

SunetC 200G blogg



SunetC 200G

Magnus Bergroth posted on May 25, 2020


It's been a while since we did anything larger with the SunetC network. The network has been running since 2016 with it's ups and downs.

We needed to upgrade the network for upcoming projects, one is the ESS project in Lund. To give the 100G connections to Denmark with redundancy we needed to add an additional connection between Sunet and Nordunet in the south and also upgrade the existing connection to 2x100GE. We stated to plan this in November - December 2019.

The current setup Sunet uses with integrated DWDM optics in the routers, Junipers Cordoba Card is not that modern anymore and getting one 100G out of one slot in a MX960 is a bit of waste of resources. We decided to go for something newer so we went for the Juniper MX10003 as routers and Adva Cloud Connect as transponder/muxpoder.

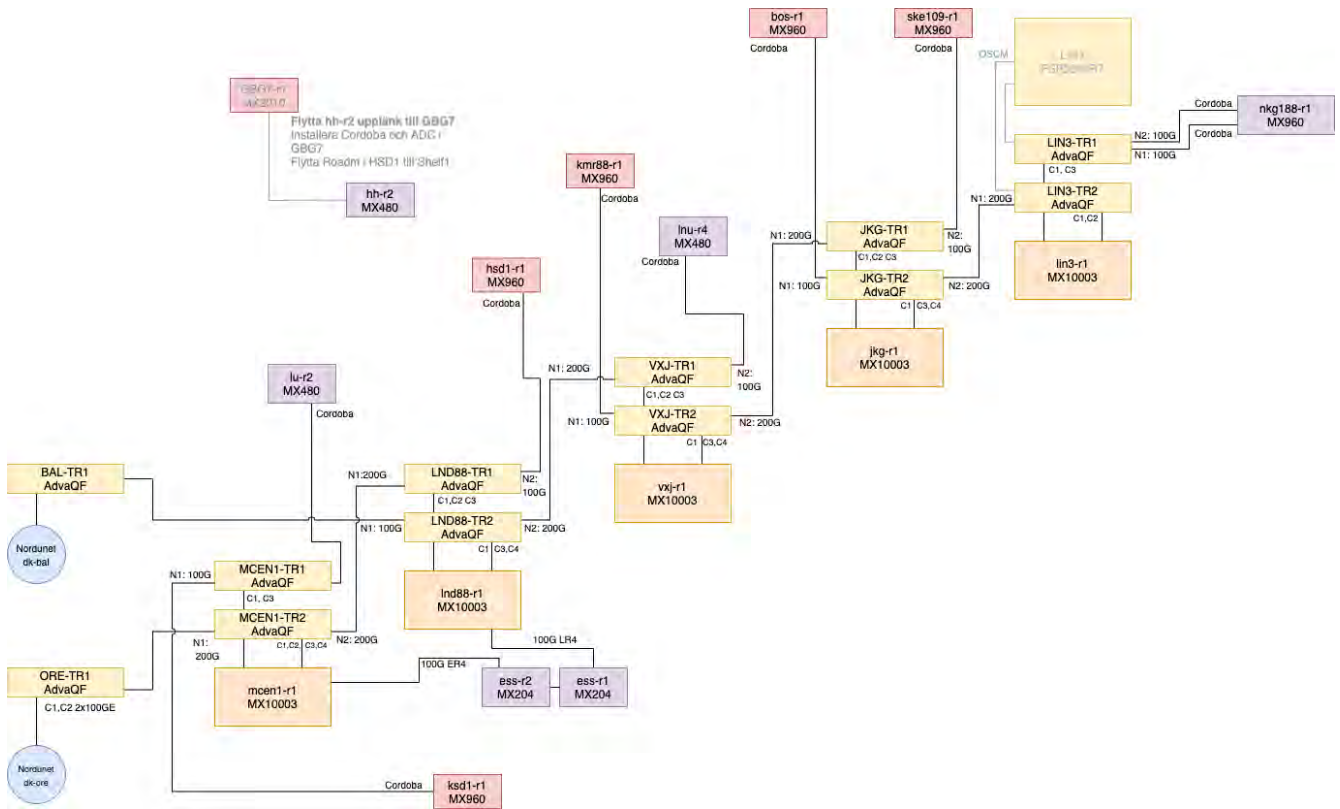
The idea is to replace the leg between Stockholm and Copenhagen with new routers and upgrade that section to 200G and then use the Juniper Cordoba cards that is freed up to upgrade another part of the network. It might be Stockholm - Göteborg or Stockholm - Umeå we will know shortly it depends on the fallout of other projects.

The cities that will get new equipment are Linköping, Jököping, Växjö, Lund, Malmö, Örestad DK, Ballerup DK.

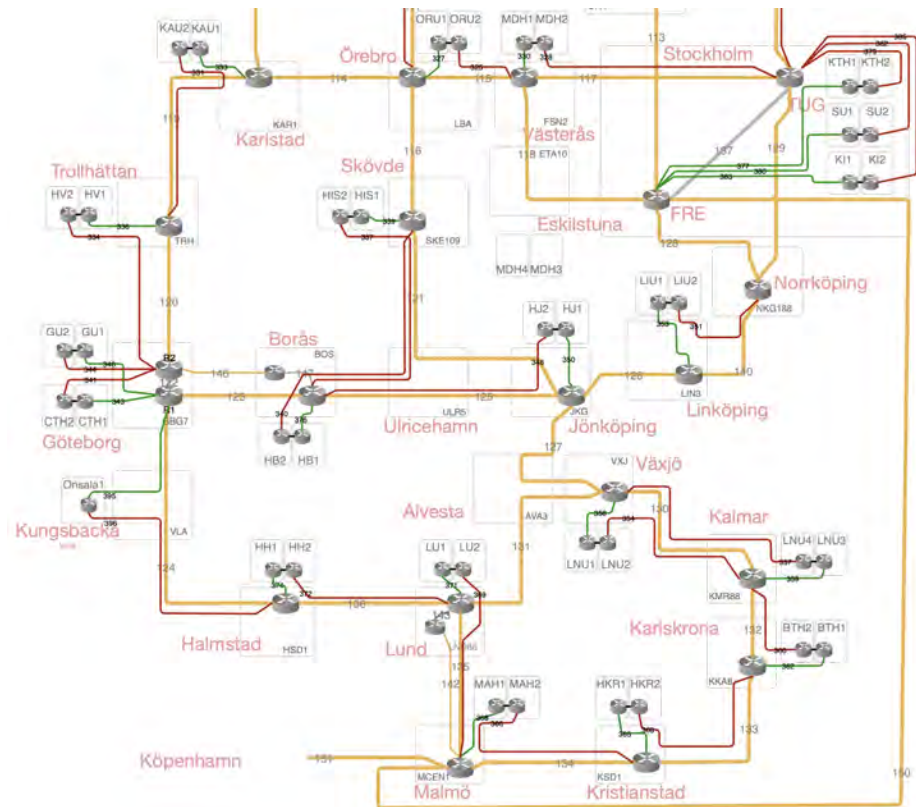


Design

We made a design where we put one mx10003 and two CloudConnect with one quadflex card in each city. The Quadflex card has two line connections that can run in 100G or 200G mode and four 100GE QSFP28 ports. The Adva box can run in a backwards compatibility mode so it can talk to the Juniper Cordoba card. It's the same chipset manufacturer in both the Adva box and the Juniper card.



The SunetC design connects the universities in each city with one gray LR4 100GE connection between the university and the City pop and one DWDM link that goes to the next city.



