# How hard could it be to greenfield a high capacity network?


hugge@nordu.net

RIPE75

- **Fredrik Korsbäck**
  - hugge
  - ~~29~~ 30
  - Stockholm
  - Network Architect for 15yrs
  - Attending all the meetings
  - First RIPE meeting!?

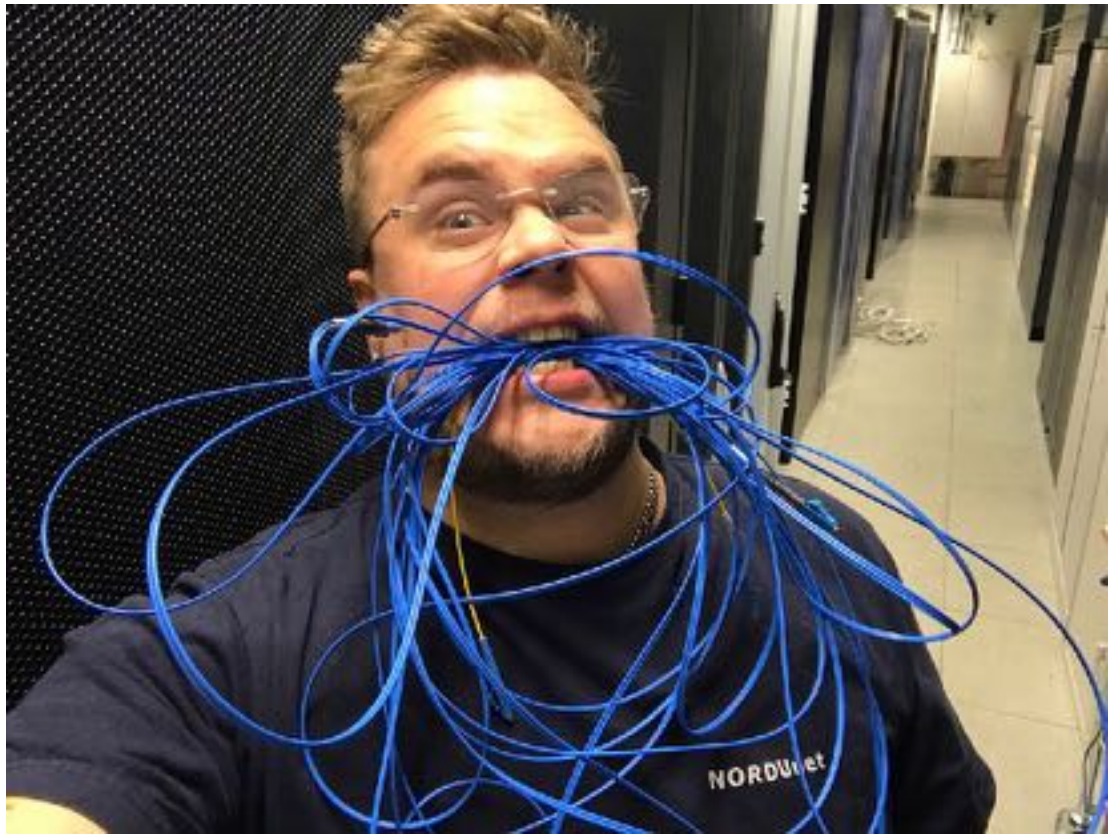

# Sweden

IKEA

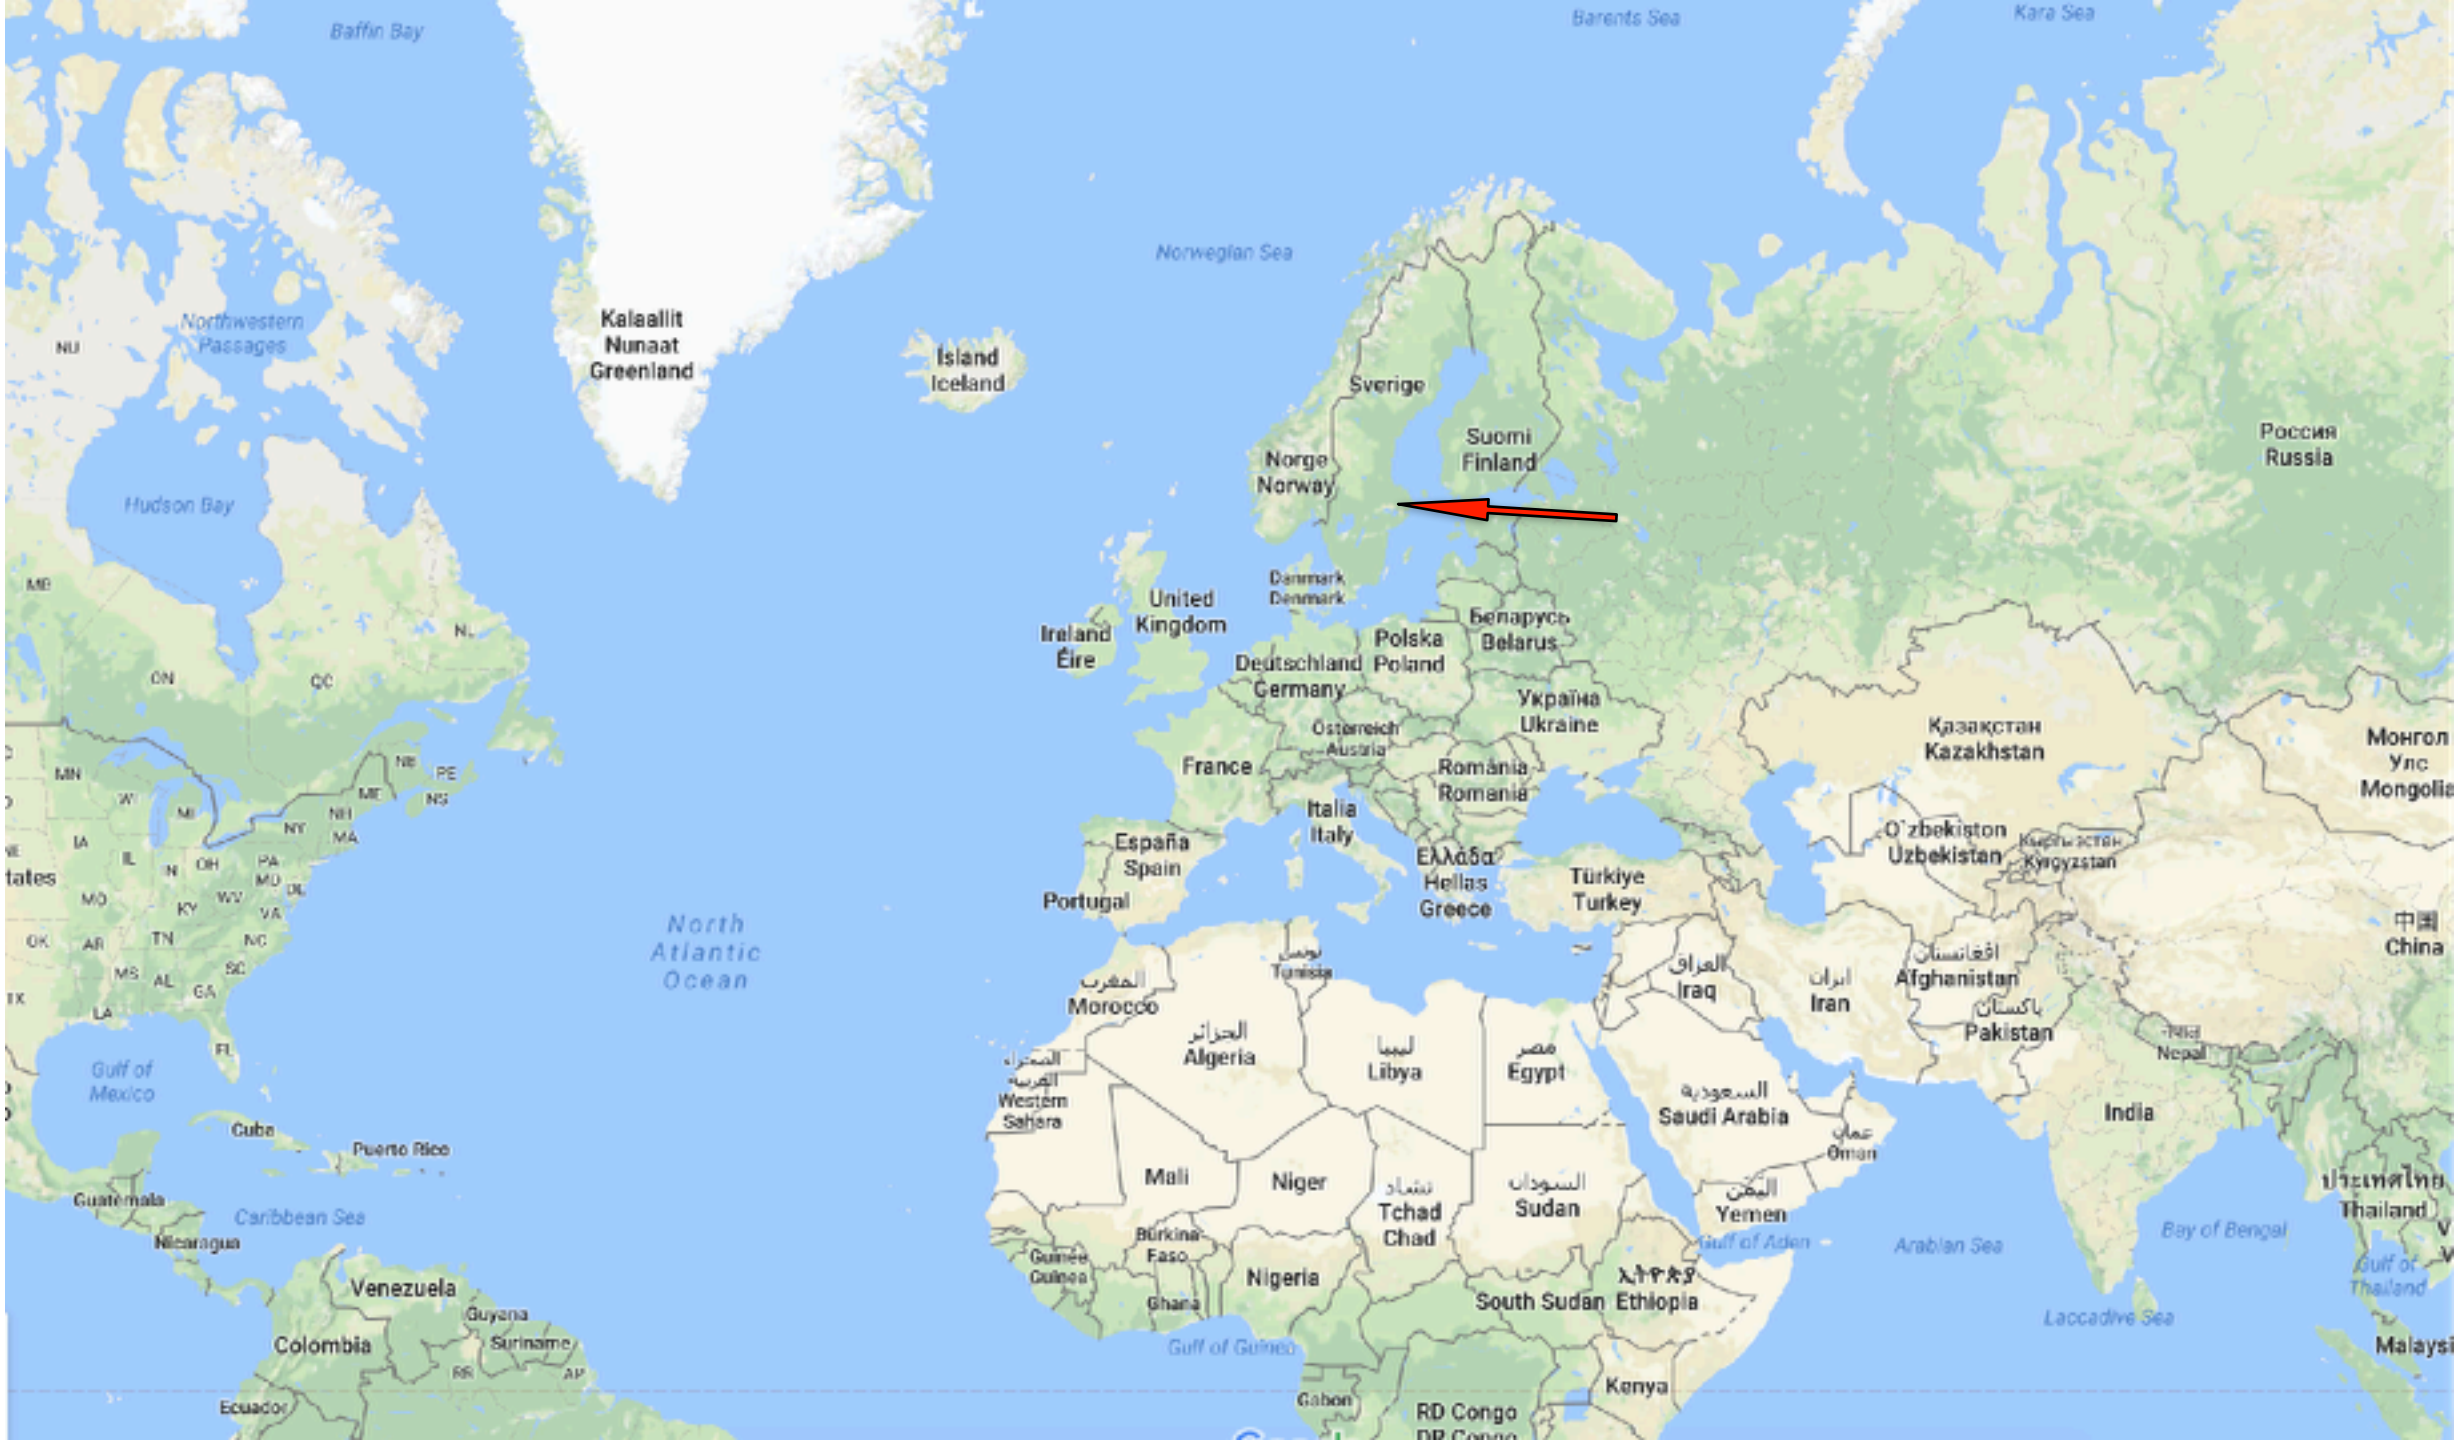
Baffin Bay
Barents Sea
Kara Sea
Norwegian Sea
Northwestern Passages
NU
Kalaallit Nunaat Greenland
Island Iceland
Sverige
Suomi Finland
Norge Norway
Россия Russia
Hudson Bay
Denmark Danmark
United Kingdom
Ireland Éire
Беларусь Belarus
Polska Poland
Deutschland Germany
Україна Ukraine
Österreich Austria
France
România Romania
Italia Italy
Қазақстан Kazakhstan
Монгол Улс Mongolia
España Spain
Portugal
Ελλάδα Hellas Greece
Türkiye Turkey
O'zbekiston Uzbekistan
Кыргызстан
North Atlantic Ocean
تونس Tunisia
المغرب Morocco
العراق Iraq
ایران Iran
افغانستان Afghanistan
پاکستان Pakistan
中国 China
الجزائر Algeria
ليبيا Libya
مصر Egypt
नेपाल Nepal
Gulf of Mexico
الصحراء الغربية Western Sahara
السعودية Saudi Arabia
India
Cuba
Puerto Rico
Mali
Niger
تشاد Tchad
السودان Sudan
عمان Oman
Guatemala
Caribbean Sea
Nicaragua
Burkina Faso
Guinée Guinea
Nigeria
Ghana
Gulf of Guinea
اليمن Yemen
Gulf of Aden
Arabian Sea
Bay of Bengal
ประเทศไทย Thailand
Gulf of Thailand
Venezuela
Guyana
Colombia
Suriname
ኢትዮጵያ Ethiopia
South Sudan
Laccadive Sea
Malaysi
Ecuador
Gabon
Kenya
RD Congo

