SUNET

FULL SPEED AHEAD!

Last few posts has been about DDOS, so lets kick it back for a while and look how the Network rollout is progressing.

Fiber:
We have recieved almost all 8000km of fibre the backbone will be built upon, two or so spans is delayed but all critical paths is delivered. We have measured and quality-assured most of the spans with various success. We do have a few challenges with the current fiberplant that we are trying to solve or workaround right now. The first major problem we have is citynetworks and RAMAN-amplifiation. If we take Stockholm for example, both of our datacenters at both Tulegatan and the bunker in Fredhäll is quite centrally placed locations in the inner-city. This means that there is quite a big stretch of citynetwork (STOKAB) we need to run through before we go into the powerlines where our long-distance fibre is placed. One characteristic of Citynetworks is that there is a lot of connectors, connectors means attenuation and reflectance. The closer to source of light we come, the more of a problem this cause for us.

Länkmätningar

Mätområdesförlust	: 22.222 dB	Medelskarvförlust	: 0.458 dB
Spännvidd	: 80.8114 km	Max. skarvförlust	: 1.138 dB
Medelförlust	: 0.275 dB/km	Mätområde ORL	: 35.07 dB

OTDR-kurva

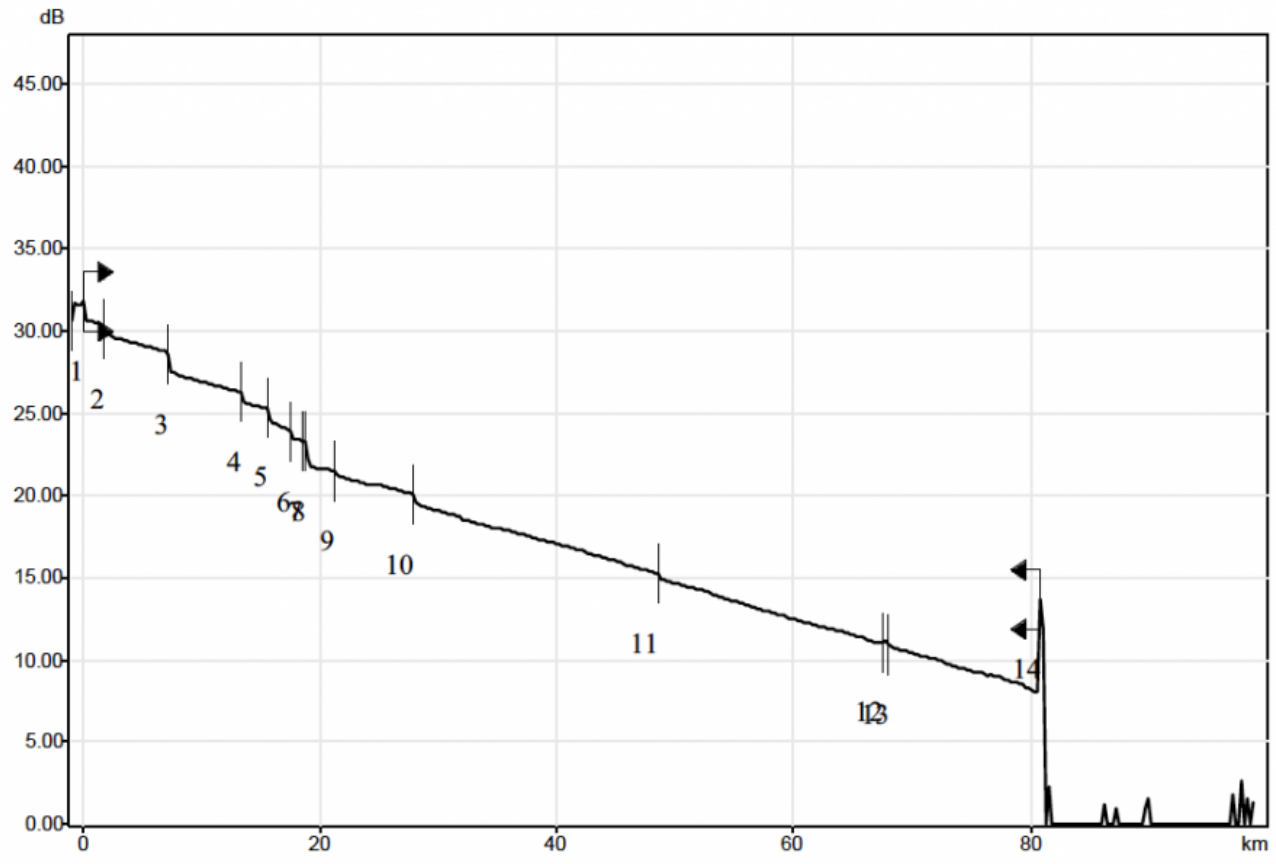

Here is an example of a OTDR coming from Tulegatan to ”Hamra” which is an amplifier-site north of Enköping that belongs to the optical span heading to Västerås. This fiber has 9 events before coming up to the 20km mark, there is very high connector-

