Redesigning an OOB Network for Resilience

Moritz Frenzel, RIPE86, Rotterdam

Globalways

whoami

Moritz Frenzel

- CEO @ Globalways GmbH
- Vice Chairman @ DENOG e.V.
- Network Architect @ Stuttgart-IX

Most of the work: Michael Meyer, Senior Network Architect @ Globalways GmbH

Globalways

Globalways GmbH AS48918

- Offering L2 and L3 Services for B2B Customers
- >360km of own dark fiber in Stuttgart Germany
- 700+ connected buildings
- 19 POPs – 15x STR, 2x FRA, 1x BER, 1x MUC
- Core & Access: ASR900X, QFX51XX, MX204
  - OSPF(v3), iBGP, LDP, MPLS
- CPE: EX3300, FSP150
  - >1.000 devices

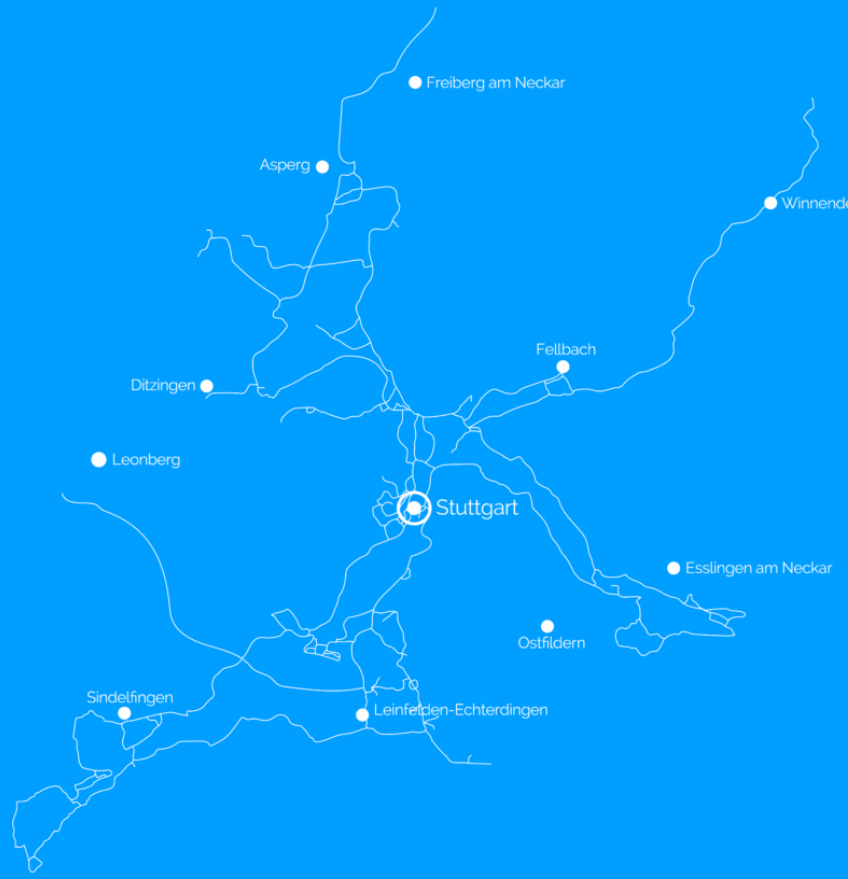

Globalways

Globalways GmbH AS48918

- Offering L2 and L3 Services for B2B Customers
- >360km of own dark fiber in Stuttgart Germany
- 700+ connected buildings
- 19 POPs – 15x STR, 2x FRA, 1x BER, 1x MUC
- Core & Access: ASR900X, QFX51XX, MX204
  - OSPF(v3), iBGP, LDP, MPLS
- CPE: EX3300, FSP150
  - >1.000 devices