# Sweden


Turkey
Turkmenistan
Syria
Lebanon
Iraq
Afghanis
Israel
Jordan
Iran
Egypt
Persian Gulf
Riyadh
United Arab Emirates
Saudi Arabia
Mecca
Red Sea
Oman
Sudan
Yemen
Eritrea
Gulf of Aden
Djibouti
Arabian Se
Ethiopia
th Sudan

# SUNET in a nutshell

- **The Swedish NREN**
  - Network - SunetC (AS1653)
  - Identity federation – SWAMID
  - Hundreds of OTT services (iaas, paas, vconf, backup, storage, hosting ++)
  - One of the oldest RIPE LIRs
  - One of the prime owner and customer of NORDUnet (AS2603)
  - Non-profit funded out of the Swedish Research Council
  - One of the real OG´s in IP-networking.
  - Every single packet in the network can be research, traffic wants to be free
    - Netflix, Google, Bittorrents, Wikipedia, Google, Twitter, LHC Particle Explosion data…etc etc…

# SUNET in a nutshell

- **Our Users – all personnel and students at:**
  - Universities and university colleges
  - Dormitories and ISPs serving dormitories
  - Art colleges, Royal library, governmental central museums.
  - Other organizations with research activities. (like the space-people)
  - Govermental entities with specific needs. (election authorities etc..)
  - MOU with the national archives.
  - All supercomputers
  - Non-profit that fights for the Internet or conduct research and or lab activities
  - And some other cool people.

# Backstory

- **All contracts ended at the same time.**
  - 15 year old fiber IRUs ending december 2016
  - 9 year old DWDM Tender and EOS ending february 2017
  - 5 year old IP/Packet Tender ending february 2017
  - All previous equipment written off, not a single $ locked up
  - Design old and needed a complete revamp
    - Dual-Star with Stockholm Origin
      - Router expensive, Optical cheap.
      - Very long eBGP links


# Nationwide network, 8.273 km

- 3 fiber routes to all regions.
  - Creates good redundancy and availability.
- Can connect customers to our equipment in ~30 cities.
  - Either ”in town” or at the university.
  - Customers outside these cities can be connected via (others) regional networks.
  - Possibility to convert another ~80 sites if needed
  - One contract, ~150 providers of fiber

# Fiberplant

- Core fibers mostly Aerial in high voltage power lines
  - New and old fibers, G.652B some G.652D
  - OTDR and attenuation data delivered as they measured
  - Kmz data
- Access fiber from city networks
  - Two separate paths to university
  - No KMZ data, because they **suck**


“Diverse” access fiberpaths

# Fiber Tender

- Only ***one*** patch the first 20km and
  - =<-42dB reflection
- .25db/km
- Measurement of water peak
- OTDR Measurement
  - Short Pulse
  - Long Pulse
- Require all connectors to pass IEC 61300-3-35
  - Splice >= 0.2dB
  - Connector >= 0.4dB

  (Proof of this with pictures of all connectors)
- Equal Spacing between sites
- Connector-Splice plan / fiber Stretch
- Attenuation/Stretch shall include all connectors and splices