attenuation with a average connector-loss of 0.458dB and a whopping 1.138dB as maximum loss. Seeing as the plan is to put a RAMAN amplifier in Tulegatan it will not even start at all with all these connectors so close to the amplifier itself. We tried solving this in a number of ways. The first idea was to have STOKAB splice a few of the connectors together (essentially splice until the amp starts), this was out of question since they want the network to be modular and a ”expresspath” out of Stockholm was not to think about. Next idea was to get another fiber to get out of the city from another vendor, but that quality on that one was even worse so we skipped that idea as well. The last idea is to put the RAMAN-amplifier directly after the worst fiber is passed through, this is where we are now and we are currently in the process of acquiring a small amplifier-site just after the 20km mark when worst part has passed, and from Stockholm we just use a regular EDFA-amplifier and not a hybrid-RAMAN to reach the new break-out site at the 20km mark which will then handle the path all out to Hamra.

This is not the only the problem we have with the fibreplant either in terms of putting RAMANs inside. We do have a few sites that is placed within a modular telehousing-facility. While these are extremely well-built, modular and luxurious sites we had a few setbacks here as well. Lets look on what the idea is...

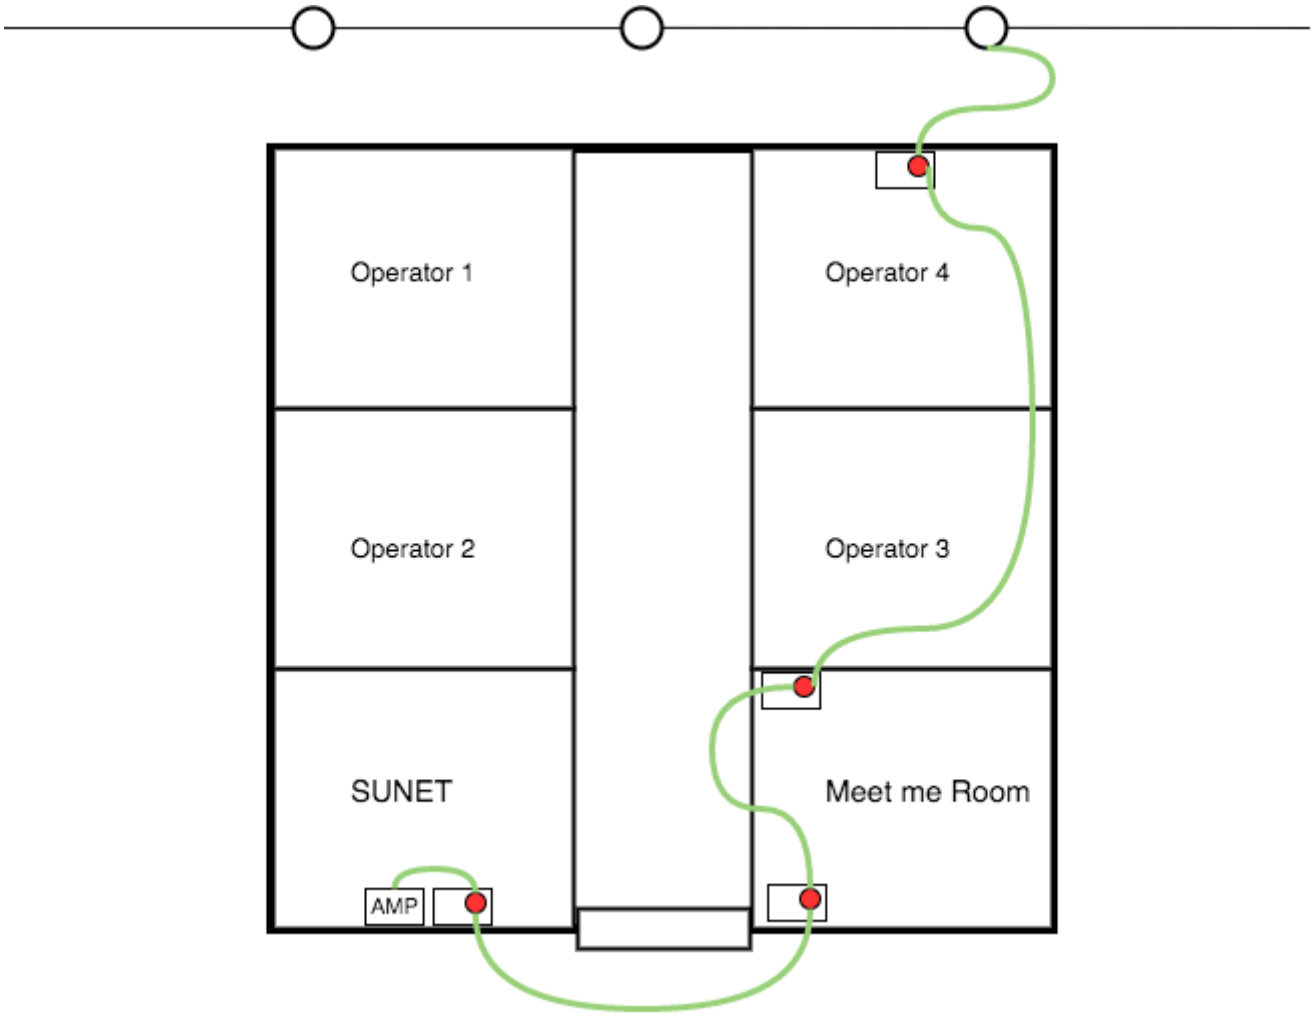