Norrköping already has 2x100G to Stockholm, sure it's to two different pops but we decided that it counts. We took a 100G DWDM card from our lab and installed in Norrköping MX960 router to get 2x100G to Linköping. The lab will get it back when we replaces the other MX960.

If you look closely on the network diagram you will see that there are 5 waves going into Lund. But we have only four DWDM ports in the Adva boxes. To solve this we are moving Halmstad universites DWDM connection in the other direction and it will got to Göteborg instead of Lund. The Göteborg router is an Juniper MX2010 and it needs an adapter card to run the DWDM card. That card will be taken from the Malmö router, but is currently borrowed from a router that is waiting to be installed in Stockholm.

Installation

The installation was planned to be made in March 2020, but the Adva gear got delayed due to what everything gets delayed this spring. We were doing the installation our self with multiple people from our staff. But as it is, it's now done in the field by me and the configuration done by my co-workers working from home.

We started by racking the routers in Stockholm and upgraded them to a Junos version that could run 1GE on a QSFP 40GE port. The Juniper mx10003 linecard has 6x40GE QSFP and 12x100GE QSFP28. We have a few connections at 1GE so we are connecting them with a QSFP to sfp adapter with a 1GE SFP.

No configuration were added to the Juniper routers, all that was done remotely. The two MX204 for ESS where pre-configured though.

The Adva boxes were also upgraded at the Sunet Office in Stockholm and got an IP-address so we could connect them to the ADVA DCN network.

We started we a test installation in Linköping, were we installed a MX10003 and the two ADVA cloud connect boxes. To figure out what could go wrong with the installation.

The first stop where in Norrköping to install the additional Cordoba DWDM card. We had that connected to the new router in Linköping to have them running in parallel.

There were plenty of space in the rack in Linköping:



The small issue is that the Juniper MX960 and the MX10003 has different distance between the front and the back rack post. And the installed equipment made it impossible to move the posts.

A trip to the local company Biltema were some M16 long nuts and long m6 bolts over bridged the 75mm gap.



Back home it turned out that the new software loaded on the ADVA boxes were not supported on the ADVA NMS that we had installed. To not fly blind it was decided to upgrade the NMS.

This got the installation delayed by three weeks.

Field trip

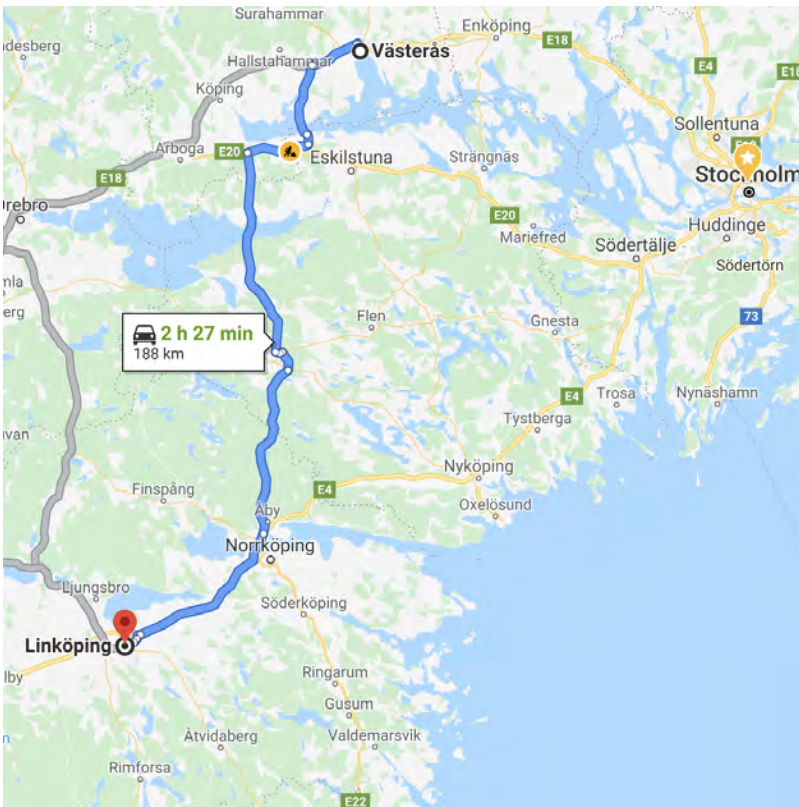
With the NMS upgraded we started the trip to upgrade the network. The routers and the Adva boxes and what ever needed was loaded into my car and to have somewhere to sleep and eat a trailer was added to the car.

2020-05-11 07:30




The first destination was Linköping to complete the migration.

2020-05-11 10:40



The pop in Linköping is a small pop



The racks are on rails so that you can go in behind them and as the PDU:s are on the back of the rack it's the place I spent the most of the time. Fibers and power are going in chain-tracks on top off the racks.

The Adva boxes and the the router had already been installed but there were still work to do.

The power needed to be completed, both routers are powered on at the same time to minimize the downtime during the migration. The new mx10003 are running on B-power and the old MX960 on A-power. The power cable were on backorder when the router was first installed so only two of six feeds were cabled.



This is the A-side of the MX10003. Those safety plastic covers are a joke, they snap on but they do not sit very tight so if get close with a hammer or any other tool they get lost.

On to the migration. We have serial out off band connection to each pop via 4G and an OpenGear terminal server that we reach over an OpenVPN tunnel. Each box has four serial and four ethernet connections. The console from the MX960 was moved from RE1 to RE0 on the MX10003 and could be configured from remote. Dennis Wallberg did the configuration from home while I was moving cables and completed the power work.

We did as we do a normal software upgrade of the router. Set ISIS overload advertise high metrics and disable all customer eBGP session and then wait for the traffic to go away. Then look one more time on the configuration as not all the traffic went away to find eBGP session inside a VRF.

Then disable all the interfaces, we decided to disable the interface instead of turning off the router. This gave us the ability to revert back fast if it would turn out that this box where a single point of failure due to some other errors in the network.

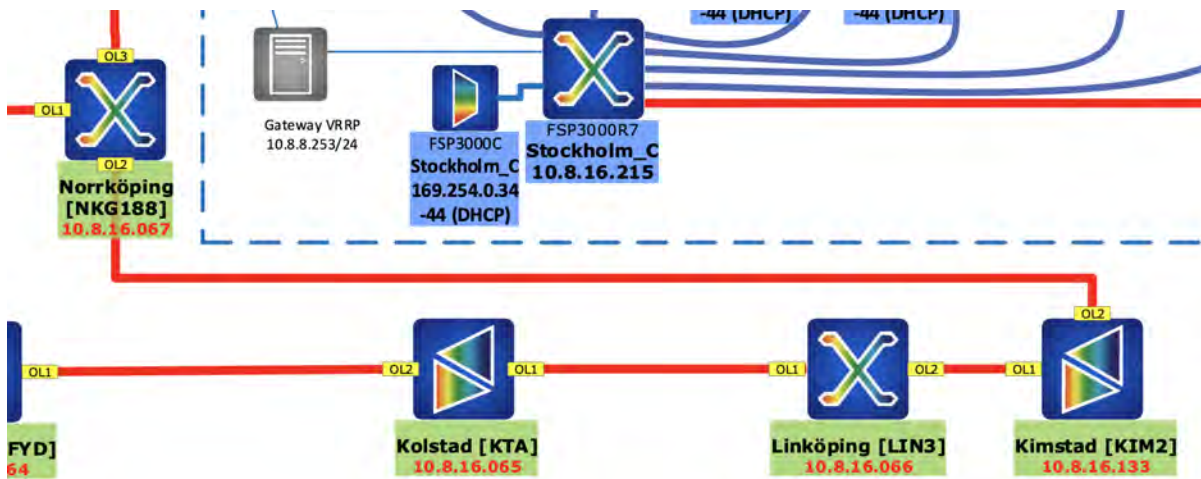
As we already had provisioned a new link between the mx10k3 and nkg188-r1.sunet.se we just setup an Lag in nkg188-r1 with that link and the box were connected to the rest of the core..

```
bergroth@nkg188-r1-re0> show isis adjacency
Interface          System          L State          Hold (secs) SNPA
ae1.1              lin3-r1-re0    2 Up              22
```