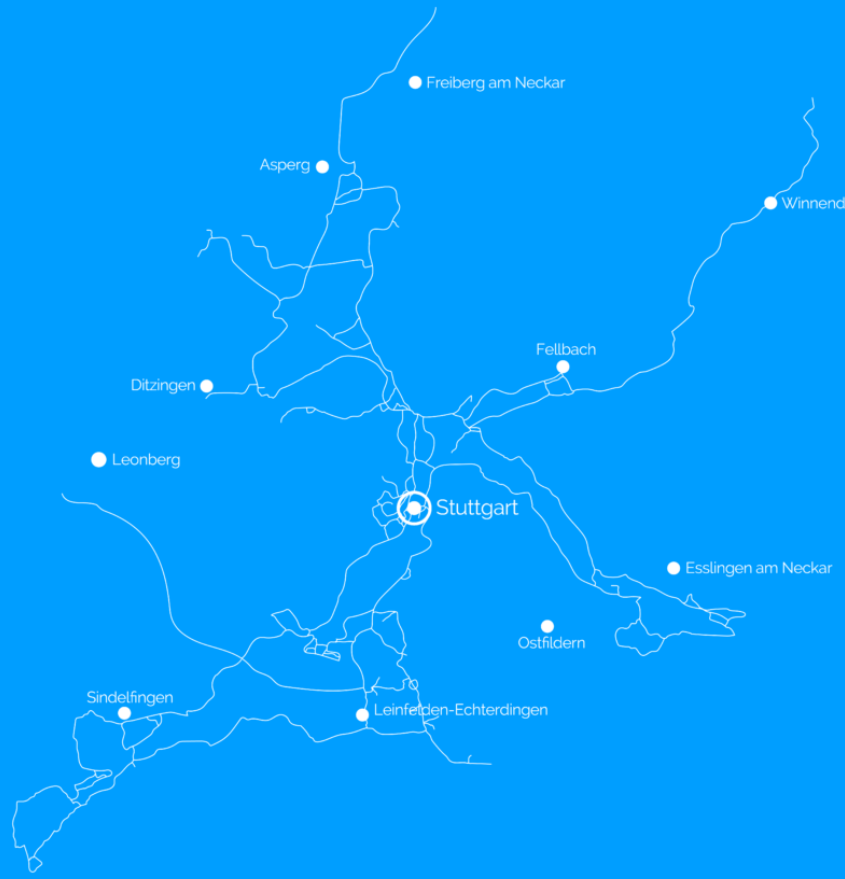

Globalways

Old OOB Network

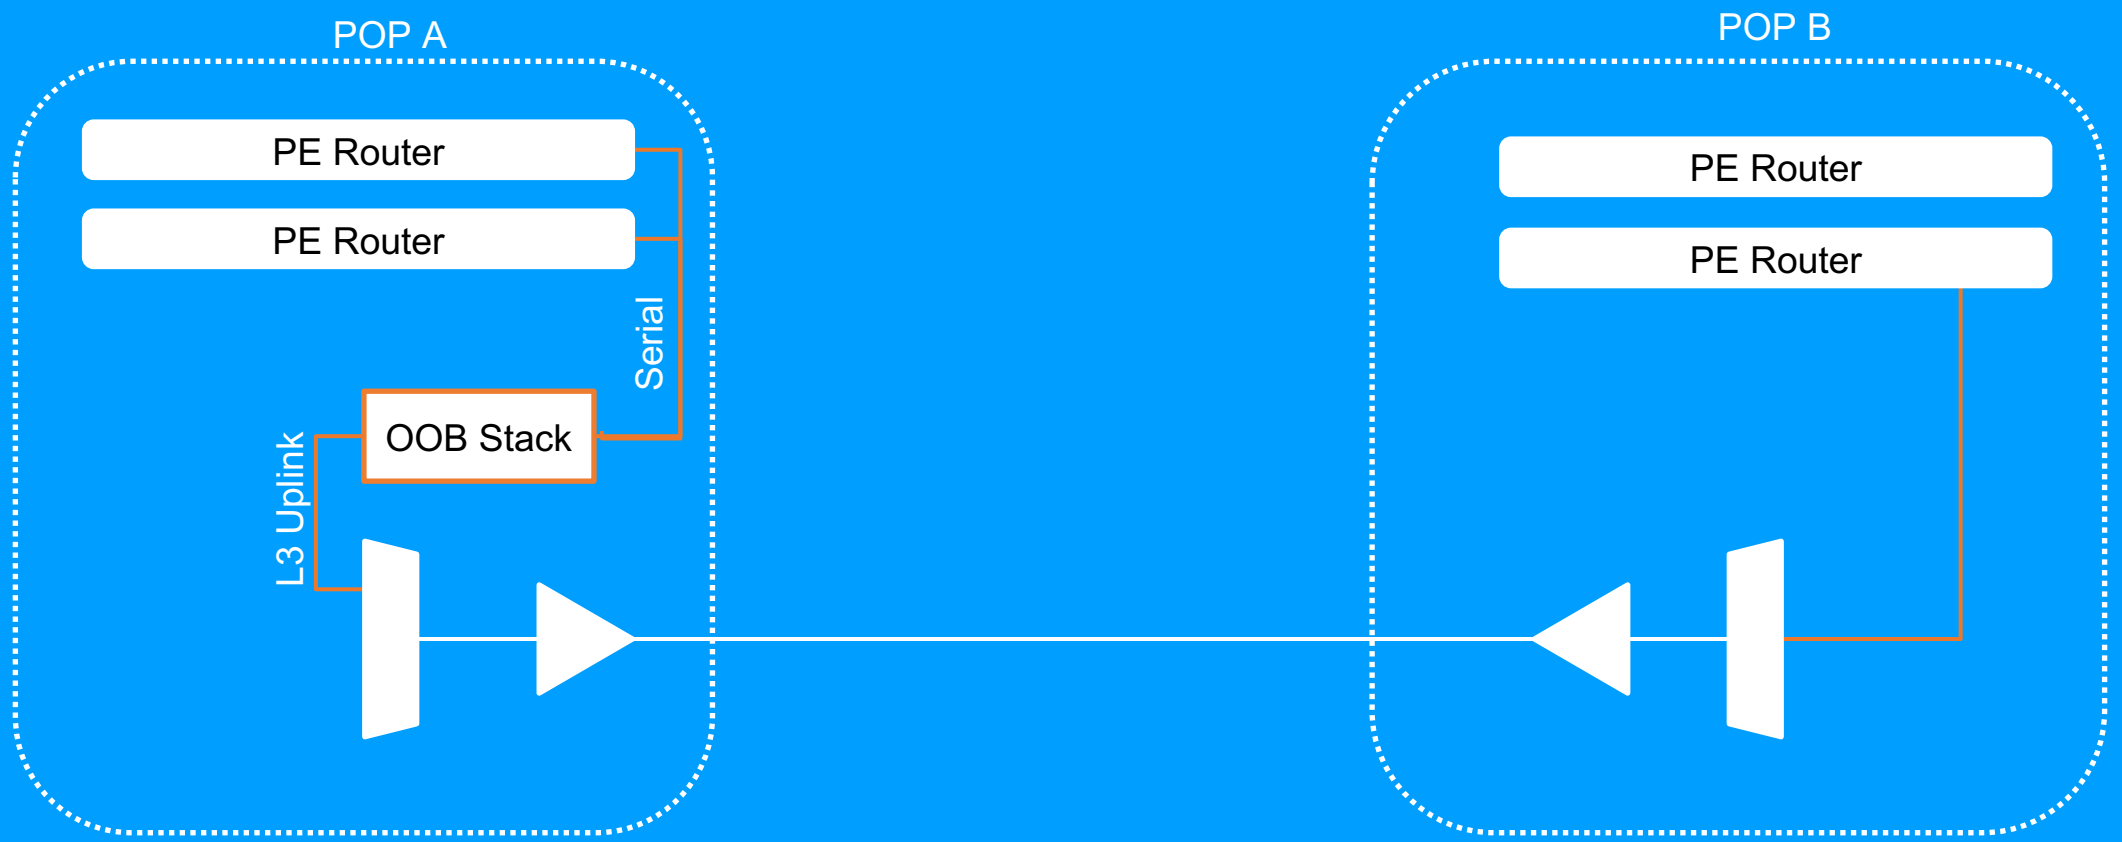

Globalways

Old OOB Network

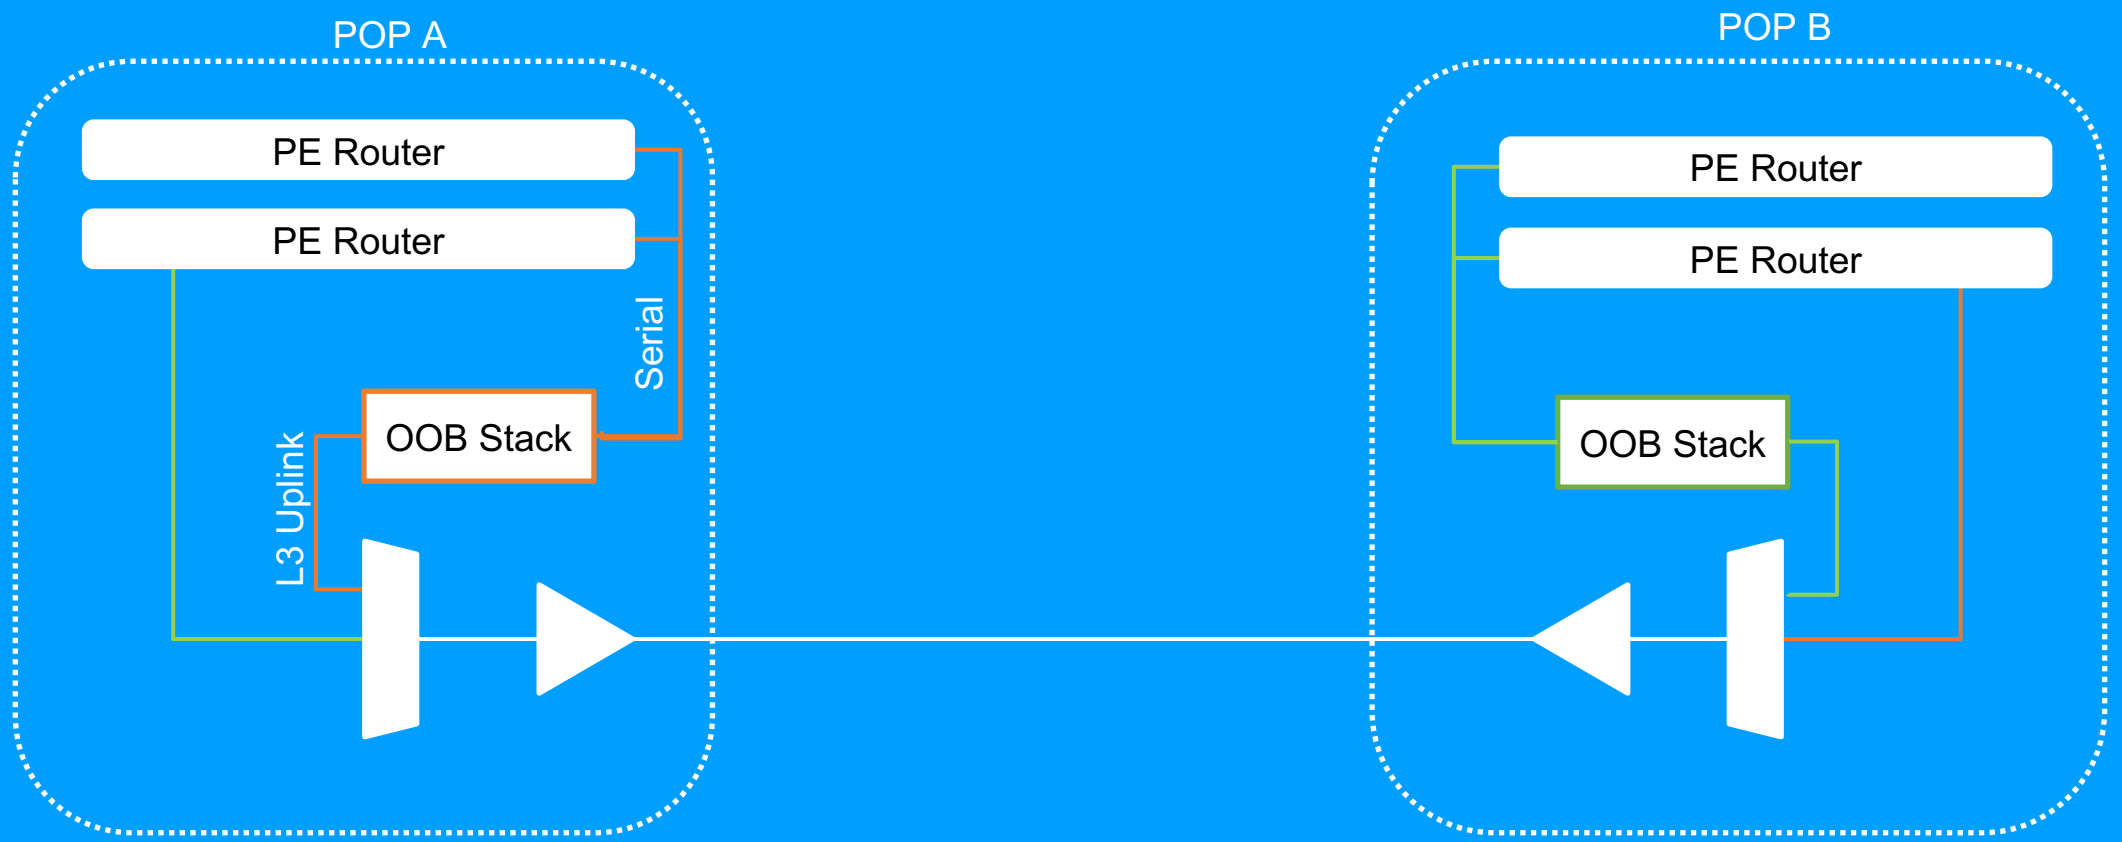