This is the general idea on how the modular sites works. This is a very smart and modular concept with small private-rooms for the operators to put stuff in, and a meet-me-room to be able to connect (and pay) to others operators without needing to enter other operators rooms. However the sharp eye can see that this generates us 4 connectors the first few meters and this is not good for us if it’s a RAMAN-qualified stretch. We have this problem south of Botkyrka on the span heading to Eskilstuna (almost 100km of fiber) where our RAMAN-amplifier is not even able to start due to the reflection produced in the ODFs, it should be noted that this is not APC-ODFs but regular SC-UPC which makes the problem even worse. Lets look at the facts...

Results

Span length:	0.7416 km	Average loss:	1.970 dB/km	Injection level:	12.4 dB
Span loss:	1.461 dB	Average splice loss:	---		
Span ORL:	40.37 dB	Maximum splice loss:	---		

Graph

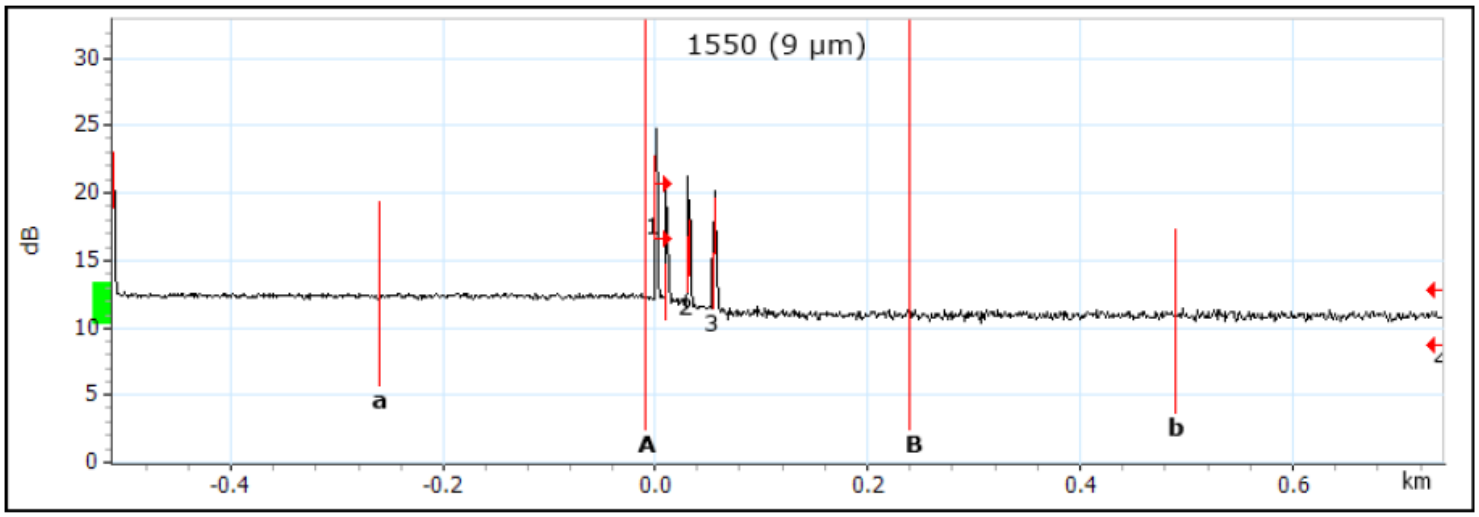

Markers

Marker	Position (km)	Value (dB)
a	-0.2600	12.521
A	-0.0100	12.477
B	0.2400	10.987
b	0.4900	10.499

A-B LSA attenuation:	12.344 dB/km	A-B average loss:	5.956 dB/km
A-B LSA loss:	0.125 dB	4-point Event Loss:	1.363 dB
A-B ORL:	41.53 dB	Maximum reflectance:	-43.1 dB

Might be abit hard to see, but it's actually four connectors even though the analyzer mark it as three here. First connector is the SUNET ODF, Second is SUNET ODF in MMR, Third is other operator in MMR and fourth is the ODF that connects to the grid-fiber. This is quite a lot of reflectance aswell since its. On this actual stretch its so much reflection that the RAMAN cant even start at all, it just shuts off and only works in EDFA-mode.

So the idea we had, was that we want to be able to also get a express-path here and bypass the MMR completely and go straight into the ODF of the line-fibre. So together with our subcontractors we made a proof-of-concept on-site with a loose patch on the floor.