The link between Norrköping and Linköping are two 100Gbps DP-QPSK 15% FEC on the line side. On the client side it's standard 100GE that are bundled into a LAG, running LACP.

```
bergroth@nkg188-r1-re0> show lacp interfaces ael
Aggregated interface: ael
LACP state:      Role   Exp   Def   Dist  Col   Syn   Aggr  Timeout  Activity
et-1/0/0        Actor  No    No    Yes   Yes   Yes   Yes    Fast    Active
et-1/0/0        Partner No    No    Yes   Yes   Yes   Yes    Fast    Active
et-9/0/0        Actor  No    No    Yes   Yes   Yes   Yes    Fast    Active
et-9/0/0        Partner No    No    Yes   Yes   Yes   Yes    Fast    Active
LACP protocol:  Receive State  Transmit State  Mux State
et-1/0/0        Current        Fast periodic  Collecting distributing
et-9/0/0        Current        Fast periodic  Collecting distributing
```

The Distance between Norrköping and Linköping is only 50km, but the city fiber was so bad so we couldn't run Raman amplifiers, so we added an amp site in between.



Looking on the optical values of this short link in the juniper router:

```

bergroth@nkg188-rl-re0# run show interfaces diagnostics optics et-1/0/0
Physical interface: et-1/0/0
...

Tx power                               : 1.000 mW / 0.00 dBm
Rx power (total)                       : 0.407 mW / -3.90 dBm
Rx power (signal)                     : 0.054 mW / -12.68 dBm
Lane chromatic dispersion              : 1109 ps/nm
Lane differential group delay           : 5 ps
Lane Q2 factor                         : 14.20 dB
Lane carrier frequency offset          : -835 MHz
Lane electrical SNR                    : 14.80 dB

```

And the same value on the Adva box:

```

admin@LIN3-TR1> show interface 1/1/n1 opt-phy pm current

pm current table:
mon-entity interval pm-profile max-intervals tca-transient suspect elapsed
opt-phy live IFQFnw 0 False not-suspect 2573734 seconds
pm-name loc thr-low -high -range -def curr-value
opt-rx-pwr nend -33 13 -33..13 global -3.5 dBm
opt-tx-pwr nend -3 9 -3..9 global 0.0 dBm
opt-lbc nend N/A global 100.0 mA
laser-temperature nend N/A global 35.6 C

```

The optical receive power at -3.5dBm is actually the power of all waves going in to the box and not only the wave that the port is tuned to.

The optical levels are not that interesting and can be compensated in the DWDM networks Roadm to some extent.

The other optical levels are a bit more interesting


```
admin@LIN3-TR1> show interface 1/1/n1/ot100 och pm current
```

```
pm current table:
```

```
mon-entity interval pm-profile max-intervals tca-transient suspect elapsed
och live RxQuality 0 False not-suspect 2573926 seconds
  pm-name loc thr-low -high -range -def curr-value
  q-factor nend N/A global 16.5 dB

mon-entity interval pm-profile max-intervals tca-transient suspect elapsed
och live RxQFnl00g 0 False not-suspect 2573926 seconds
  pm-name loc thr-low -high -range -def curr-value
  chromatic-dispersion-compensation nend -70000 70000 -70000..70000 global -1170.0 ps/nm
  carrier-frequency-offset nend -2.5 2.5 -3.2..3.2 global -0.85 GHz

mon-entity interval pm-profile max-intervals tca-transient suspect elapsed
och live ImpQFnl00g 0 False not-suspect 2573926 seconds
  pm-name loc thr-low -high -range -def curr-value
  signal-to-noise-ratio nend 6 6..16 global 17.6 dB
  differential-group-delay nend 100 0..100 global 3.0 ps
```

Adva sees the Q-factor as 16.5dB and the juniper as 14.2dB. SNR Adva 17.6dB and Juniper 14.8dB.

My guess that the integrated optics in the Adva QuadFlex compared to the Juniper CFP2 ACO makes a big difference here.

CD Adva reports it as a negative value and Juniper as positive, but at least they agree on the number.

Mikott had configured new optical tunnels in the Adva DWDM network with new wave-lengths for all the new paths. The other link between Linköping and Norrköping were brought up by changing the wavelength in the Juniper router and activate the new tunnel in the Adva DWDM node. The reason for new wavelength has nothing to do with the upgrade project it's just that we would like a better wavelength plan, the current plan starts with the first wave of the spectrum and that's not optimal.

The core link from Linköping to Jönköping were also configured as a 100G link for the moment, it will be reconfigured to 200G 16QAM when the JKG site has been upgraded.

The two client interfaces were moved along with the SFP, the 1GE port needed some extra config to for the QSFP to understand that's a one gig.

```
bergroth@lin3-r1-re0> show configuration interfaces xe-0/0/0:0
description "SU-S003542 link to SFHM, sunet-sfhm";
gigether-options {
  no-auto-negotiation;
  speed 1g;
}
```

And it's called xe- and not ge- as 1ge normally are done in Junipers.

Left to do on the site where to get the mx960 out of the rack and into the car.

The Line cards are a bit heavy so they were temporarily located in the driver seat.



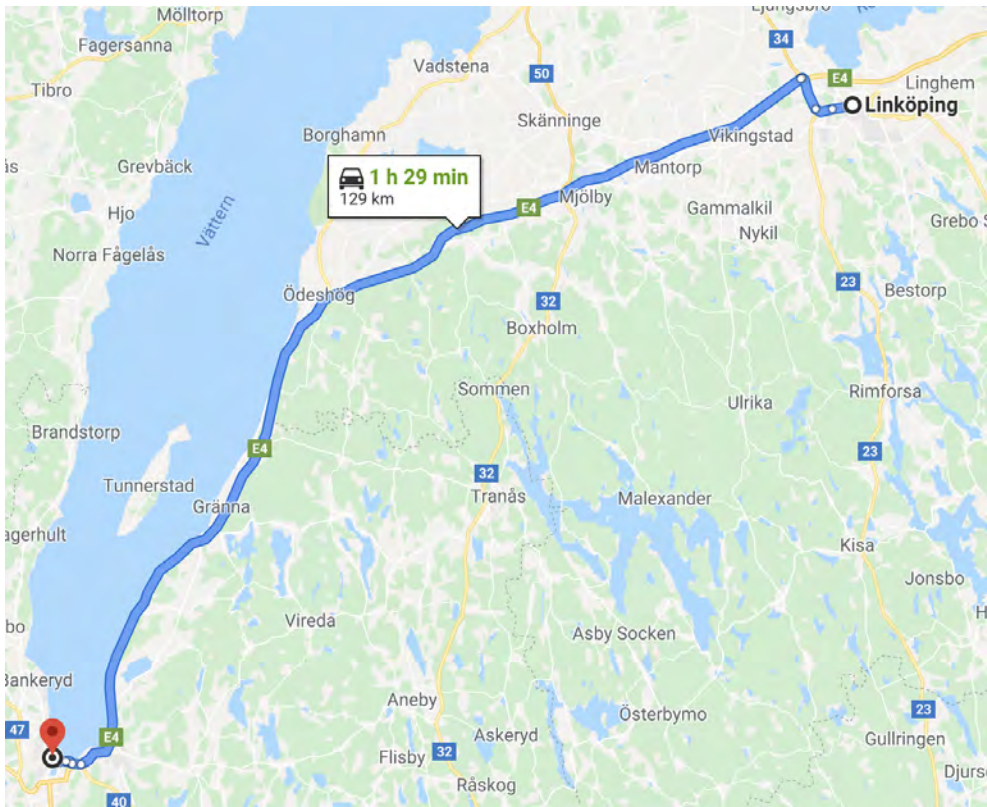
and then moved back to the chassis in the back



from the Juniper mx960 documentation "Standard chassis with components removed: 150 lb"