Globalways

Old OOB Network

- Hardware
  - Juniper EX3300-24T(-DC)
    - Provides ports for UPS, Temp monitoring, …
  - Cisco 2509 / 2511 as console Server
    - Yes, with modern 10m half-duplex via AUI
  - PCEngines APU2-C4 running debian

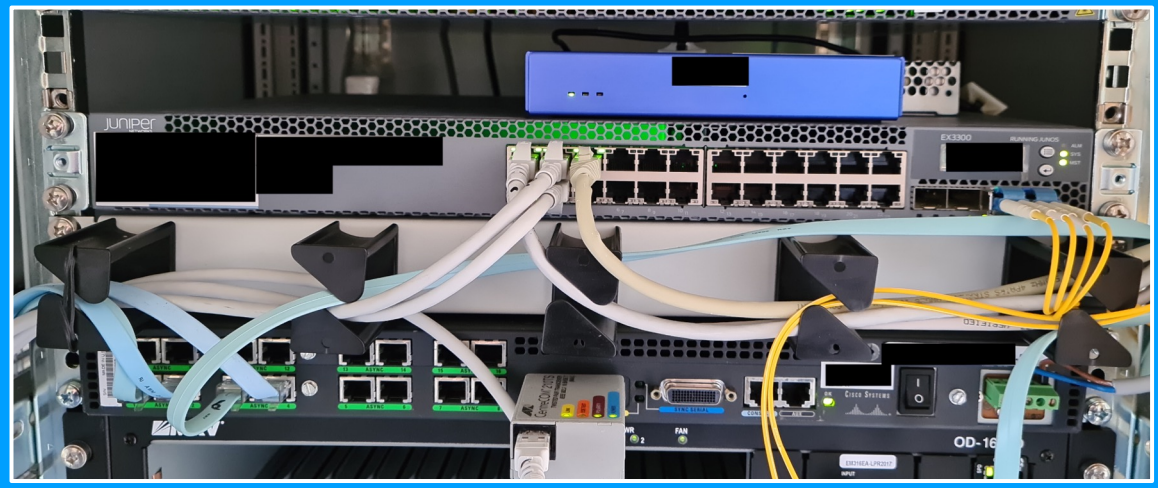

An old OOB network installation

- Connectivity
  - DWDM Wave or DF to the next POP
    - To ensure connectivity if the local routers fail
  - Where feasible: circuit from other providers
  - OpenVPN to a virtual (redundant) concentrator (also debian)