Results

Span length:	0.7415 km	Average loss:	0.711 dB/km	Injection level:	12.4 dB
Span loss:	0.528 dB	Average splice loss:	---		
Span ORL:	42.61 dB	Maximum splice loss:	---		

Graph

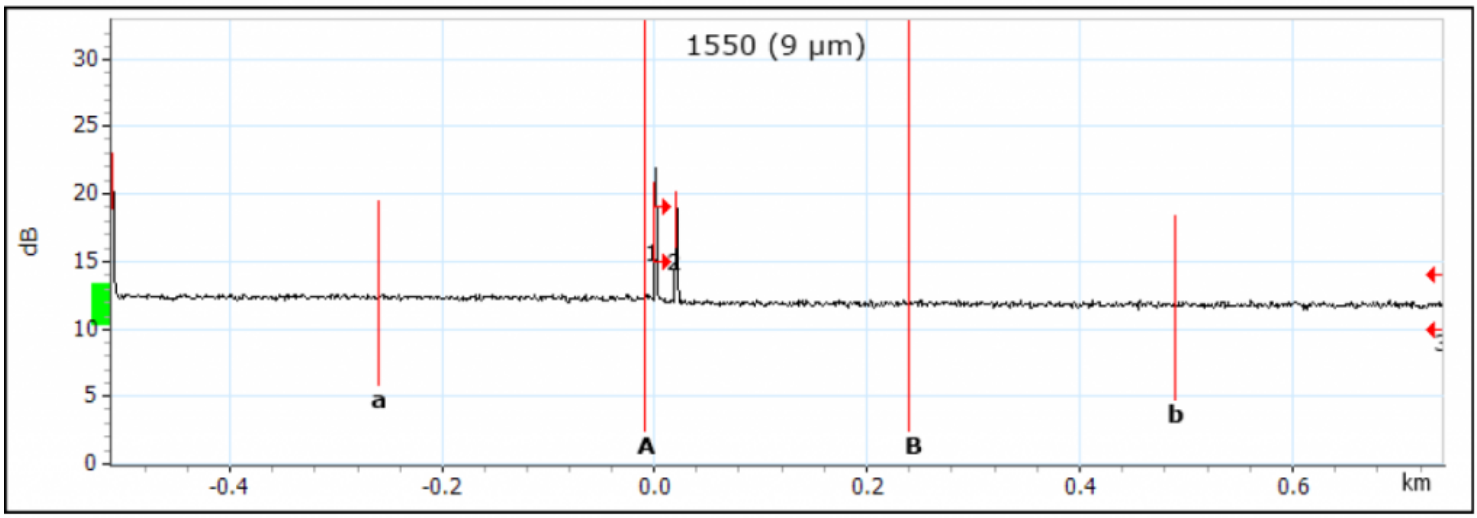

Markers

Marker	Position (km)	Value (dB)
a	-0.2600	12.651
A	-0.0101	12.648
B	0.2399	11.897
b	0.4899	11.634

A-B LSA attenuation:	4.161 dB/km	A-B average loss:	3.004 dB/km
A-B LSA loss:	0.125 dB	4-point Event Loss:	0.271 dB
A-B ORL:	46.03 dB	Maximum reflectance:	-48.8 dB

Now we have two connectors, one in the amplifier and one in the ODF leaving the site up to the grid. This lowered the reflection (-48 instead of -43) which makes the RAMAN real happy and we also gained a full dB just on the first 500m of fiber, and with long-haul networks that is a significant amount. This also caused the RAMAN to not fail its self-test and it booted properly and achieved almost full efficiency immediately, which is what the design is based around. From now on – this is the solution we are proposing for these type of sites.

We are not exactly best-friends with site-owners and installation crew since we are very anal about people making things correct. We are going to be living with this fiberplant for 10-15 years and we have one chance of making it good, and that is now. Even if most of the network will actually fully work with all of these small performance degradations everywhere we want the ride to be as smooth as possible for the years to come and breaking apart fiber for cleaning, patching, and measurements when it's already taken into full production is very tedious.

Any optical connector in our long-distance network that has not been cleaned, scoped (and passed through a connector-tester), photographed, documented and verified with OTDR will not be accepted as delivered.

Download the before-and-after OTDR-reports below.

OTDR Before

OTDR After

Sites

We have visited and inspected 86 out of 108 sites that's been delivered, with most of the un-visited sites being along the E45, which we usually call the "third leg" so we are not going to prioritise that path just now since we have more pressing matters to attend do. The results of our travels has been very mixed. A few sites has been extremely well built there was not a single thing to complain about and we could accept the delivery of them directly. Some other sites have we not been so lucky with. We have encountered everything from missing labels which can be categorized as a trivial fault, to possible life-threatening installations of the power-distribution which immediately declines the site as delivered. The idea on what we do is that we are always trying to stay a week or so in front of our installation-team from sub-contractor IPNett that drives around with trucks full of equipment. If we find the faults before they do, and whatever we find can be fixed, we save both a lot of money and time since we need to have the installation-team to be on time and not be delayed.