2020-05-11 14:55

Left Linköping for Jönköping, dropped off the mx960 on a small detour.



2020-05-12 06:30

Jönköping site. The access to this site is inside a loading tunnel so I couldn't bring the caravan to this site.



The steps up to the door made it a bit fun to get the Juniper mx10003 loaded. There is also no where to park your car so you have to unload your gear and hope that you don't forget anything, which I of course did.

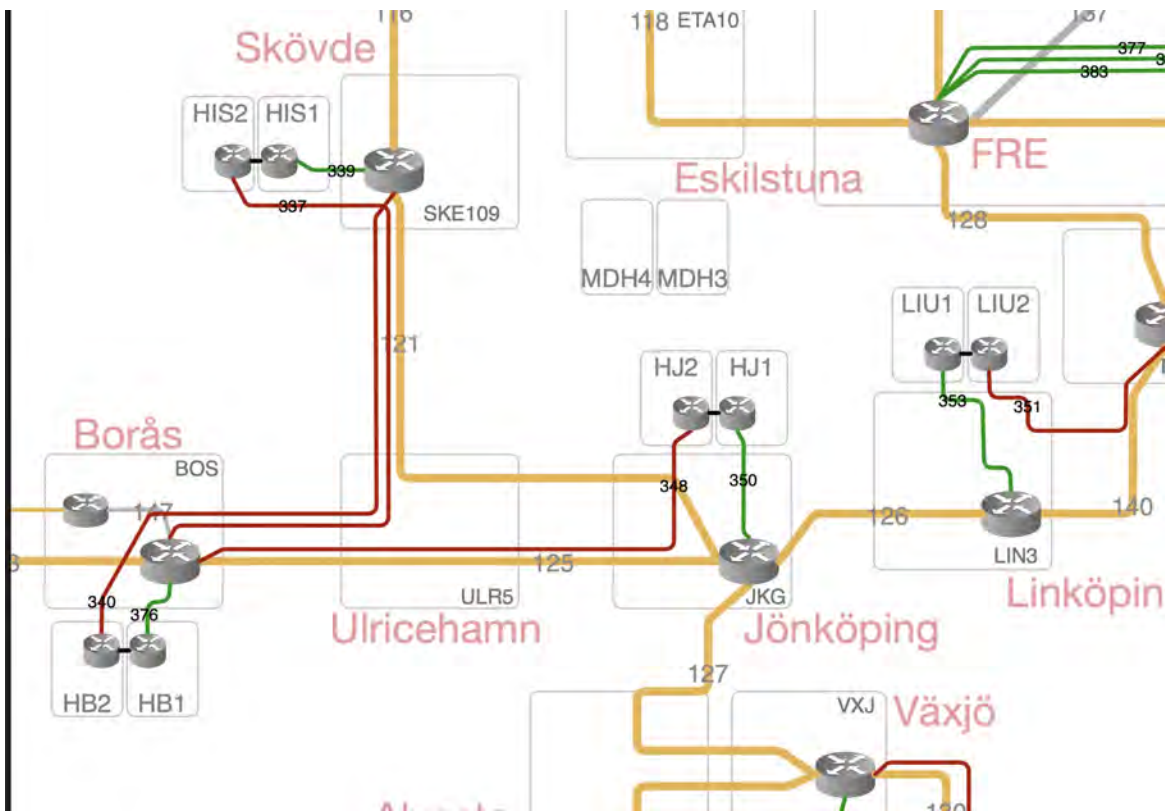


This site is bigger so there are no rails on the rack.

The rack had the same problem as the rack in Linköping, the posts were to close together. The "shims" were upgraded to m15 black iron pipe nuts.

But I forgot to clean them, they are rust prevented with oil, so it got a bit messy.





Jönköping has four core dwdm connections and no dwdm university connection. The links to Skövde and Borås are branches to the 200G Stockholm Copenhagen path remains on 100Gbps mode.


When we configure the Adva DCI boxes we set the link going south, in this case towards Växjö, to run on the wavelength that the juniper dwdm cards already uses. That makes the switchover faster as we then don't need to change anything in the DWDM system or on the router on the other end.

The south link will also be the first link to be migrated to the new router as it only requires the fiber to be moved from the mx960 to the adva QF port for the new mx10003 to get connectivity yo to the core. It gives the routers enuff time for iBGP to converge before we connect the rest of the links. When all the BGP prefixes are loaded in FIB we start migrating any single homed users, Jönköping didn't have any of those so the next step are the other 100G links.

There are three steps to get a link up and running.

1. Move the fiber from the local mx960 router to the correct Adva network port and clean the fiber port. One could think that fiber ports with nice metallic flaps could all be clean from the factory. Most were but some looked like this.





Mantra to learn: Scope, clean, Scope

2. Change the wavelength on the juniper dwdm card. Adva specifies the wavelength as frequency (Hz) and Juniper in nm, so you need to get them right.

Defined on the Adva DWDM system: Channel 19380

| Remote Device/ Channel | Identifier | Channel Bandwidth | Rate | Frame Format | Forward Error Correction | Far End Location | User Label |
|---------------------------|---------------|-------------------|--------|--------------|-----------------------------|------------------|----------------|
| et-0/0/0 | EOM-1-9-C1 | | | | | | |
| JKG-TR2 N2 | EOM-1-9-C2 | | | | | | |
| et-1/0/0 | EOM-1-11-C1 | | | | | | |
| JKG-TR1 N1 | EOM-1-11-C2 | | | | | | |
| et-10/0/0 | EOM-2-9-C1 | | | | | | |
| et-11/0/0 | EOM-2-9-C2 | | | | | | |
| JKG-TR1 N1 | EOM-2-9-C3 | | | | | | |
| 19380 | VECH-2-9-C3-1 | 50 GHz | 111809 | Not Defined | None | BOS | JKG-TR1.bos-r1 |

Defined on the cli of the Adva dci, THz:

```
admin@JKG-TR1> show interface 1/1/n1 opt-phy

opt-phy:
  frequency: 193.80000 THz <---
  bandwidth: 37.500 GHz
  laser-state: True
  force-laser-active: 0 minutes
  force-laser-active-seconds: 0 seconds
  setpoint: 0.0 dBm
```

Defined on the same box but the web interface, MHz:



And then on the Juniper router, nm:


```
bergroth@bos-r2-re0# show interfaces et-0/0/0
description "SU-S002125 Link to jkg-r1, bos-r1.jkg-r1";
flexible-vlan-tagging;
optics-options {
    wavelength 1546.92; <--- nm
    tx-power 0;
}
```

But the show interface at least shows both:

```
bergroth@bos-r2-re0> show interfaces et-0/0/0
Physical interface: et-0/0/0, Enabled, Physical link is Up
...
Wavelength      : 1546.92 nm, Frequency: 193.80 THz
```

3. Bring up the wavelength in the DWDM system.

TNL-WDM-JKG-TR1.bos-r1

State Fault Fit Log Prov Oper Transp Recovery Paths PAF

Admin State In Service
 Signaling

Operational State: Normal

Secondary States:

Cancel Refresh Apply Abort Delete

Dependencies...