# OTDR RAMAN SECTION

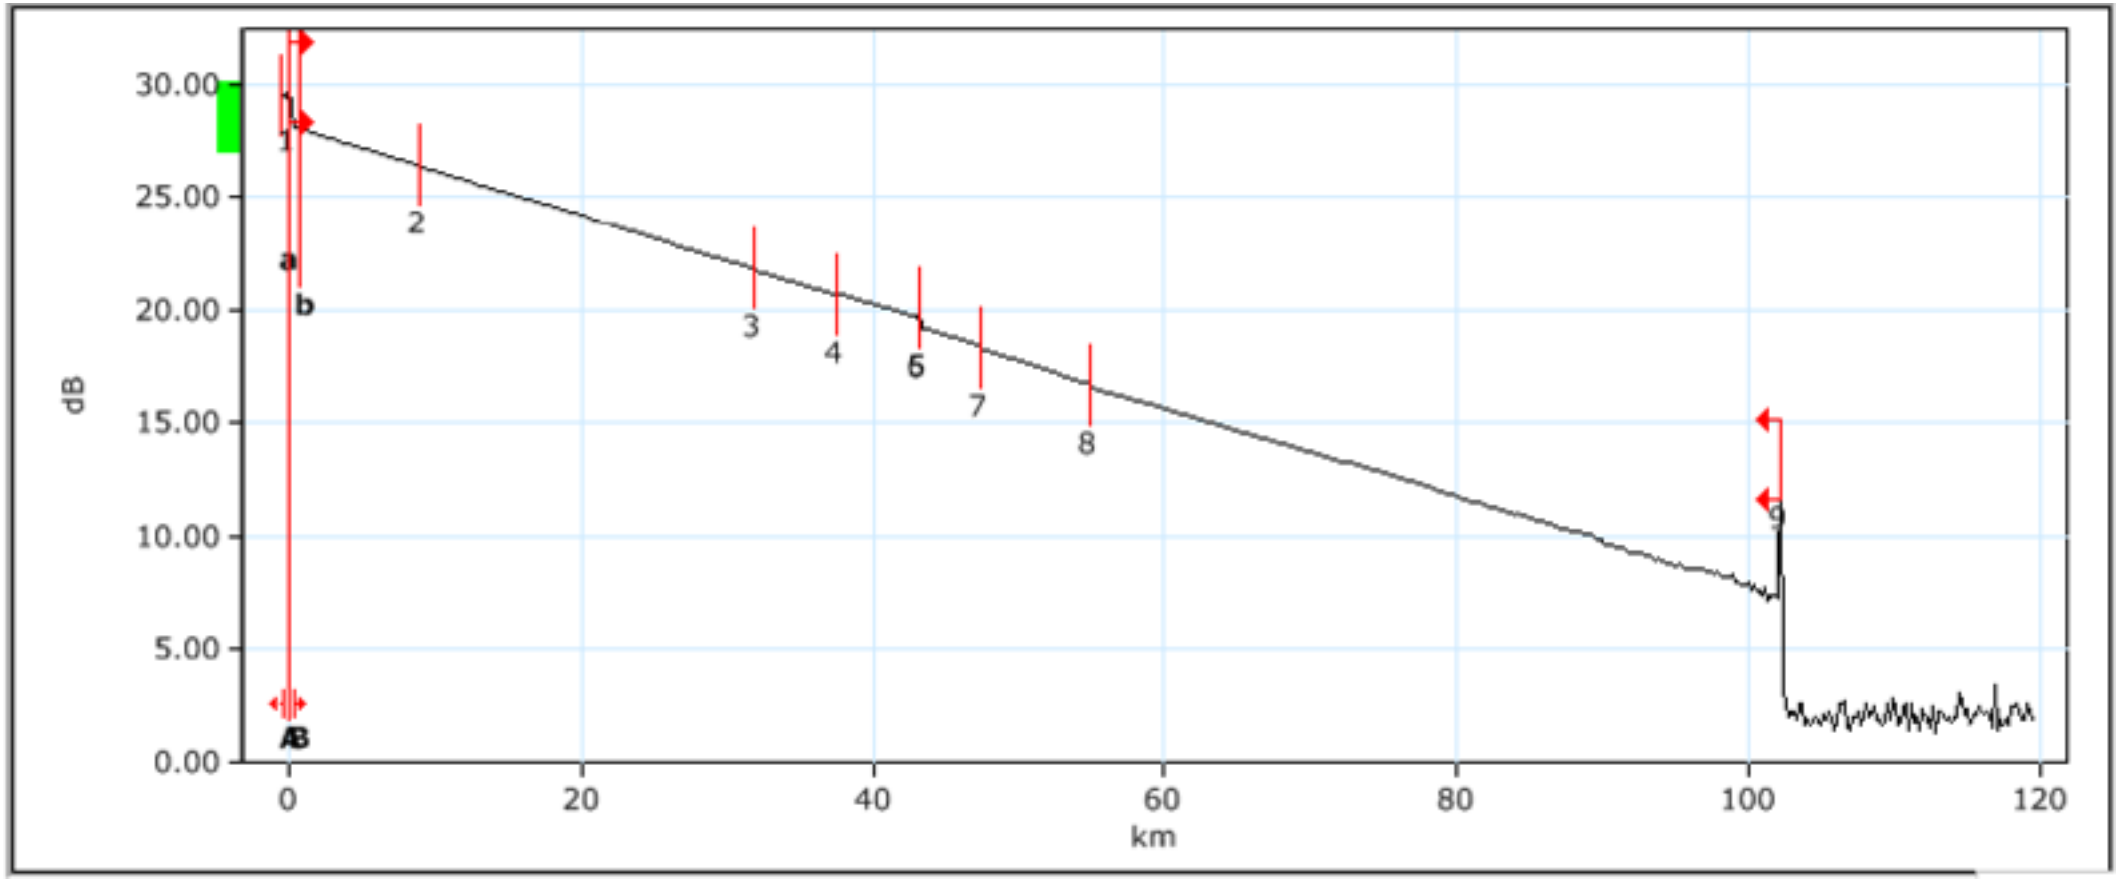

# OTDR RAMAN SECTION
# SHORT DISTANCE

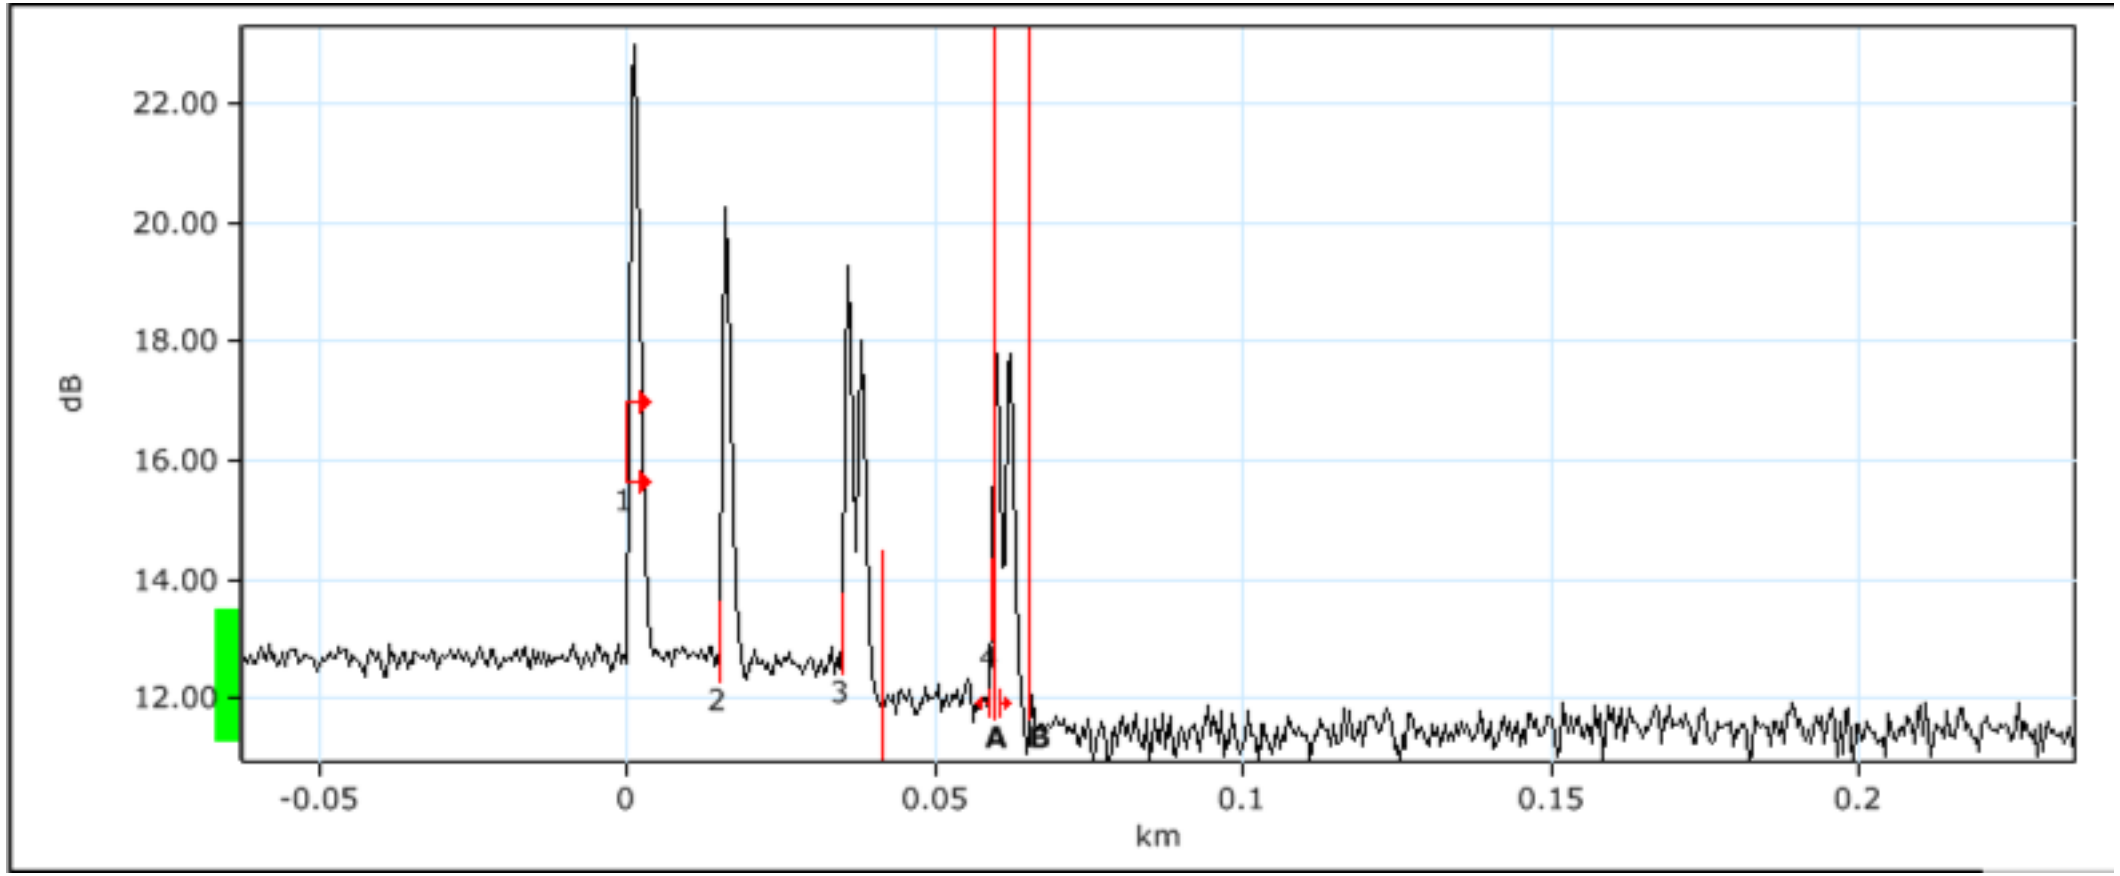

# A (new-ish) way of building networks

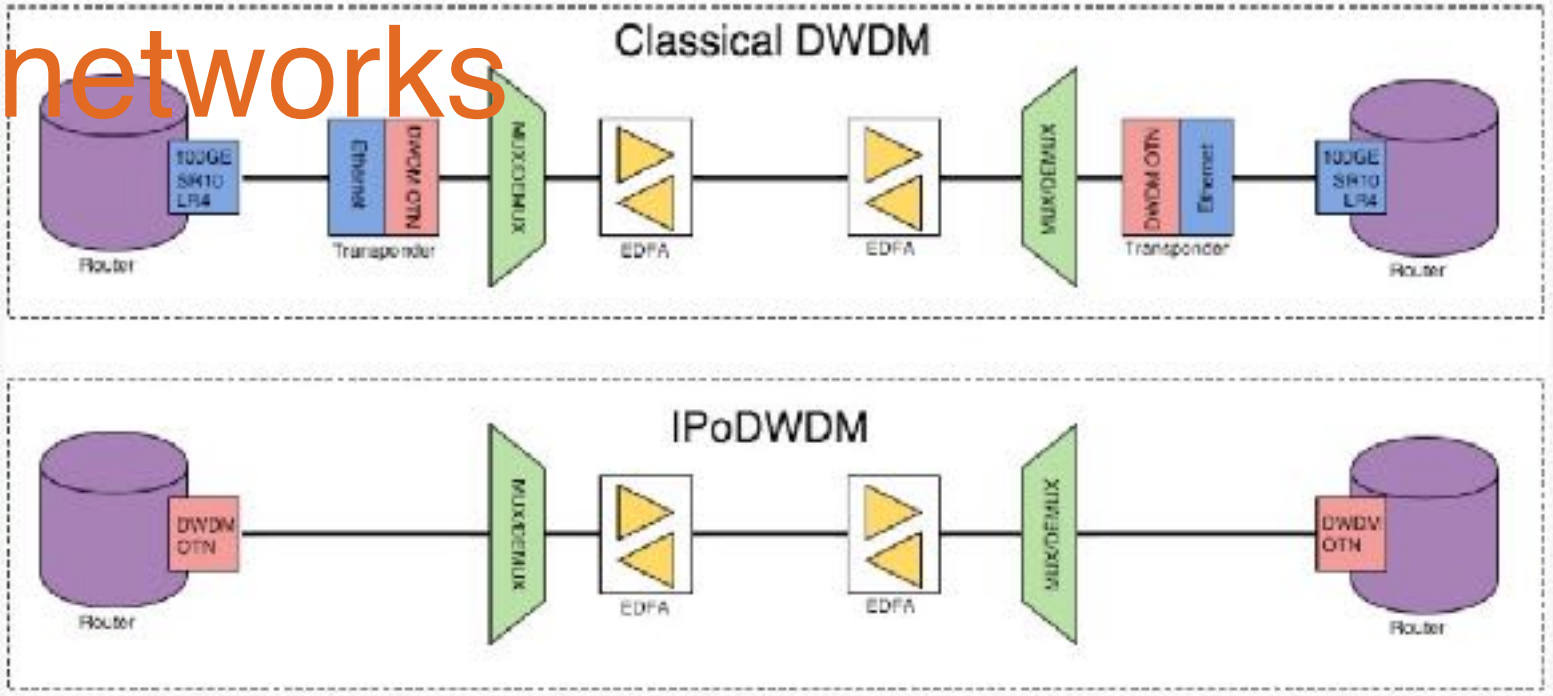


- IP over DWDM (IPoDWDM).
  - Eliminates unnecessary and expensive components => less components that fails
  - Reroute traffic before fail
  - Much lower CAPEX but also lower OPEX
- **Gridless**

# Equipment

- DWDM
  - ADVA FSP3000R7
    - Amps and ROADM
      - Gridless
      - Colorless
- Routers
  - Juniper MX
    - 100G DWDM Coherent
    - 100GE – 1GE

# Routers

- 15 MX80
- 53 CPE, MX480
- 29 Core-PoP, MX960
- 5 Core-PoP, MX2010
- 2 Core-PoP, MX2020

# Optical

- 35 ROADM FSP 3000
- 80 ILA FSP 3000

CPE

POP

Malmö
Göteborg Luleå

Stockholm


MWA FWA FWA OOH MAH FUH.
I CAN'T UNDERSTAND A WORD YOU SAY.
www.dilbert.com scottadams@aol.com
AND YOU'RE POORLY DRESSED. YOU MUST BE SOME SORT OF TECHNOLOGY EXPERT.
OR A RODEO CLOWN.
3/8/99 © 1999 United Feature Syndicate, Inc.
THERE'S A FIFTY PERCENT CHANCE I HIRED A DATA NETWORK ENGINEER.
I SMELL HAY.