This is an example of a site which was filled to the brim with crap. Unacceptable of course and it needs to be cleaned out before we can install valuable equipment in here. **FAIL**

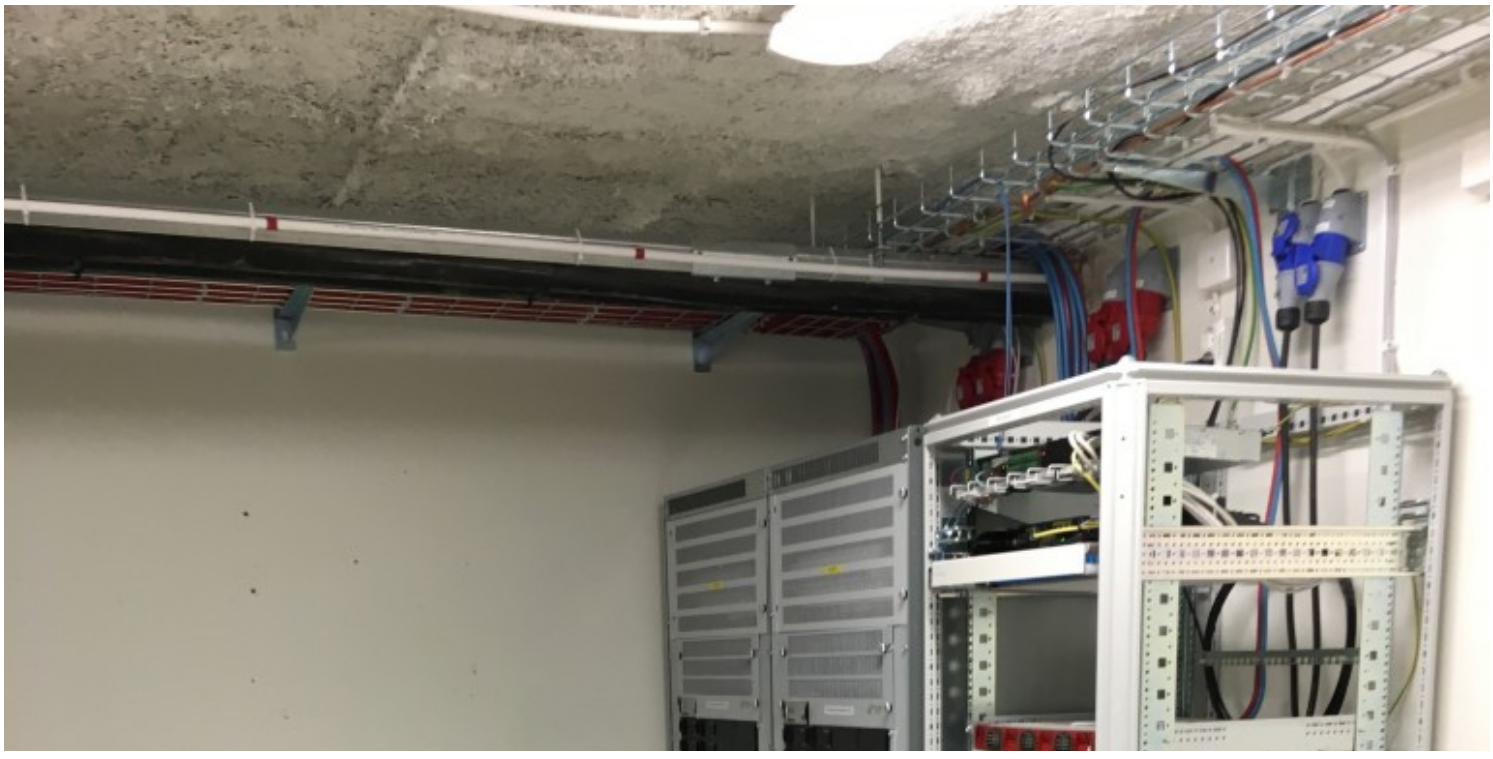

This is a site is in Borlänge, here we couldn't find a single-fault and site got accepted as delivered and we feel comfortable paying for it.

At the moment there is about 100 or so outstanding issues but we are getting them solved one by one and hopefully the critical ones will be solved swiftly to not delayed equipment rollout.

Equipment

Deliveries is coming and going everyday. At this very moment (thursday 30th of march) we have three different teams out in the country with trucks and equipment. There is one team installing around Sundsvall, one in the Stockholm region and one around Jönköping. These teams are installing the optical equipment primarly but we have also started with the IP-equipment aswell. The router is primarily being delivered this and next week but we have gotten a few items already.

