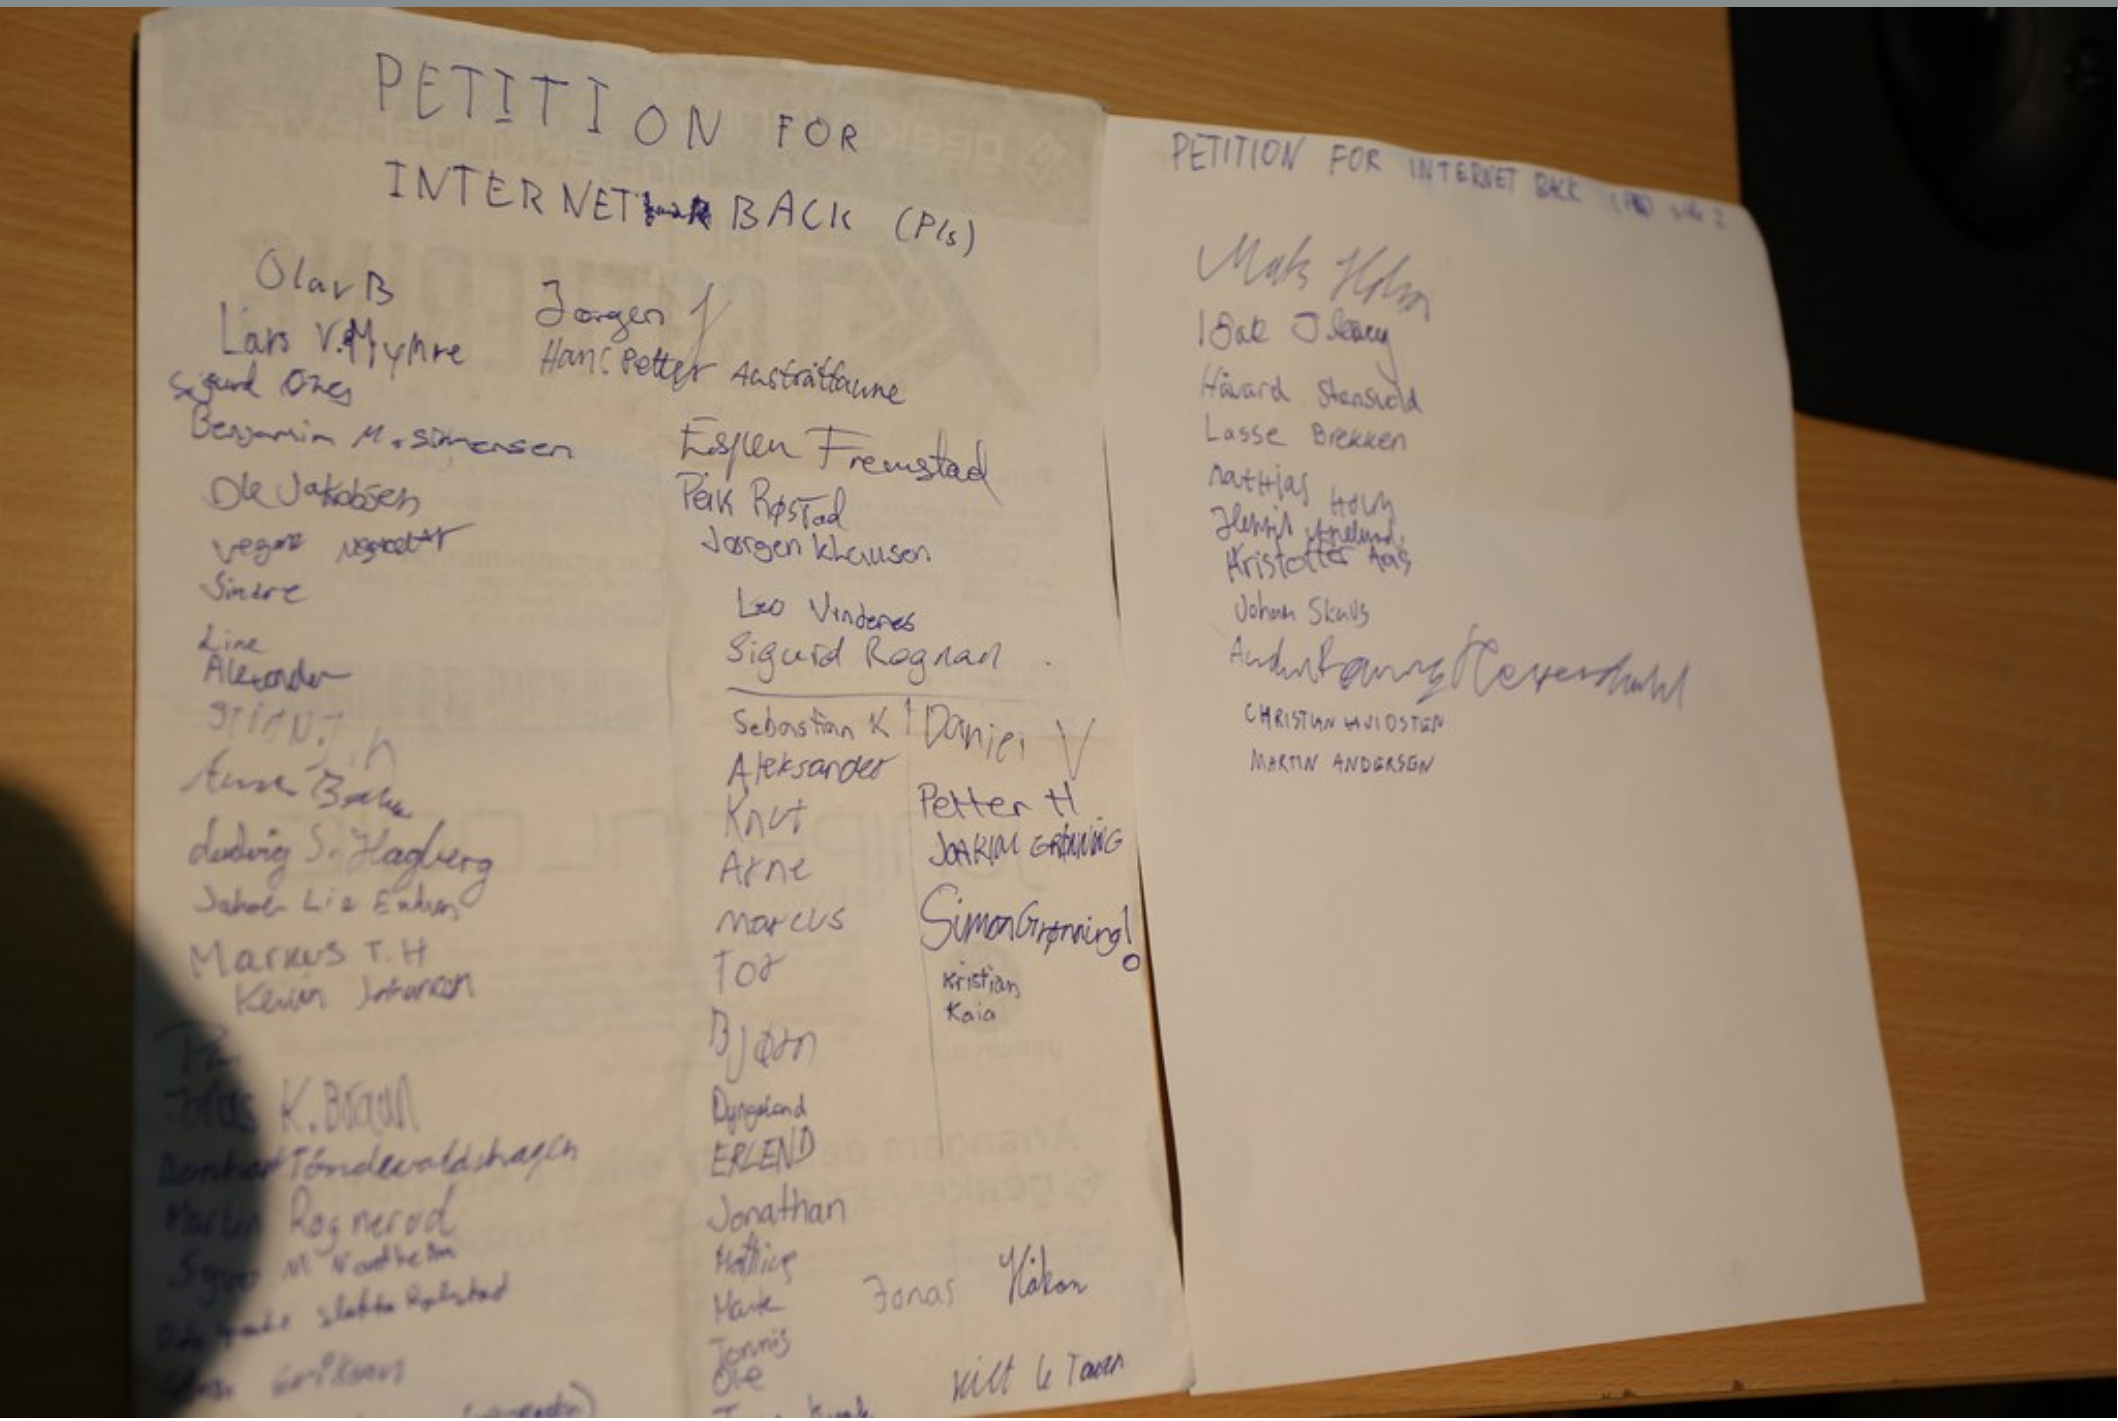

Gondul helps with...