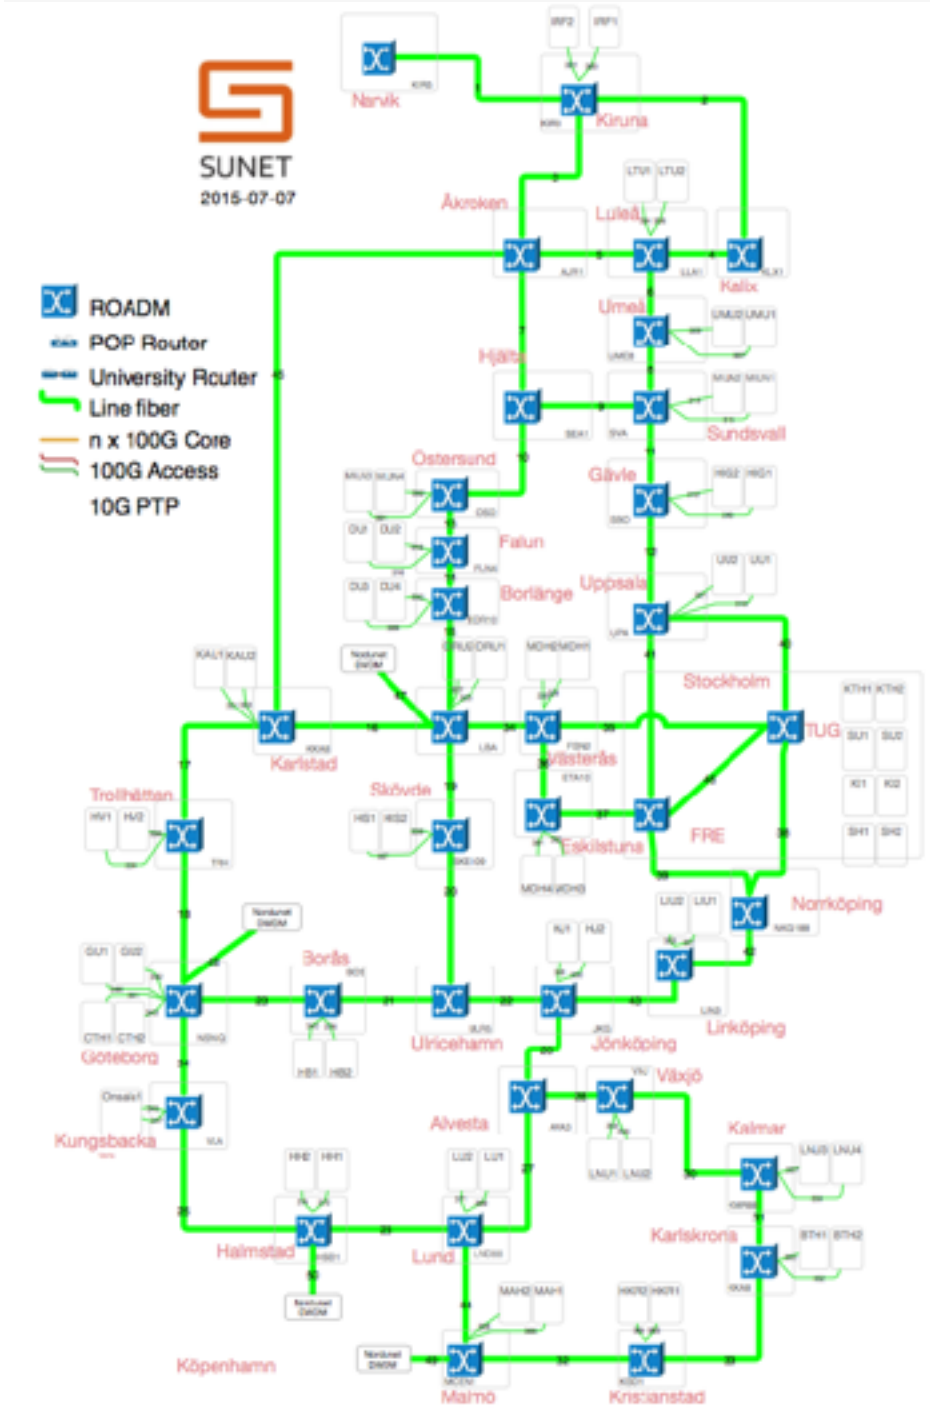


# Optical Design

- ROADM site in each University city
- 31 UNI-CITY 9degree ROADM
- 4 pure ROADM sites
  - Where no traffic is dropped

# Hybrid Raman EDFA Amplifiers from ADVA

- Raman at all places with more than 17db span loss
  - Problem with Raman on some stretches => rebuild by adding an intermediary site
- Optimize for high OSNR
- Low Noise EDFA
- Preparation for the future, 400G/1T
- Gain 9-15dB, depends on fiber and connectors

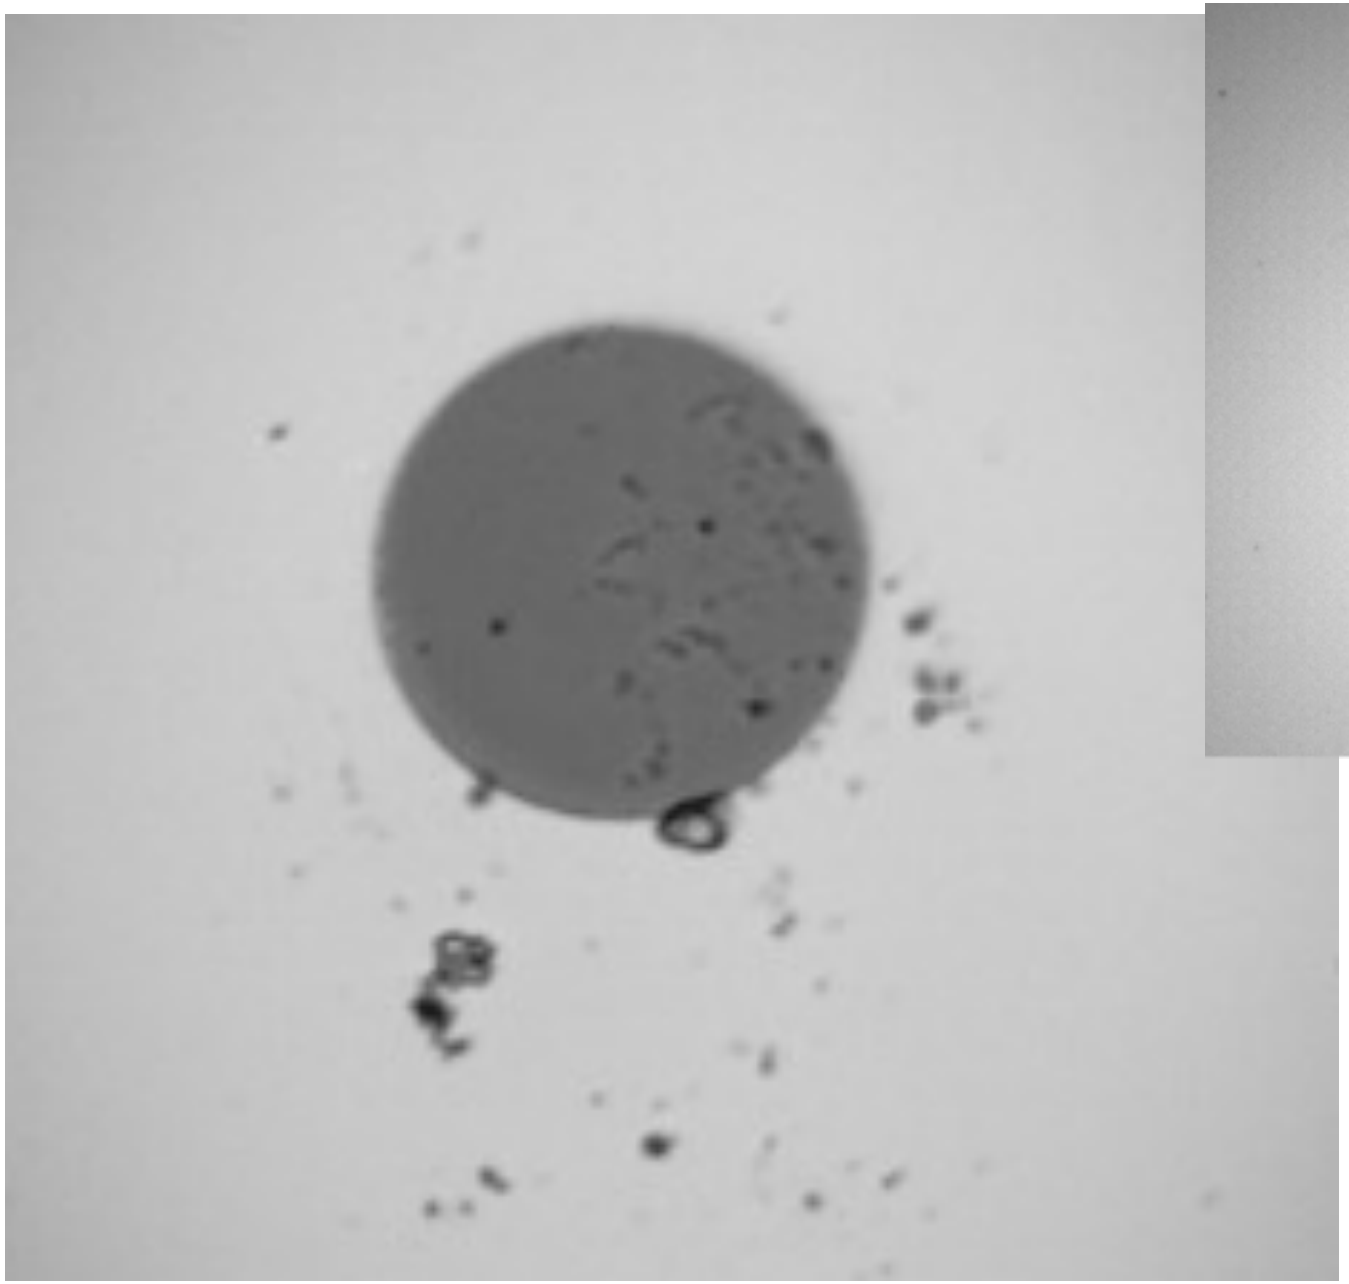

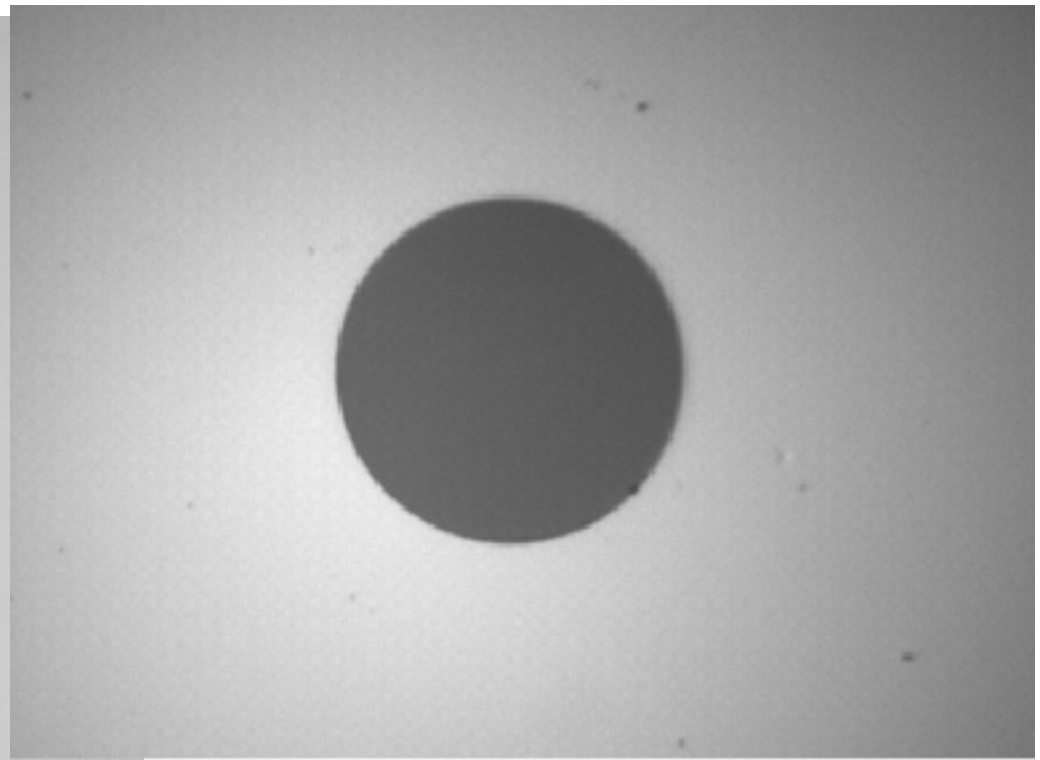

# Router Design

- Connected via ER4/LR4
- Coherent interface to the next city/ pop
- In (supercore sites) Stockholm, Göteborg and Luleå
  - Two routers at the PoPs