The pre configured Tunnel is set from "Signaling" to "In Service", what that does is that the system are using GMPLS to create a wavelength tunnel between the defined ports on the end nodes.

| | | | | | | | | |
|-------|-------|---------|------|------|--------|----------|-------|-----|
| State | Fault | Fit Log | Prov | Oper | Transp | Recovery | Paths | PAF |
|-------|-------|---------|------|------|--------|----------|-------|-----|

Tunnel Name: JKG-TR1.bos-r1
 Tunnel Number: 0
 Tunnel Number Extension: 168300607
 Tunnel Type: Point to Point
 Tunnel Scope: N to N
 TE Switch Level: Lambda
 Tunnel Template: No
 OIF UNI: No
 Alias:

Source Target ID: JKG
 Source Node: 10.8.16.63
 Source Equipment: VECH-2-9-C3-1

Destination Target ID: BOS
 Destination Node IP: 10.8.16.35
 Destination Equipment: VECH-2-7-C1-2

In this case a 8 port splitter combiner. VECH-2-9-C3-1 is a description for External Channel Service, shelf 2, slot 9, port C3, wavelength number 1.

After the tunnel is built, the system tries to set the correct optical levels for the channel, each Roadm block sets its wavelenght selected switch (WSS) to the attenuation that gives the pre configured power per channel on the outgoing port. It starts from one end of the tunnel path and does this over each Roadm in the path to it comes to the destination and then does it in the other direction. This is called equalisation. If the end points are tuned to correct wavelength and RX and TX are not swapped, aka light comes in to the system, the the equalisation successes and you get a link up.

TNL-WDM-JKG-TR1.bos-r1

| | | | | | | | | |
|-------|-------|---------|------|------|--------|----------|-------|-----|
| State | Fault | Fit Log | Prov | Oper | Transp | Recovery | Paths | PAF |
|-------|-------|---------|------|------|--------|----------|-------|-----|

| No. | Timestamp | AID | Condition | NC | Status |
|-----------------------|----------------------|-----------------------|--------------|----|--------|
| 83497 | 20-05-19,07:09:06.77 | CNX-WDM-JKG-TR1.bos-r | EQLZ-COMPL | NA | n/a |
| 83496 | 20-05-19,07:09:06.73 | CNX-WDM-JKG-TR1.bos-r | MODFCN-COMPL | NA | n/a |
| 83495 | 20-05-19,07:09:06.73 | CNX-WDM-JKG-TR1.bos-r | MODFCN-START | NA | n/a |
| 83490 | 20-05-19,07:09:02.26 | CNX-WDM-JKG-TR1.bos-r | EQLZ-START | NA | n/a |

The last step is to reconfigure Linköping to 200Gbps line side and 2x100GE on the client.

On the Juniper in Linköping you need to re-enable LACP on the LAG as we where running a 2x100GE LAG in Linköping with one interface down and a single normal no LAG 100GE interface on the MX960 in Jönköping as a intermediate step.

On the Adva DCI you need to put all interfaces related to the 100G link to "Out Of Service", Delete them and build a 200Gbps OT200 interface for the N port.

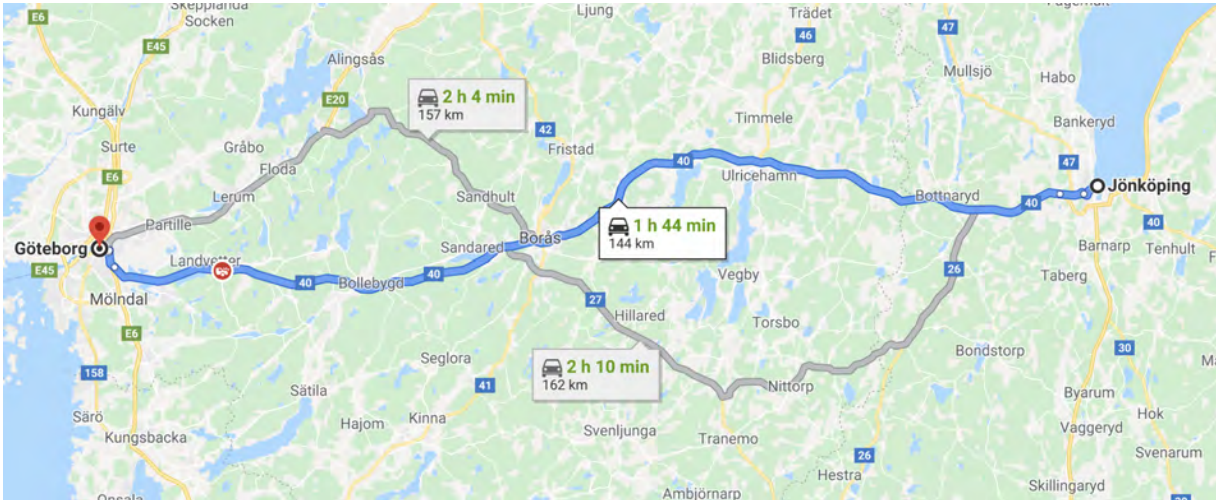
When all is up and running it's just to load the old mx960 in the car to be installed in Borås.

I found a loading scissor lift at the hotel loading dock that made it a bit easier to get the mx960 into the car.



2020-05-12 14:00

Left site, heading to Göteborg



2020-05-12 16:30

Meeting up with Per Anderson at Chalmers.

Rikard had last week fixed so that Chalmers got 100GE client interfaces. This was done by replacing the two Juniper mx480 they had with two new mx204. Two of the cards in the mx480 routers will be used for the install of the router bos-r2 in Borås and the chassis will go to Linne Universitets that will move to a new site in Kalmar.


Per and I unscrewed one of the mx480 from a rack and the other was already on a pallet. Both routers were then stacked in the backseat of my car.

2020-05-12 17:15


Site GBG7, As pointed out above Halmstad universities DWDM link needed to be moved from Lund to Göteborg. The dwdm card in the router gbg7-r1 was taken from the router in Jönköping and the adapter card from Stockholm.




This was supposed to be a fast in and out. Install a card, pull a fiber, move interface config from Lund to Göteborg, update IPAM, update DNS, update inventory, update map, bring up DWDM wave, take down old DWDM wave.



But the new wave didn't want to equalize, tried to swap RX and TX still no luck. Tested with a light meter and it showed -16dBm into the 8PSM card, it should be atleast -1dBm.