- Catching when a participants unplugs an uplink port or plugs into a spare uplink port
- Catching when an entire section is missing DHCP (why?)
- Detecting bandwidth issues (e.g.: access switch needing an other uplink)
- Quickly detecting when switches go down
- Configuring access switches used to replace “broken” switches
- Debugging when things just don’t work (that NEVER happens)

Redpill Linpro

Instant insight

Instant insight

- Unplug a switch and we'll know it before you put the cable down. Assuming we're watching.
- The goal: A single map of the entire infrastructure.
- Nobody needs any introduction to the system beyond login/password and an instruction to go nuts.
- Click a switch / server for more information.

Deployment

- Planning (e.g.: Ip ranges etc)
- planning.cpp – makes final decisions on IPs for access layer. Run once. Also generates patch plans for distro switches and calculates ideal cable lengths to minimize number of cables.
- Feed planning-output to Gondul
- Feed Gondul-information to scripts that generate DHCP configuration
- Access switches get configuration over DHCP + Gondul
- Distro and core has hand-written(typed) configuration

How?

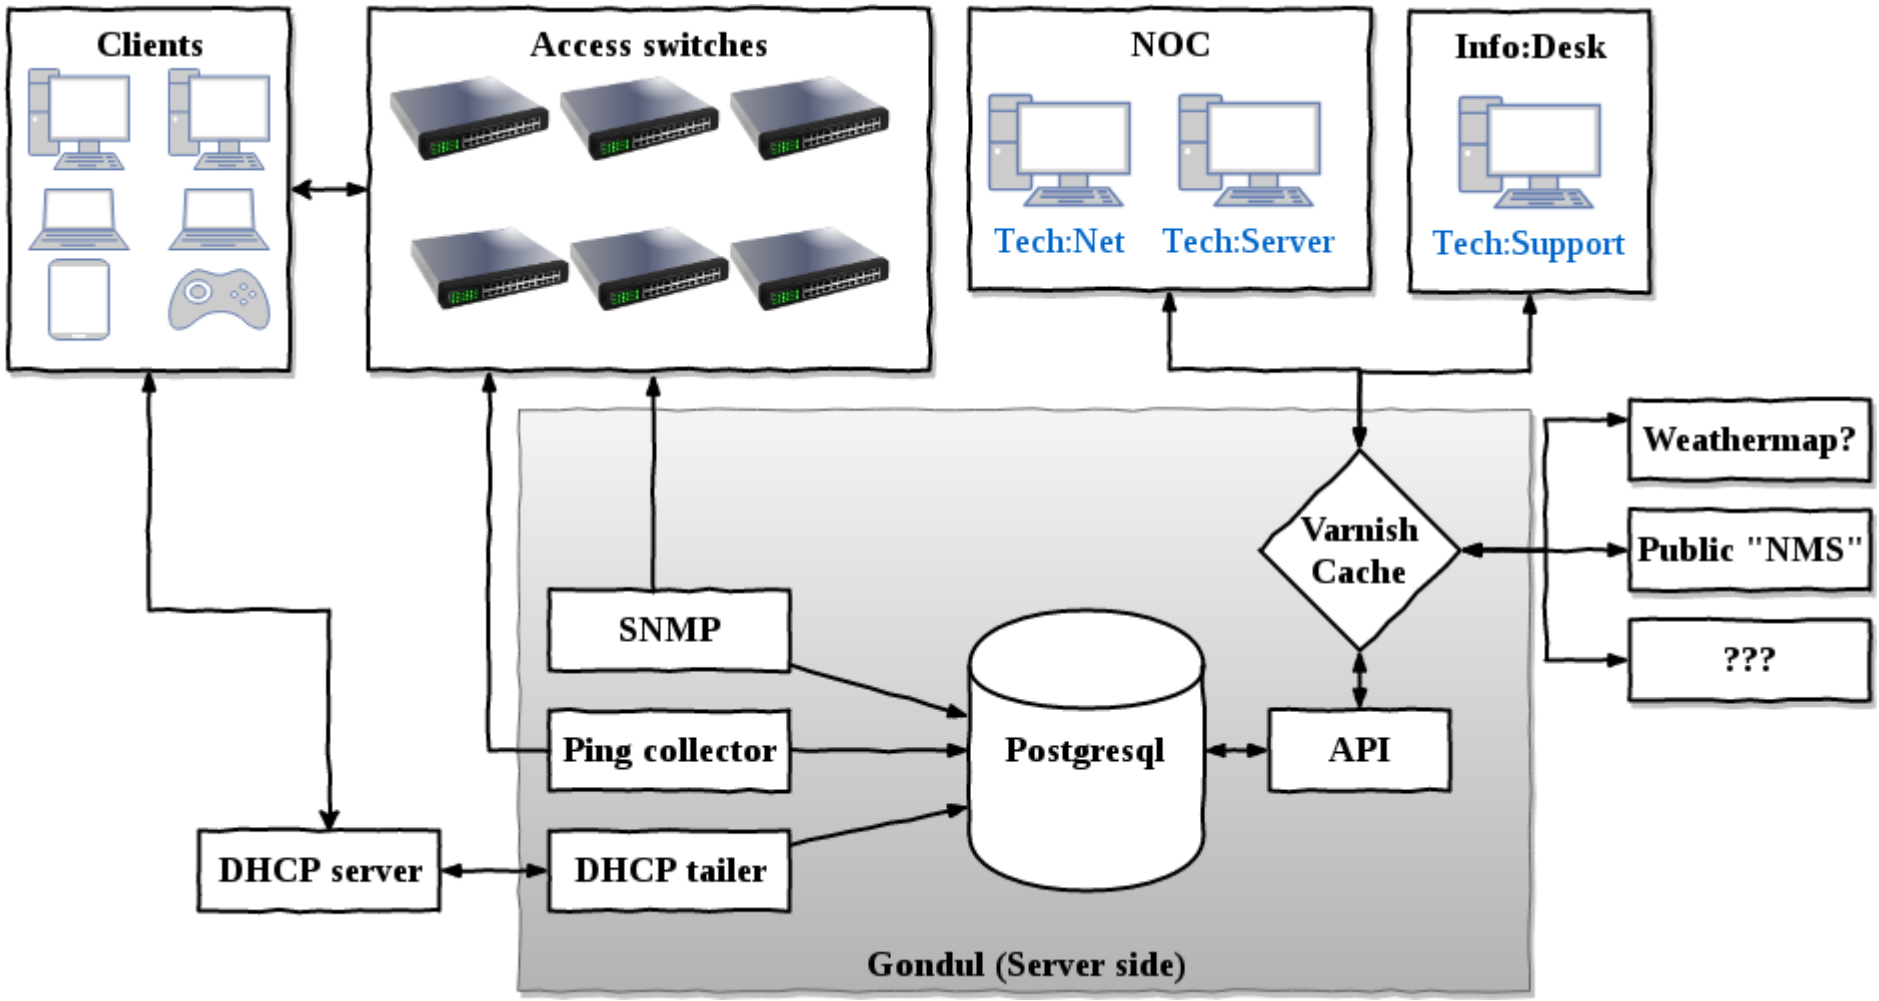
Clients
Access switches
NOC
Tech:Net
Tech:Server
Info:Desk
Tech:Support
Weathermap?
Varnish Cache
Public "NMS"
???
SNMP
Ping collector
Postgresql
API
DHCP server
DHCP tailer
Gondul (Server side)