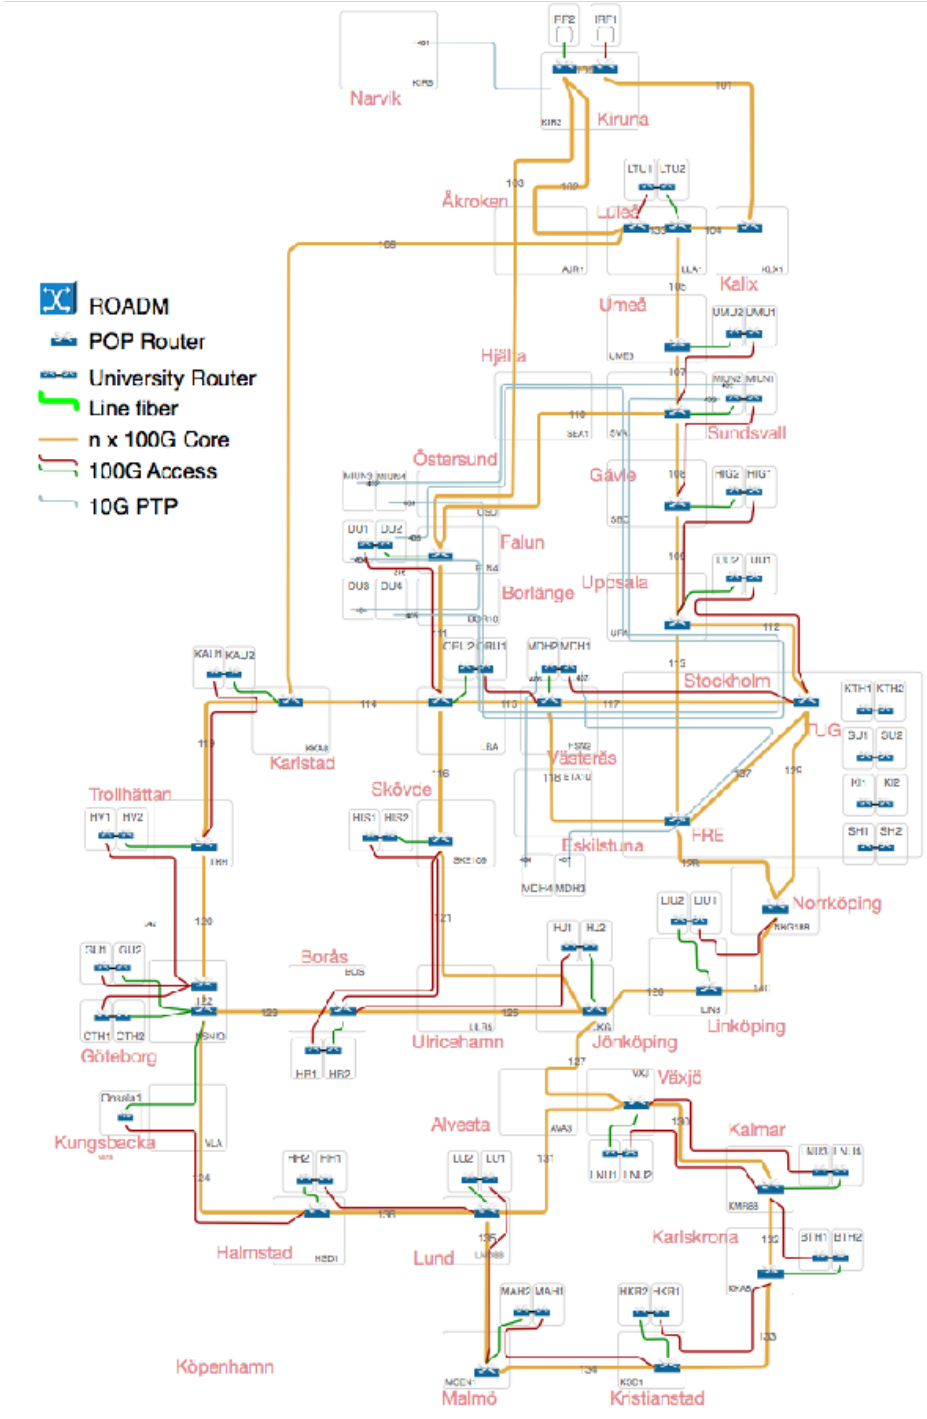


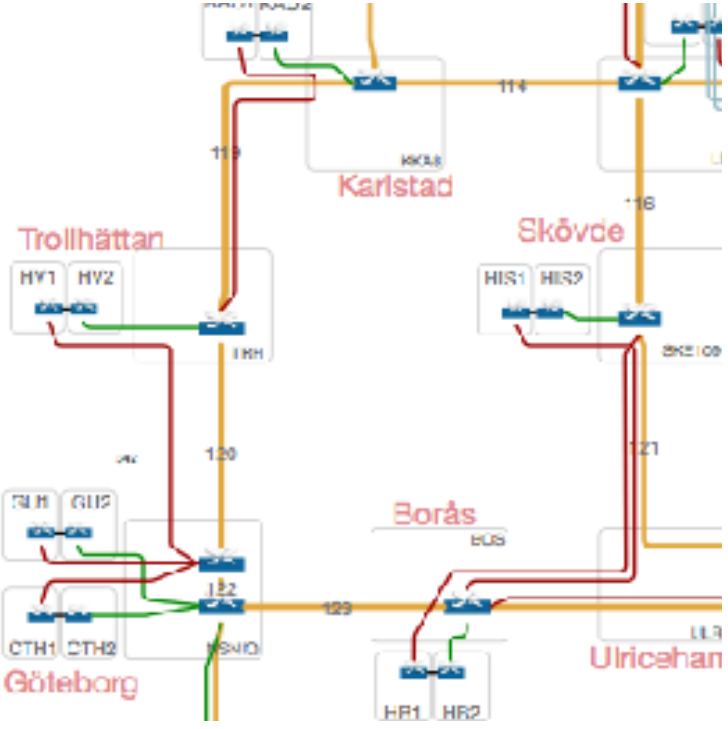

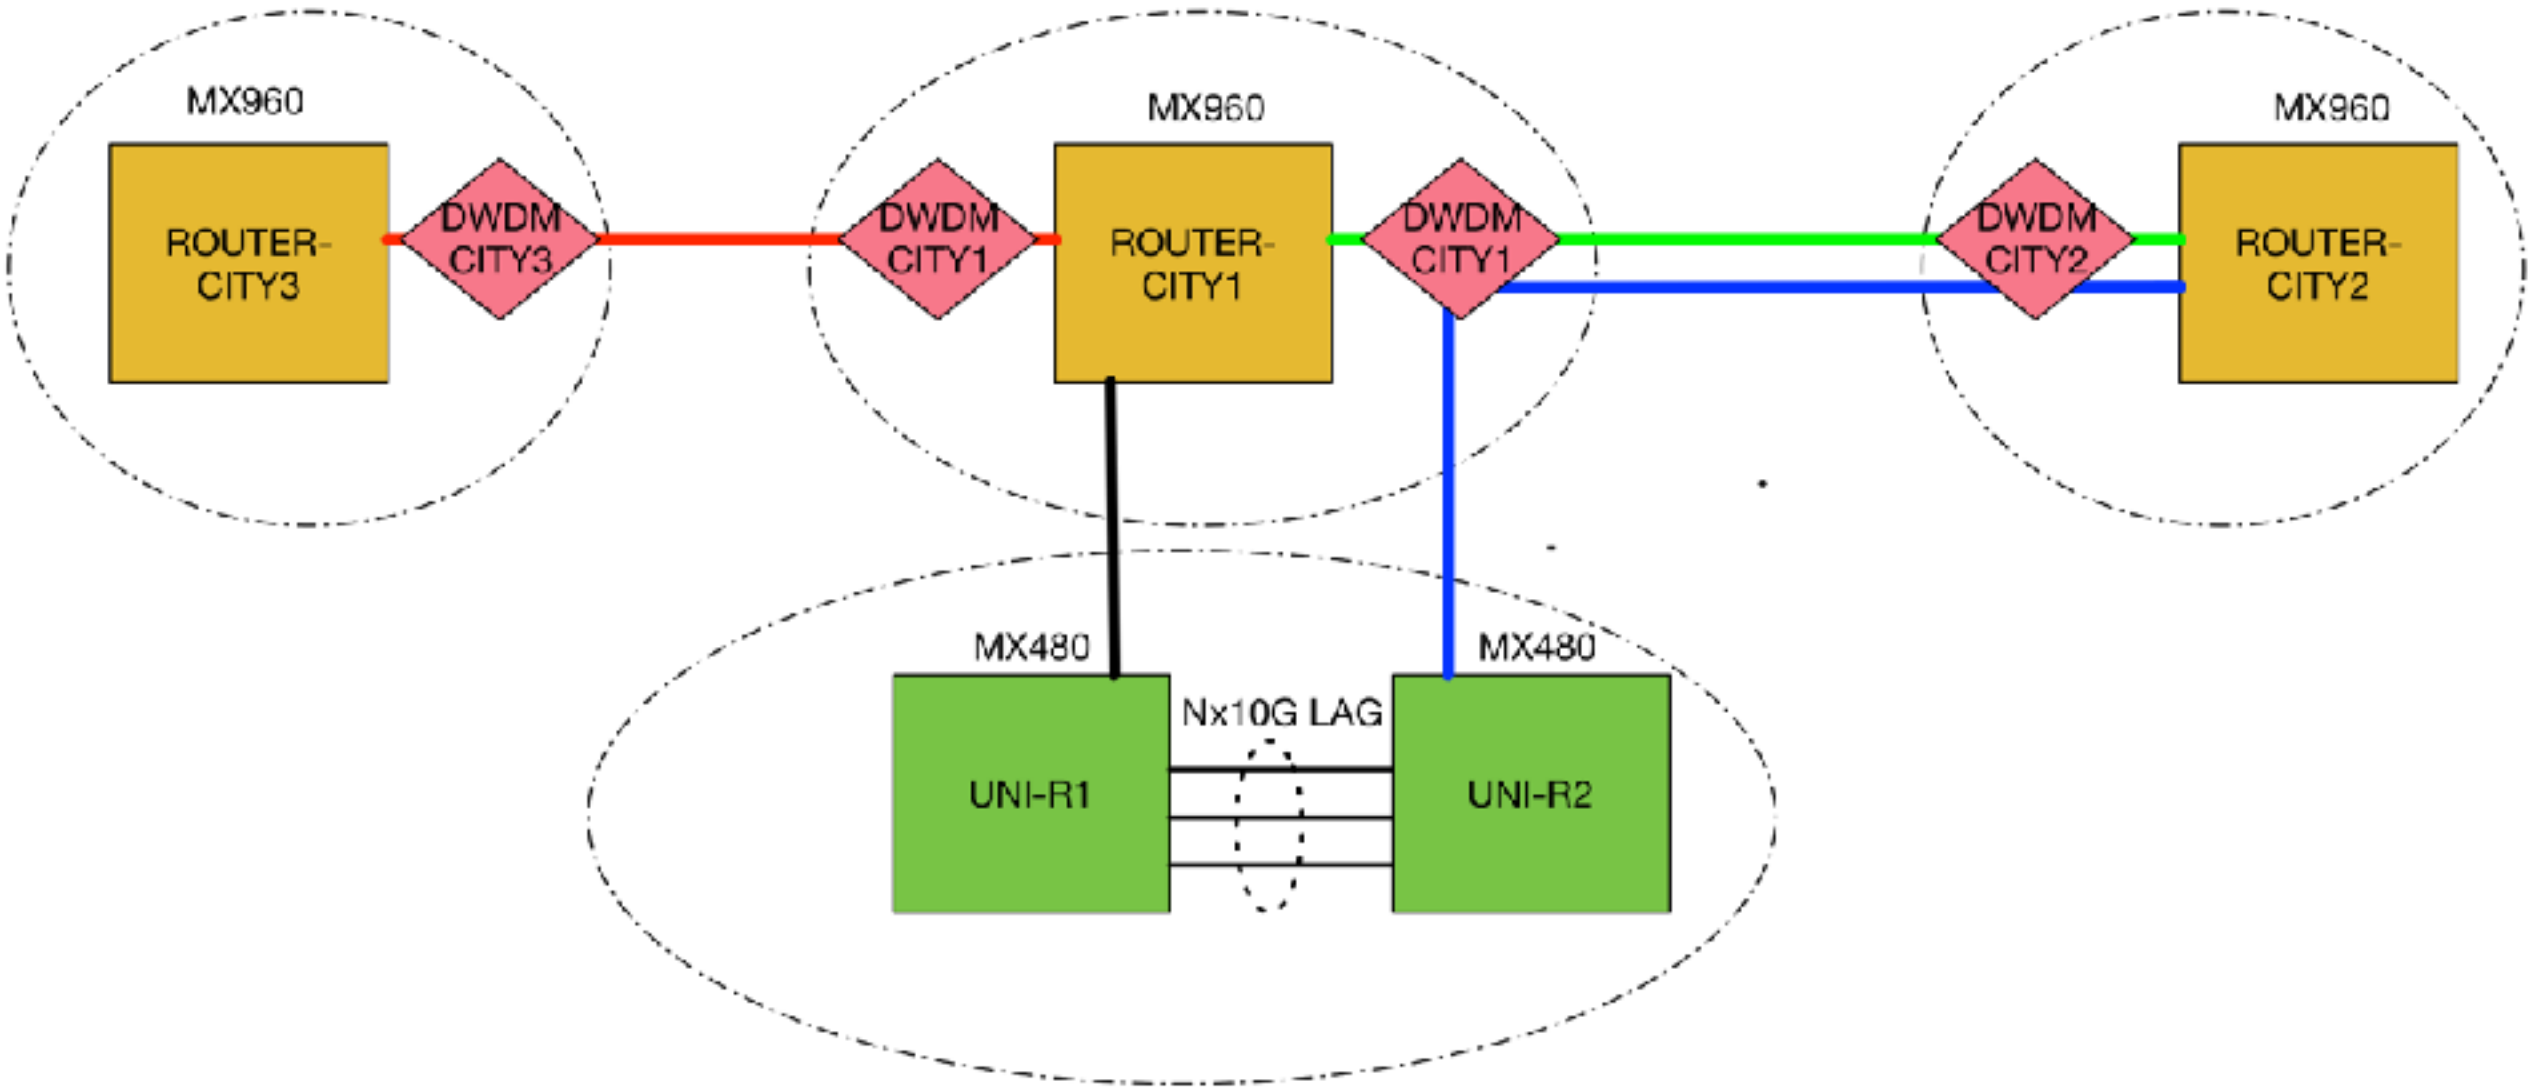
MX960
ROUTER-
CITY3
DWDM
CITY3
MX960
DWDM
CITY1
ROUTER-
CITY1
DWDM
CITY1
MX960
DWDM
CITY2
ROUTER-
CITY2
MX480
UNI-R1
Nx10G LAG
MX480
UNI-R2