Turned out that the fiberpatch LC connector had a plastic burr from the manufacturing process that prevented the LC connector to be seated correctly in the CFP2 plug and the airgap created the extra attenuation.



solved with a leatherman

with proper levels as

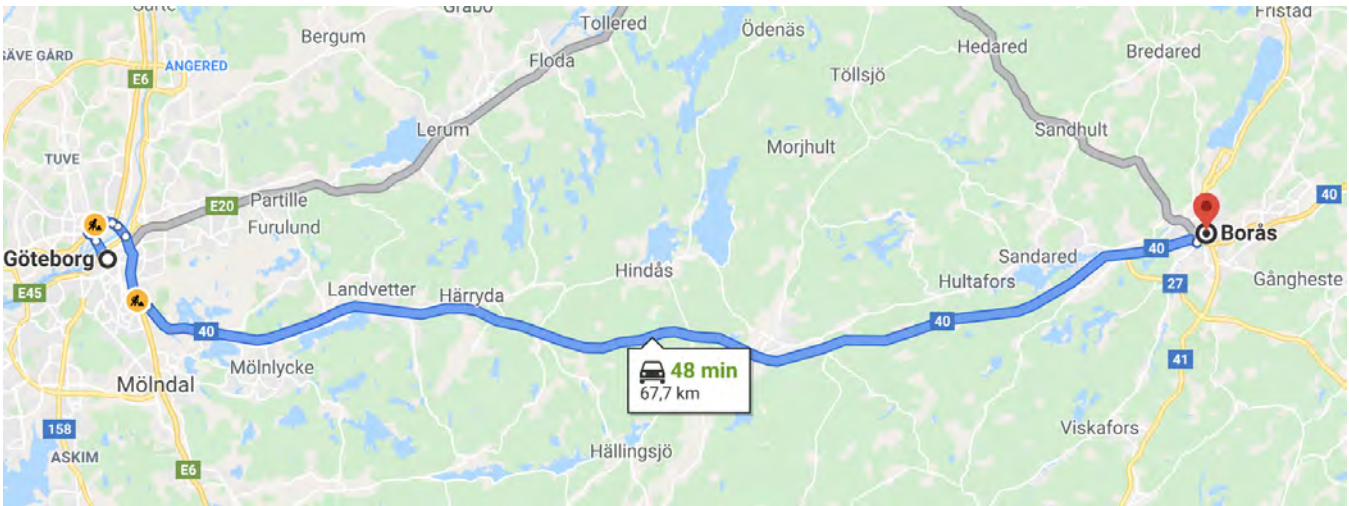


result

and a link up.

2020-05-12 18:30

Left Göteborg for Borås




The plan was to install the MX960 from Jönköping in Borås to get redundant 100GE router connections for an existing user.

2020-05-12 20:00

Borås, the site is located inside a hotel. Unfortunately couldn't the the hotel receptionist Emily find the key to the room the site is located in. After trying all the keys she had we gave up. To solve the logistical problems with moving routers around I had to stored the MX960 in their office.



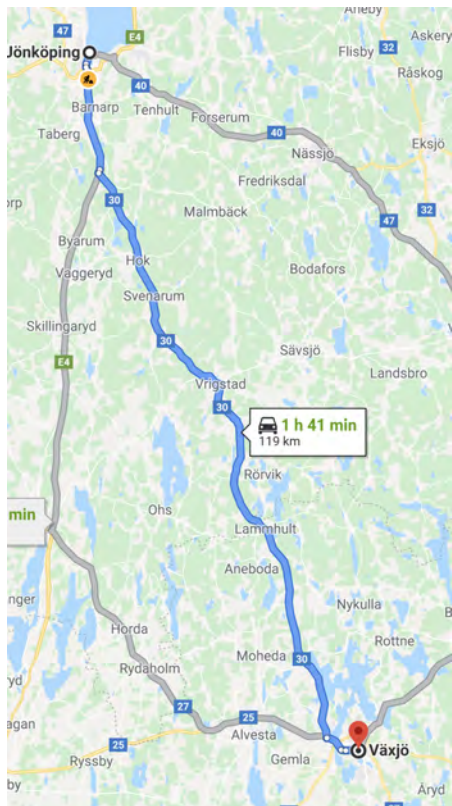
2020-05-12 22:00



Back in the Caravan in Jönköping

2020-05-13 08:30

Left Jönköping for Växjö



2020-05-13 11:30

At the VXJ site, parked next to the tracks inside the gates and fences




and some lunch



before entering the site




We are using DC power on all the core sites, it makes creating backup power very easy as it's just a bunch of batteries in series across the power feed. The Voltage is -48DC, the positive polarity is connected to the neutral/ground like in an old British MG with Lucas power system. The Adva DCI boxes has two power feeds A and B power




There are three cables into each power block + - and Ground. The colors for DC cables is a bit of a mess. This cable loom from Adva has black for positive +, blue negative - and Ground Yellow and Green (this is a standard in Europe at least).


After racking the boxes the first thing is to connect those ground wires to the ground bar in the rack. The ground bar is the bare copper strip in the left corner of the rack.



the cables are cut to length



stripped




and crimped with shoes





and then placed between the tag washer and the normal washer

Then repeat for all the others cables



When the ground is done it's time to connect the power feeds. The PDU for the DWDM gear is feed with a 63A breaker and 25mm² dual insulated red(+) and blue(-) cable. That feed is then divided to the loads via smaller breakers, the DWDM shelf gets an 40A breaker and the DCI boxes 16A breakers.




The blue and black cable from the ADVA DCI box are dual insulated cables as required by local installation laws. It means that the insulation is rather thick and that the end protection sleeves that we use on the cable will not fit over the isolation. To fix that you have to cut the insulation twice.




First the correct length of the bare wire


then the length of the outer insulation




then crimp the end sleeve




and you get a protected cable end to put



in the breaker



The black(postive) cable goes into the common bar in the back of the PDU and the blue cable to the breaker.




You can connect the cable without a sleeve like the red cable in the picture but there is a risk that some of the strands are not seated correctly in the breaker and makes a short circuit.

2020-05-13 18:00

Left the site with the caravan to go to the university and hand them the two MX480 from Chalmers, they also got two cards from the vxj-r1 router.

Back to the site, we had problem with getting one of the ADVA DCI boxes to transmit light, the only indication was an "non" traffic affecting alarm saying low Tx power at -40 dBm. Warm reboot or cold reboot didn't help so we replaced the card with one that I had in the car and that solved the problem.

The rack with the mx960 removed and all the gear up and running



and the router on a wheel-board




Bumpy unload



2020-05-13 20:50

Left the Växjö site for Lund



2020-05-13 23:30

It's been chilly during the day and that affected the temperature in the caravan.



degrees. but there is 3kW of electrical floor heating so it only takes a few hours to get to 20

Not that many others that do camping in May



It was windy so there was people doing kitesurfing



2020-05-14 09:00

Off to ESS to give them their new MX204. I was not allowed to enter the building site, but Rikard had prepared all the configuration so all the ESS staff needed to do was to install the boxes and swap the fibers.



2020-05-14 09:50




At the LND88 site in Lund

The joy to realise that there are no elevator to the



basement.



This is a smaller site

so the rack rails are back.

The OOB opengear router was broken in this site so we started to replace that one, I did the cabling and set the ip-address on the wan interface and Dennis did the rest of the openvpn server configuration on remote.



The Mx10003 has a separate grounding point on the side of the box that needs to be used when you run the router on DC power.



But as you can see it's inside the rack rails, so you can't mount the cable until you have mounted the router in the rack and if you have a rack with closed side you will not fit normal a screwdriver. Stupid design.

To get the router into the rack I mounted a small shelf to let the router rest on until it get screwed in.



The MX10003 has six DC PEMs, you can run it in two but it will complain and give you a chassis alarm error.