Redpill Linpro

Building: Docker and ansible

- Run some things in docker (but not all – docker still doesn’t do IPv6 sensibly)
- Ansible to orchestrate it and provide developer environment
- Postgres runs outside of docker on production

Collectors: Simple, fast, flexible

- Ping collector is by far the most critical: If it stops, we assume the net is down. Needs to “just work”.
- SNMP collector can be a performance challenge. Can be run in paralell. Needs to be easily (re-)-configured.
- DHCP log tailer equally simple. If we mess it up, we can just delete the entire DHCP table and restart the DHCP collector – it will re-read the log as if nothing was wrong.

Backends: Simple, generic

- Backend endpoints are really simple: Just provide data and don’t try to be too smart about it.
- Tons of data ….
- Some central scaffolding to add cache headers and a general structure.
- Idempotent writing where relevant: Changing and adding a switch is the same thing.

Cache Layer: Generic, smart

- Backend sends cache-control headers
- Cache layer doesn’t have any overrides (except to allow caching content for authenticated users)
- 99.999% cache efficiency – despite short cache durations.
- ETag, max-age, stale-while-revalidate… All used both by Varnish and the browser.
- Each client adds 2req/s or thereabouts
- Server load is almost completely isolated from how many clients we have. 1 client or 1000 makes no difference.

Frontend: All client-side logic

- All logic is in the front-end.
- The browser always has basically all information available about the infrastructure.
- Intelligent modularization. Business logic separated from data-handling, graphics, etc
- Multiple HTML Canvas
- Development is easy and fast: Testing on a laptop then deploying to live – basically no risk involved.

Redpill Linpro

Replaced/rewritten/refactored n 2.5 years:

SNMP walker.
Ping daemon.
DHCP tailer.
Build system.
Entire frontend.

Graphing removed, done externally, redone, etc.
(list goes on)

Redpill Linpro

Simple stuff is simple to replace.
So replace it.

Redpill Linpro

KISS