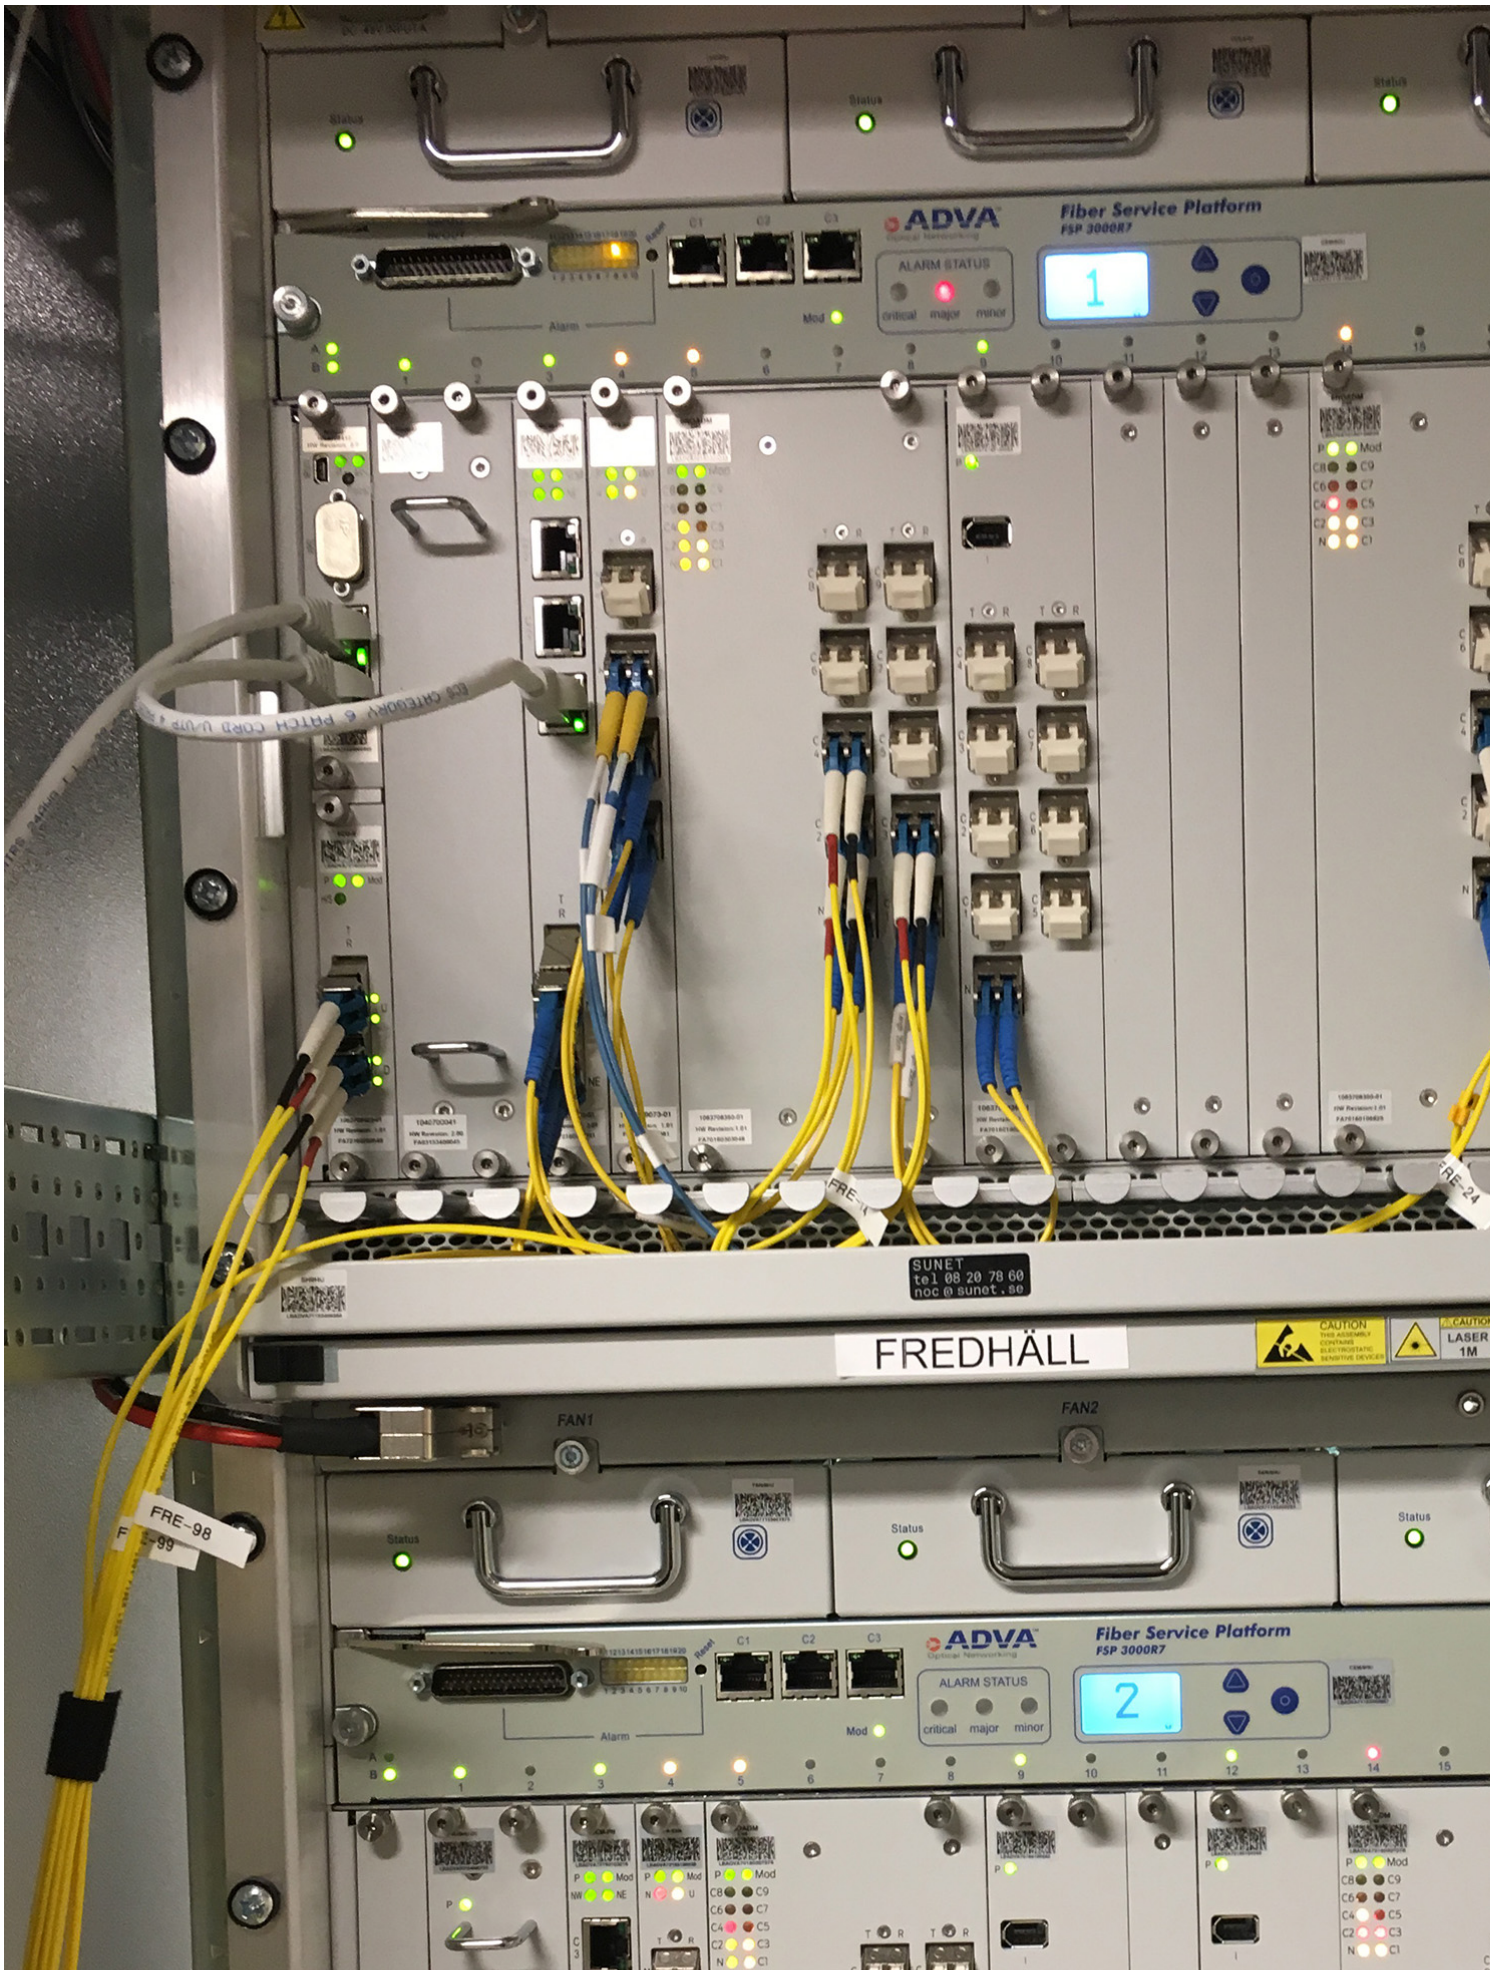
ADVA
Fiber Service Platform
FSP 3000R7
ALARM STATUS
critical major minor
SUNET
tel 08 20 78 60
noc@sunet.se
FREDHÄLL
FAN1
FAN2
FRE-98
FRE-24
CAUTION
LASER
1M