PEM 0, 2 and 4 on the left side

and PEM 1, 3 and 5 on the right




Here we use the same colors as the site feeds, Red + (positive) Blue - (negative)

To have the both routers running at the same time the B-feed is disconnected from the PDU for the MX960 and the MX10003 are connected instead.

The feed to the PDU is 2 x 63A A-feed and 2 x 63A B-feed. The MX10003 accepts 3 x 32A A-feed and 3 x 32A B-feed. Additional breakers were installed into the PDU to divide one 63A feed to two 32A feeds and the other 63A feed to a 32A breaker so the 6mm² cables to the router can't be overloaded.



As the PDUs are not built the same, all the interconnect cables had to be made on site.

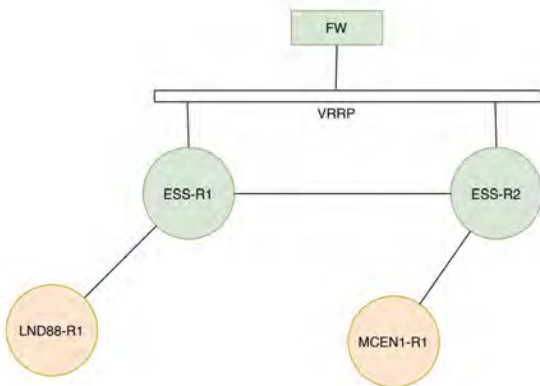


When the MX960 is turned off and disconnected can the A-side of the MX10003 be connected to the PDU.

Migrating ESS

ESS has two MX80 and each router has a connection to Sunet, a link between the routers and an interface towards ESS internal network. On the link to their internal network are they running VRRP. eBGP runs between ESS and Sunet and iBGP between the two routers. They also has a bunch of L2VPN services.

Traffic was removed from the ess-r1 router by making ess-r2 VRRP master and taking down the eBGP session between ess-r1 and Ind88-r1..



Ess staff was then asked to power off ESS-R1 and move the fiber going to LND88-R1 to the new mx204, the port was changed at the same time in Ind88-r1 to a new 100GE LR4 port. When the connectivity was established between the routers could the ESS staff move the client interfaces to the new router one by one.


ESS-R2 was set to be moved the next day after we had migrated Malmö mcen1-r1 from a MX2010 to a MX10003. The link between ess-r1 and ess-r2 was never established as it now had a 10GE in one end and a 100GE on the other end.

Then the fun of getting the router out of the rack,



Removing all the cards from the router makes it light enuff

to drag it out on roller board



and roll it out to the stairs

w



here it magically jumps up the steps

a



nd into the car

2020-05-14 18:10

Left site for an early evening




With a view of Denmark in the distance

and some pasta for food



2020-05-15 07:30


MCEN1 Malmö site



A larger site where we have three racks., 120cm deep compared to 100cm on the other core sites. The idea was that we should have the ability to place servers on four sites, Stockholm Göteborg Malmö and Luleå. So far we only have servers in Stockholm and Luleå. We also have dual AC 230V power feeds two that extra rack.



The MX2010 takes up the space of one rack



the DWDM the other rack



and the last rack will be used for the new gear

There are no PDU in the empty rack so the power cables had to be pulled from the MX2010 rack over the rack tops



The MX2010 was connected with 5 x 63A but only 2 of the feeds were used for the new mx10003



Ran out of red and blue wires and the only thing that was in stock on short notice was single dual insulated red and black cables.



Configuration of the Adva boxes

The Adva DCI can run in standalone mode and integrated mode together with the DWDM shelf. We are running them in standalone mode.

In standalone mode the configuration is done via CLI and some basic web interface. We been told that the web interface will be better in the upcoming releases, let see.

When we got the boxes they were all running release 2.2.3, Jonas did the job of upgrading them to 3.1 at the stockholm office and configuring ip-addresses on them.

The cli on the box is very much like doing API calls, it takes some time to get use to.

This is how you do a ping, success just means that the ping command worked.

```
admin@MCEN1-TR1> ping address 10.8.16.4

[Success]

[node 1]
admin@MCEN1-TR1> show dcn 1 ping-result

ping-result: [
  address: 10.8.16.4
  result: PING 10.8.16.4 (10.8.16.4) 56(84) bytes of data.\n\n--- 10.8.16.4 ping statistics ---\n5 packets
transmitted, 5 received, 0% packet loss, time 10ms\nrtrt min/avg/max/mdev = 12.096/12.155/12.183/0.103 ms\n
time-stamp: 2020-05-19 13:42:09.1033 UTC
]
```

After you have done the upgrade, set NTP and loaded the box with interface licenses you need to configure the interfaces.

Esisest is to start with the client interfaces and configure the plugs first, the adva box is locked to ADVA QSFP28. You can apparently run with 3 party optics for a hefty licensing fee.

```
admin@MCEN1-TR1# set plug 1/1/c2 plug-name
```

```
Possible completions:
QSFP28-10X10G-850I-MM-MPO
QSFP28-10X10G-1310S-SM-MPO
QSFP28-103G-CWDM4-SM-LC
QSFP28-103G-CWDM4E-SM-LC
QSFP28-103G-LR4-SM-LC
QSFP28-103G-PSM4-SM-MPO
QSFP28-103G-SR4-MM-MPO
QSFP28-112G-AOC-0100
QSFP28-112G-AOC-0300
QSFP28-112G-AOC-0500
QSFP28-112G-DAC-0037
QSFP28-112G-DAC-0100
QSFP28-112G-DAC-CR-0100
QSFP28-112G-DAC-CR-0300
QSFP28-112G-ER4F-SM-LC
QSFP28-112G-LR4-SM-LC
QSFP28-112G-PSM4-SM-MPO
QSFP28-112G-SR4-MM-MPO
```

It needs to be the correct type, or it will not work.

If you as we run DAC cables, Juniper has FEC enabled on the DAC as default but Adva has it disable as default. You need to enable FEC on the client ports on the adva box to get a link up on the juniper router. You could also disable FEC on the juniper router but that seems a bit stupid..

```
admin@MCEN1-TR1# set interface 1/1/c3/et100 is-substates append mt
admin@MCEN1-TR1*# set interface 1/1/c1/et100 ety6-100g fec

Possible completions:
802-3bj          IEEE 802.3bj FEC
nofec            FEC not enabled

admin@MCEN1-TR1*# set interface 1/1/c3/et100 ety6-100g fec 802-3bj
admin@MCEN1-TR1*# delete interface 1/1/c3/et100 is-substates all
admin@MCEN1-TR1*# commit
```

Then it's time to configure the line side. Step 1 is to configure the line speed 100G or 200G


```

admin@MCEN1-TR1# set interface 1/1/n2/

Possible completions:
1/1/n2/ot100      Choose interface
1/1/n2/ot200      Choose interface

admin@MCEN1-TR1*# set interface 1/1/n2/ot100
admin@MCEN1-TR1*# commit

```

Step 2 wavelength and to do that you first need to set the interface in maintenance mode

```

admin@MCEN1-TR1# set interface 1/1/n1 is-substates append mt
admin@MCEN1-TR1*# set interface 1/1/n1 opt-phy frequency 193.80000
admin@MCEN1-TR1*# delete interface 1/1/n1 is-substates all
admin@MCEN1-TR1*# commit