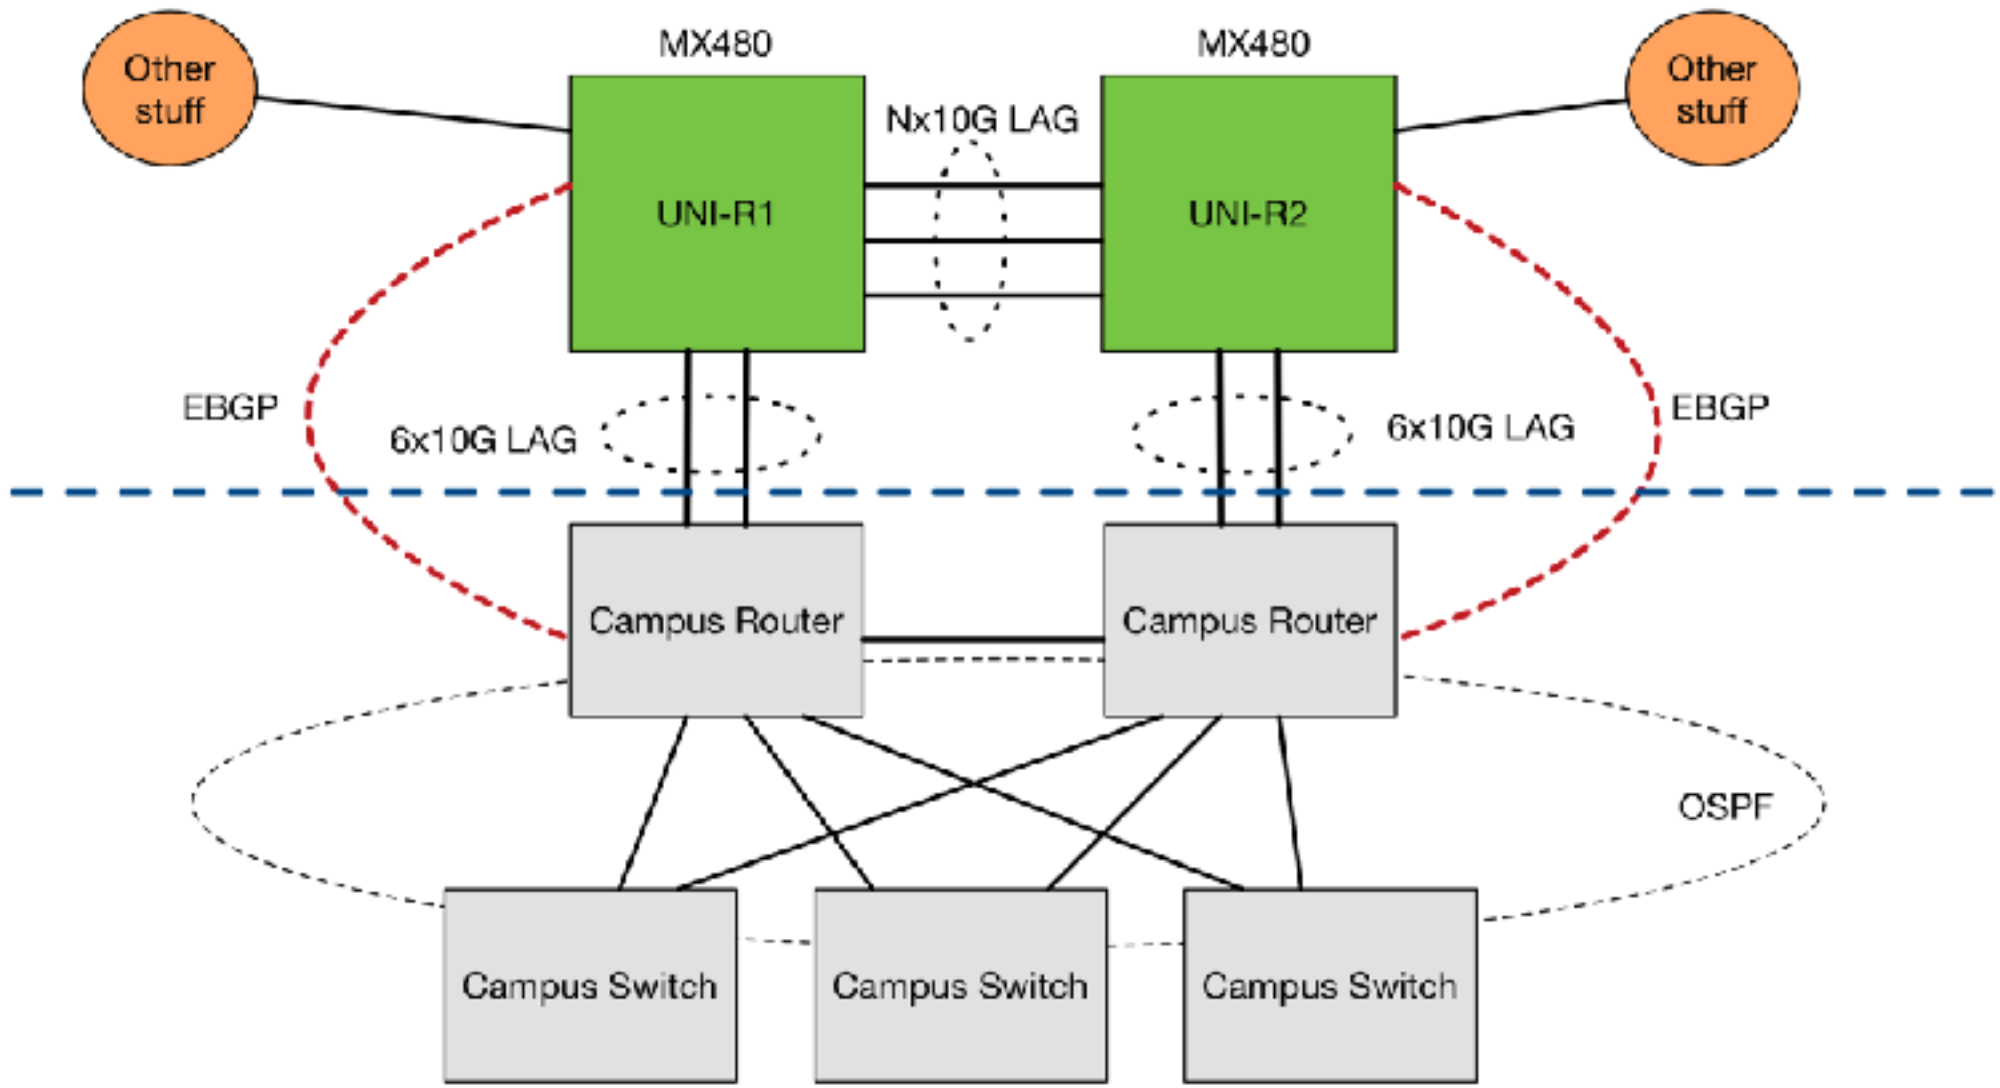
Other stuff
MX480
MX480
Other stuff
Nx10G LAG
UNI-R1
UNI-R2
EBGP
6x10G LAG
6x10G LAG
EBGP
Campus Router
Campus Router
OSPF
Campus Switch
Campus Switch
Campus Switch