Globalways

Issues

- 3 devices to manage & maintain
  - 3 wasted RUs
  - 100W+ power draw
  - 3 separate DC fuses
- VPN failover is not reliable, hassle to maintain & orchestrate
  - If the VPN concentrator fails, everything goes dark
- Link to other POPs requires working EDFAs
- Not resilient against catastrophic IGP failure

Design Goals

- Reduce footprint
- Reduce operational toll
- Reduce power draw
- No central VPN concentrator
- Resilience against catastrophic routing failures
- 3G/4G/5G for OOB access for sites with no 2nd provider
- Add perimeter security

Globalways

opengear ACM7004-5-L

- 1x SFP or 1GBaseT Uplink
- Dual SIM slot & 4G capability
- 4x Ethernet Ports
- 4x Serial Ports
- 4x USB A, which can use FTDI-Cables
- 2x Digital I/O ports
- Linux Based with full CLI access
- 11.5W power consumption

Source: opengear.com

Globalways

Cellular Connectivity

- Most of our POPs are in underground train stations
  - Luckily, most of them also house 4G/5G base stations of various operators
  - Site survey determined that no carrier is available in all locations
- Introducing: whereversim.de
  - 2G, 3G, 4G & LTE-M, 5G in roll-out
  - One SIM for all German cell service operators
    - Automagically selects the best carrier
  - Data Pooling
  - IPSec & Private APN upon request

wherever SIM

Source: whereversim.de

Globalways

VPN

- We were generally evaluating new VPN solutions for us
- OpenVPN was just too much work to maintain
  - We’re an ISP, after all, not a VPN operator
- WireGuard full mesh sounds perfect
  - Good:
    - Point-to-point connections remove the need for a concentrator
    - Really fast speeds as of Kernel 5.6 without any tuning
    - Reduced complexity compared to OpenVPN
  - Meh:
    - Key Management is still a thing
    - Commercial support
    - Still needs $firewall for ACLs
    - Full mesh config requires clever automation

tailscale

- WireGuard as data plane, tailscale as control plane
- The control plane ticks all our boxes:
  - Magically builds the full mesh
  - Automatic key rotation
  - Simple, JSON-Based ACLs
  - Audit-Compliant logging
- Yes, it costs money, but we've never had to worry about VPN since.
- Yes, we rely on some cloud software.
- Yes, RFC6598 might be a thing
- Most of it is open source.

Source: tailscale.com

Globalways

Putting it all together

- Installing tailscale on opengear was fairly easy
  - No Kernel support, but we don’t need speed
- Whereversim just required a custom APN to be set up
- Add in some python & ZTP to populate everything from netbox
- Monitor everything via snmp_exporter
- File an opengear bug as RFC3021 is not supported
  - Fixed in 4.13.2