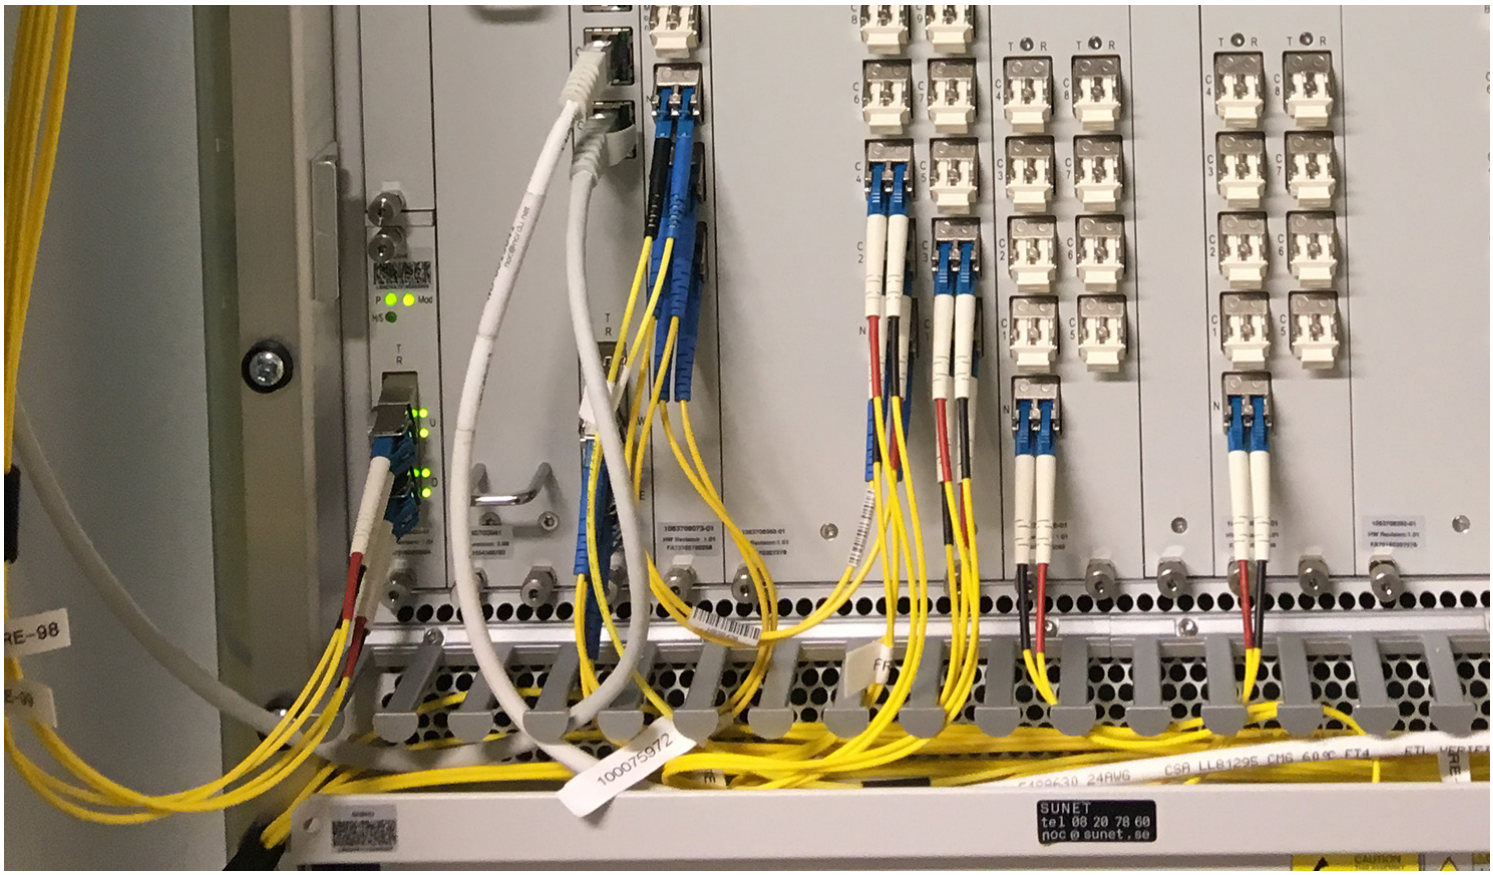

First router in place in the new network is the MX2020 that's being placed in Fredhäll, a real monster of a router. We have not been given the go-ahead to power it on just yet from the electricians so even if its the first installed router, it will not be the first router to be powered on in the new network.

Plenty of space in a Volvo V70 when hauling Juniper MX960s out in the outbacks.

In the second image the cards is installed into the router. The two leftmost and the two rightmost cards is the 100G DWDM cards, the first two cards will connect to the core-router in Örebro, aswell as shoot a 100G directly to Örebro University as well, there bypassing all active equipment in the Örebro site. The other two DWDM cards will go to Stockholm over two different optical spans, one coming in north-west of stockholm along E18 and one coming in south along E20. This means that Västerås will have three different paths to choose from. The third card is also a MPC3E-NG card but with a 100G MIC with regular LR4, this will connect one of the routers at Mälardalen University (the other uplink will be a coherent wave going directly to Stockholm from the campus, bypassing active equipment on site). In the 12th slot there is a mix of 10G and 1G interfaces for some general use to produce smaller

connections to other customers, such as branch-offices, student-organisations and on-site equipment (switches, servers, monitoring, etc)

The access to this service is forbidden since the specified partner is blocked

This is a movie from Västerås, the router is the production-router in the network to be powered on. Im also showing here how we solved the problem of small sites and full back-and-front reachability into the rack with sliding rails and cable-chains coming in from top.

Next week we hope to be able to roll out 6 more of these core-sites and then keep a steady pace so the full core-network can be completed before summer.

Skriven av

FREDRIK “HUGGE” KORSBÄCK

Network architect and chaosmonkey for AS1653 and AS2603. Fluent in BGP hugge@nordu.net