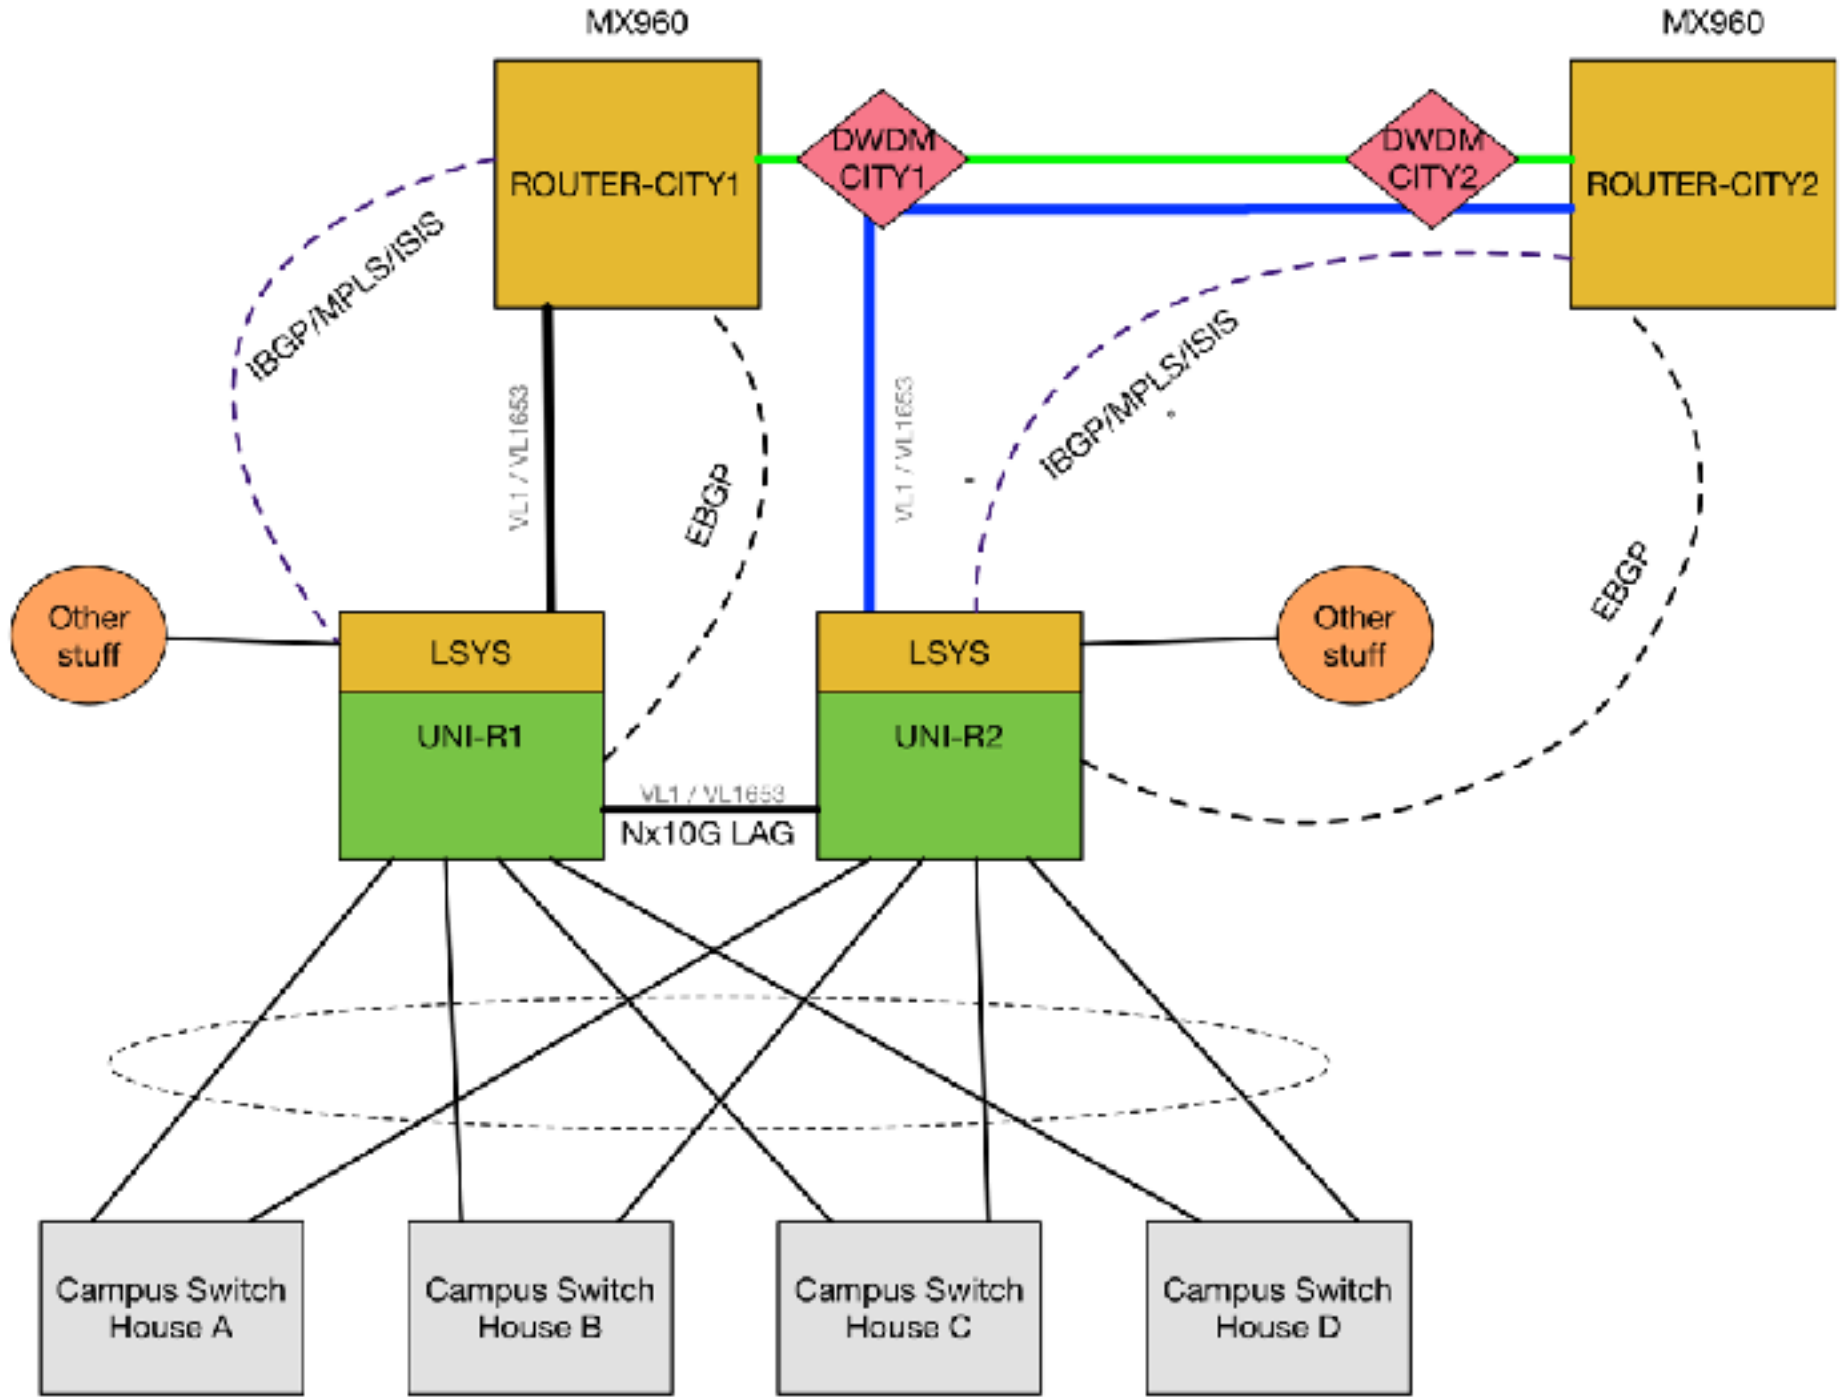
MX960
MX960
ROUTER-CITY1
ROUTER-CITY2
DWDM CITY1
DWDM CITY2
IBGP/MPLS/ISIS
IBGP/MPLS/ISIS
VL1 / VL1653
VL1 / VL1653
EBGP
EBGP
Other stuff
Other stuff
LSYS
LSYS
UNI-R1
UNI-R2
VL1 / VL1653
Nx10G LAG
Campus Switch House A
Campus Switch House B
Campus Switch House C
Campus Switch House D

# Site

# Site

# Backup Power

# Site

- Two racks per site
- Dual -48V DC battery bank feed
- Generator backup
- Free cooling
- 3G/4G Out of band