A new ACM5004-5-L at I/P/B Lützowstr. 105

Globalways

Grafana Dashboard

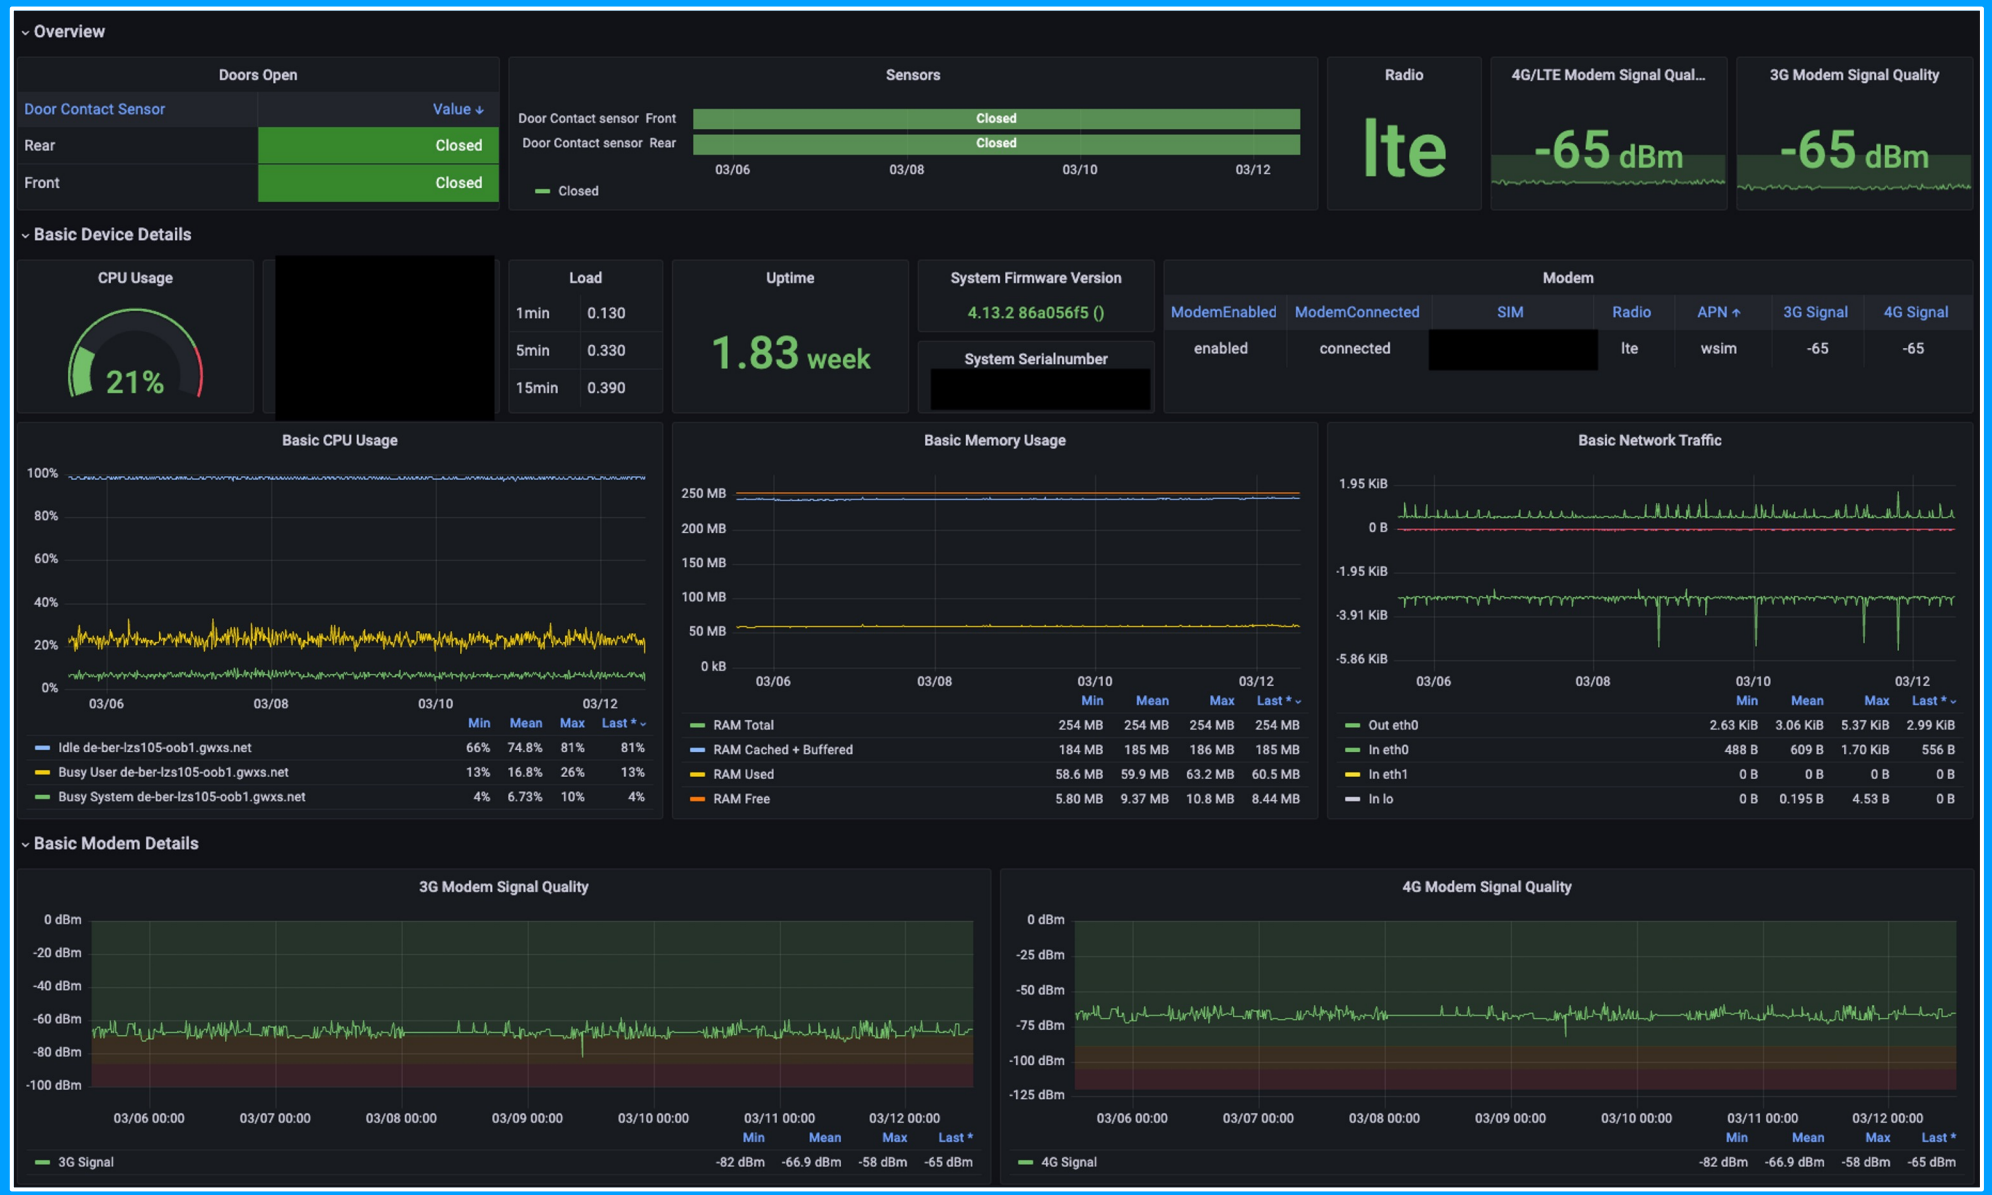

Globalways

Lessons learned

- Tailscale is amazing, opengear just works, and whereversim just works
- A fully meshed VPN is amazing, even in a catastrophic IGP failure scenario
- Door Contacts are great for verifying your remote hands bills
- An ASR9001 can crash when you add a faulty cable to its serial port
- 88.5W less x 19 POPs = 1.68kW less
  - That’s ~6.2t of CO2/year (assuming 420g/kWh)
  - also 7.300,00€/year (assuming 0,5€/kWh)
  - Hardware Break-even in < 3 years

Next Steps

- Complete the rollout
- Opensource documentation for tailscale on opengear
- Opensource snmp_exporter config and sample dashboards
- Maybe opensource our console tool
  - Uses netbox to find existing correct oob device and directly connects

DENOG

GERMAN NETWORK OPERATORS GROUP

DENOG15
November 19 – 21 2023
Berlin

Thank you!

Questions?

Moritz.Frenzel@globalways.net

Globalways