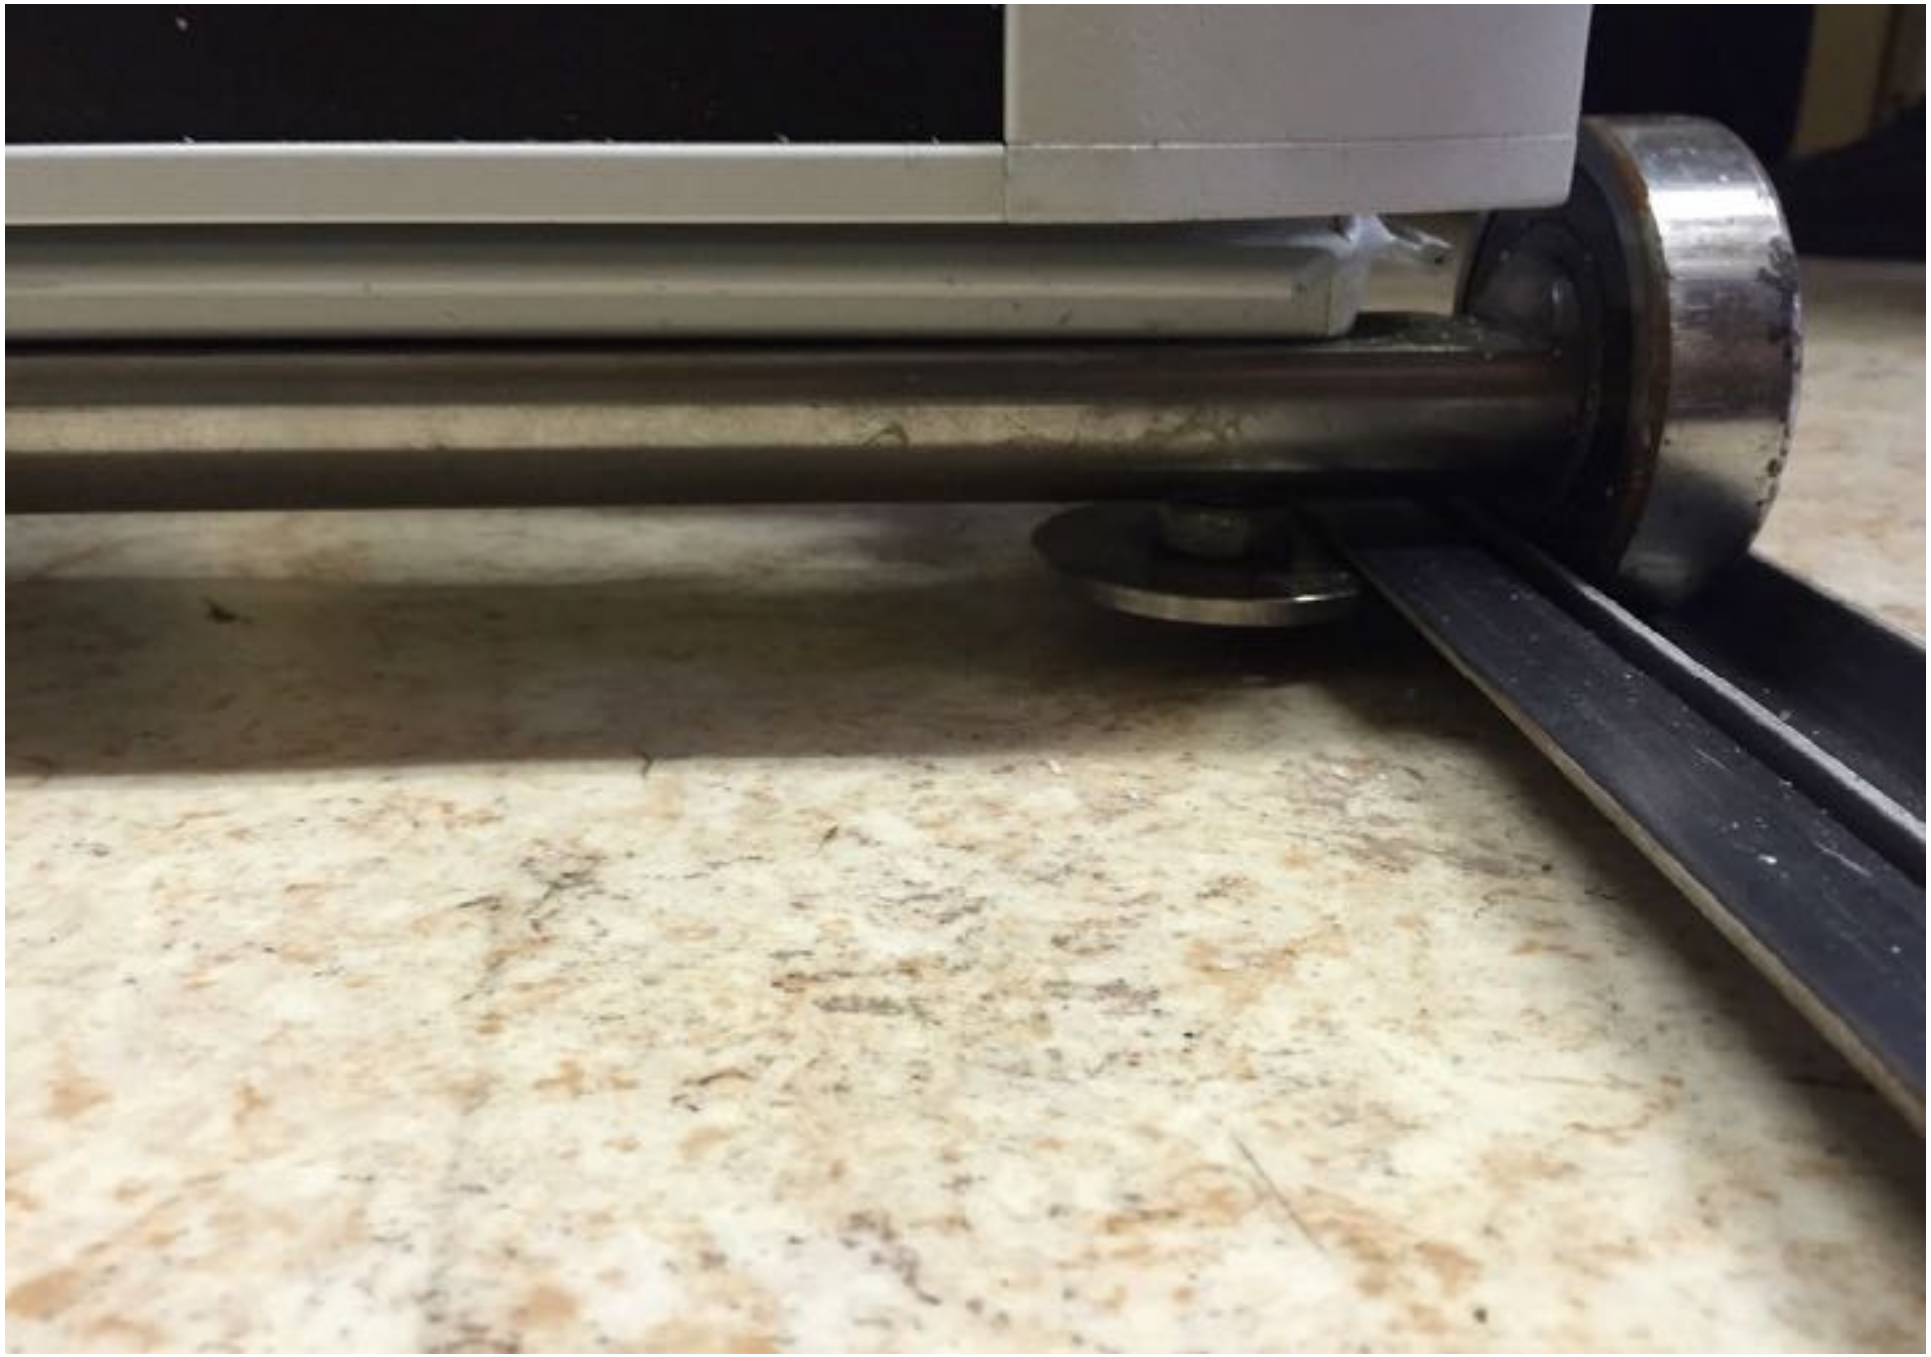

# Racks

- Custom made
- Slide bearings
- Tip safe

# Rack

- Cable chains

# Load test

Juniper mx960 fully loaded
158.76 kg

# Console server + SW

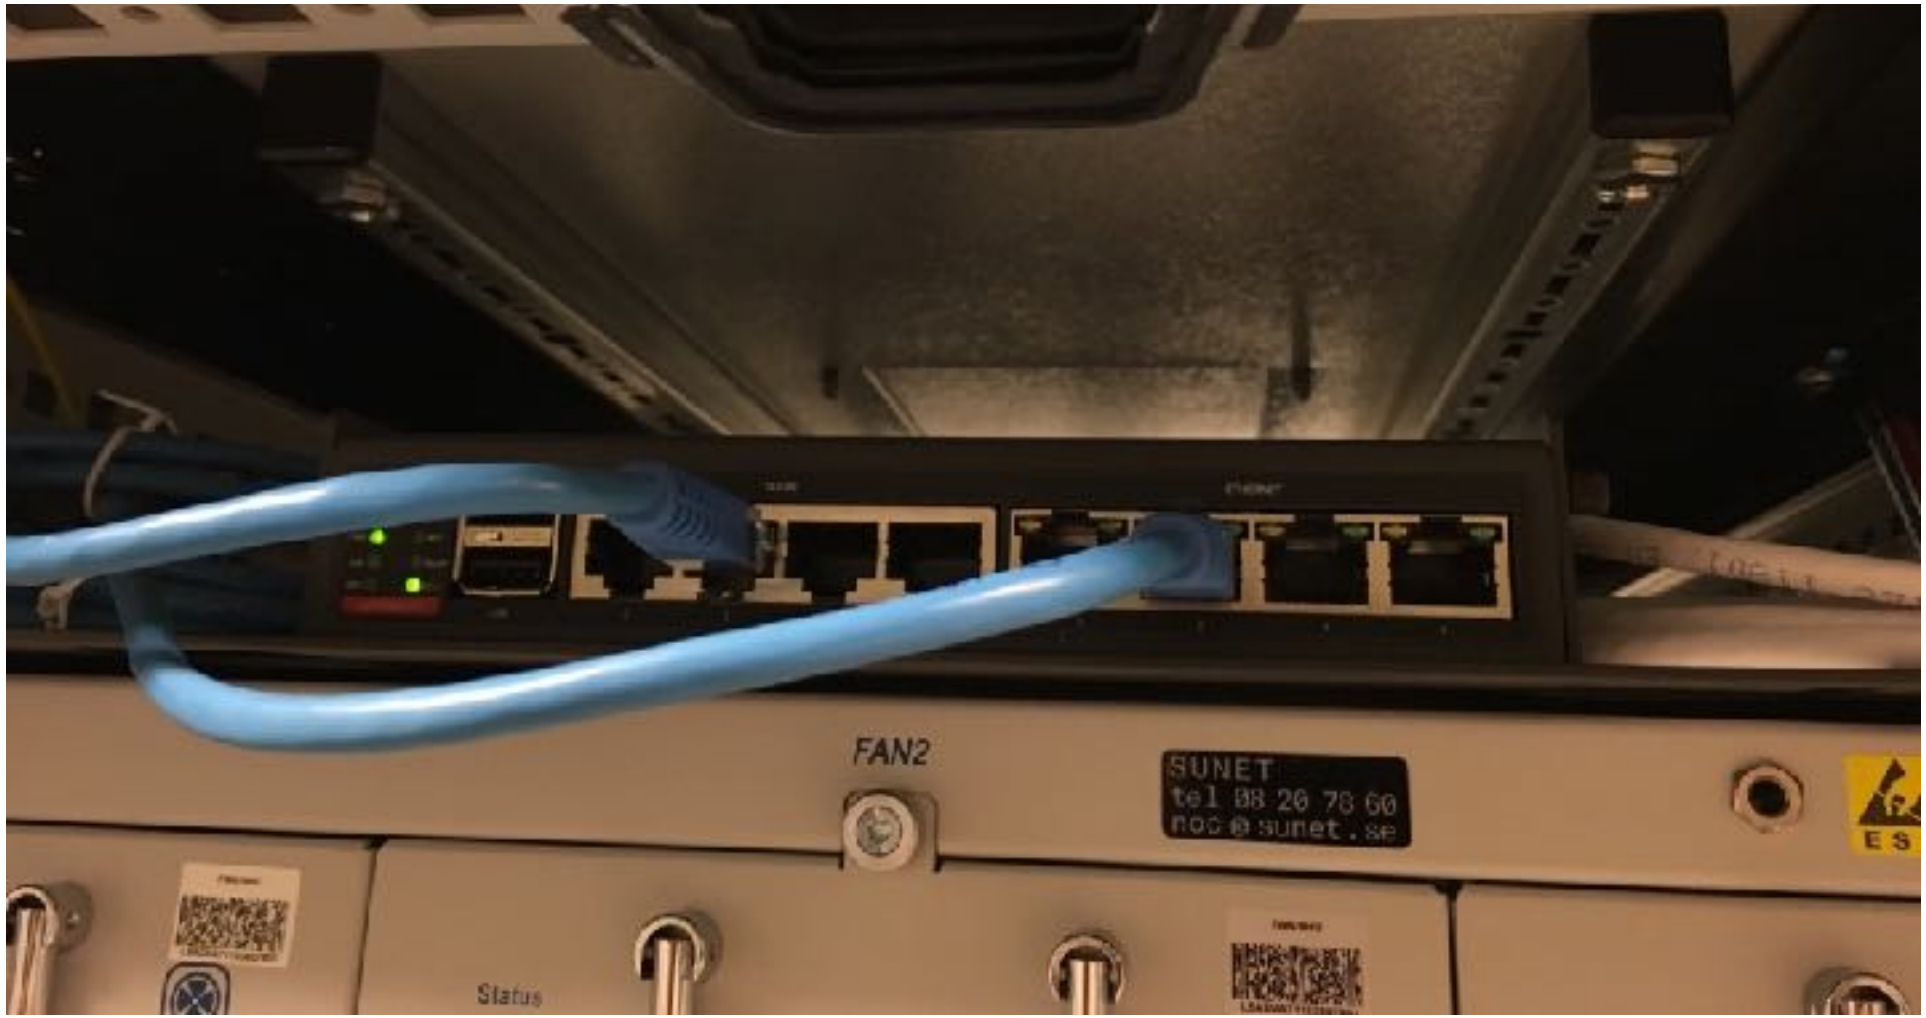


- **Opengear ACM5500**
- 4G / 3G Cellular
- Own APN, L3VPN
- Seriel Console
- Ethernet router
- 4GB Junos over X25 :(
- 4GB Junos over 45Mbit 4G
- Connections to DWDM

# Management

- Waves provisioned using GMPLS inside the ADVA network.
- ADVA NMS supports control of Juniper MX interfaces, work in progress to support PM extraction.
- NETCONF YANG support “this year” to be able to control ROADM from Cisco NSO, like we do with everything else.
- 90% of the configuration in the routers is not owned by a human but rather central orchestration.

# War stories

- Fiber/sites ☺️
  - Succesful tests of 4000km 100G optics, successful tests of 2300km 200G
  - Maybe the best in-production optical network in the world (but still pretty crappy)
  - 0 Soft or hardware related bugs/outages in the ROADMs
- Fiber/sites ☹️
  - 3x destroyed RAMAN-amps due to fat/residue on fiber = welded connector
  - 94 Fault-tickets opened on 132 backbone-spans. 80% fixed with cleaning/new patches
  - Near death experiences…
  - Optical vendors not at the forefront when developing new functions..


G5TA2 VT W401
G5TA2 VT W402
G5TA2 VT W403
G5TA2 VT W404

# War stories

- Routing/Packets ☺️
  - MPC3-DWDM-MIC = Zero issues
  - Optical performance typically above spec.
  - 0 Issues due to Logical Systems
  - 4G OOB 99.4% uptime

- Routing/Packets ☹️
  - MPC3-NG: during commissioning 10% DOA ☹️
  - Lots of software bugs related to 15.1-fX
  - 26x Prio1 PRs found, some fixed, some workaround, some waiting.
  - Disabled features to reduce PR-vulns (BGP-PIC, Flowspec)
  - 174 JTAC cases excluding RMAs


SUNET

# QUESTIONS?

hugge@nordu.net
(or find me on any IRC network)

https://www.sunet.se/bloggdel/teknikbloggen/