```

Step 3 select the fec for the line, to work with the Juniper DWDM cordoba card you need to select the AC-100 mode.

```

admin@MCEN1-TR1# set interface 1/1/n1/ot100 is-substates append mt
admin@MCEN1-TR1*# set interface 1/1/n1/ot100 otu4 fec

Possible completions:
sd-hi-gain          High gain SD-FEC 25%, 5 iterations
sd-hi-gain-lit      High gain SD-FEC 25%, 1 iteration (low latency)
sd-low-gain         Low gain SD-FEC 15%, 5 iterations
sd-low-gain-ac100  Low gain SD-FEC 15%, AC-100 compatible

admin@MCEN1-TR1*# set interface 1/1/n1/ot100 otu4 fec sd-low-gain-ac100
admin@MCEN1-TR1*# delete interface 1/1/n1/ot100 is-substates all
admin@MCEN1-TR1*# commit

```

Step 4 create cross connects between the line port and the client port, for 200G line side you need to do two crossconnects.

```

admin@MCEN1-TR1# set connection 1/1/odu4/1 a-end 1/1/c1/et100/odu4 z-end 1/1/n1/ot100/odu4
admin@MCEN1-TR1*# commit

```

Juniper MX10003

The MX10003 has 12 100GE interfaces and 6 40GE interfaces on one linecard. If you are using all 6 40GE interfaces you can only use 9 of the 100GE interfaces.



The interfaces on the MX2010 in Malmö

```

bergroth@mcen1-r1-re0> show interfaces descriptions
Interface      Admin Link Description
et-0/0/0       up    up    SU-S002135 Link to lnd88-r1, mcen1-r1.lnd88-r1
et-1/0/0       up    up    SU-S002369 Link to lu-r2, mcen1-r1.lu-r2
et-2/0/0       up    up    SU-S002134 Link to ksdl-r1, mcen1-r1.ksdl-r1
et-4/0/0       up    up    Link to Nordunet NU-S000466, sunet-nordunet2
et-5/0/0       up    up    SU-S002368 Link to mah-r1, mcen1-r1.mah-r1
xe-5/2/0       up    up    SU-S003622 Link to ess-r2, mcen1-r1.ess-r2
xe-5/3/0       up    up    SU-S003526 Link to Netnod-IX Malmo, sunet-netnod-malmo
ge-6/0/0       up    up    SU-S003519 Link to SLU Alnarp, sunet-slualnarp
ge-6/0/1       up    up    SU-S003518 Link to Nordiska Genresurscenter (NORDGEN), sunet-nordgen
ge-6/0/2       up    up    SU-S003520 Moderna Museet (MODM), sunet-modm2
xe-6/2/1       up    up    SU-S003574, netnod-mcen1-r1.sthb-r2.l2c-2
xe-6/2/2       up    up    Link to Ricoh cloudprint, ricoh-print

```

With ESS moved to 100GE we needed two 1GE interfaces and four 10GE interfaces, three standard LR and one DWDM, that will use all the QSFP slot



With Junos 19.4R1 you can run a 1GE SFP in a QSFP slot with one of these adaptors

It's a waste of bandwidth but if you only need a few sfp slots it saves you the time and space to install a switch.

It seems though that the MX10003 can not run autonegotiation when running in 1GE mode on the QSFP.

We are using a DWDM plug towards Netnod but the Menara DWDM sfp+ plug did not accept the wavelenght settings for some odd reason. We need to get a fixed DWDM plug to get the netnod connection running again.

All interfaces came up and looked okay from our side.

2020-05-15 15:30

Left Malmö MCEN1 for Copenhagen.

Talked to my co-worker Carsten in Copenhagen and he agree to meet me on the Örestad site so I could install the ADVA DCI box.



Safety on the bridge

when I got past the high part of the bridge



I got to the Danish control



The nice control person wanted passport and papers. I didn't have any papers so I asked who should issues the paper, and it turned out that it's the one that you work for in Denmark that needs to issue the paper. I called Lars Lange Björn at Nordunet and asked if he could create an official paper allowing me to work in Denmark. 10 min later I had a PDF that let me pass the control.

The site in Denmark is really boring




Sunet already has a DWDM nod here so the only thing to do was to install the ADVA DCI box so we can connect with 2 x 100GE to Nordunets MX2010 router



We have ordered the AOC cables from ADVA to make the last connection but in the meanwhile are we running a Juniper DWDM in the Nordunet router to the ADVA DCI in Malmö.

The ADVA box fitted between the ADVA dwdm nod and the nordunet Corriant DCI box.



And some more power works at the back of the rack.

This PDU uses breakers on both positive and negative feed. In DK they use black and red as the colors for DC powers.



No configuration where don on the ADVA box, only checked that we had connection to the box from our DCN. The ADVA DCI box for the other DK site dk-bal was left at the Ore site for my Danish co-workers to install when the ACO cables arrives. That adva box has AC power so it's simpler to connect the power.

2020-05-15 18:10

Left the site to go back to Sweden. No checks in this direction, just pay the bridge fee and go.

Paulo from our noc called will I was in the car and it turned out that one of the Museums internet connection didn't work. The Museum asked if we could go to the site and restart their firewall.

Why not, I drove to the site and waited for two security guards to let me in.

2020-05-15 19:40

Moderna museets basement.



A 24 port linksys switch was used as a media converter. The fibre port flashed as one could expect. Connected my laptop to the switch and I could ping their FW but not our router.

Restarted the switch, swapped ports, no luck. Found a DB9 connector and logged into the switch.


```
Engine 0 REV 14 750-074276 CAN10003 JNP10003 [MX10003]
Engine 1 BUILTIN BUILTIN RE-5-1600x8
REV 35 750-067071 CANK8919 Control Board
REV 16 711-066896 CANM5201 Control Mezz Board
REV 35 750-067071 CANK8890 Control Board

Port Configuration
-----
Port Enable Auto Neg. Spd/Dplx Ctrl
cvr 1
cvr 2
cvr 3
cvr 4
cvr 5
g1 ENABLE On Auto Off
g2 ENABLE On Auto Off
Kcvr 0 g3 ENABLE On Auto Off
Kcvr 2 g4 ENABLE On Auto Off
Kcvr 4 g5 ENABLE On Auto Off
Kcvr 5 g6 ENABLE On Auto Off
g7 ENABLE On Auto Off
g8 ENABLE On Auto Off
g9 ENABLE On Auto Off
time=1.377 g10 ENABLE On Auto Off
time=1.289 g11 ENABLE On Auto Off
time=1.272 g12 ENABLE Off 1000E Off
time=1.307
time=1.439 Action-> Quit Edit Save
time=1.406
time=2.437 ArrowKey/TAB/BACK=Move SPACE=Toggle ENTER=Select ESC=Back
time=1.585
time=1.650
time=1.287
time=1.357
time=2.215
time=1.439
time=1.840
time=1.972
time=1.595
time=1.015
```

Turns out that the Linksys switch wasn't able to get the logical port up if the other side didn't talk auto-negotiation. Disabled auto-neg and we had connection,

2020-05-15 20:00

Back to MCEN1 to do the B-side of the power and clean up.



The delivery of the fibers could have been done better but, I guess it works.

2020-05-15 22:50

Left the site to go back to the Caravan.

2020-05-16 08:30

Heading towards Borås, stopped at Lund to fit and a missing coverplate for the DC PEM, the stupid the designed plastic thing.



and Växjö to correct an error that was found on the DCI boxes.



VXJ-TR2 has alarms on the frontplate but not in the admin interface, odd.


2020-05-16 16:45

Borås started the work of replacing the MX80 with and MX960

First moving fences to get into the site




and more junk



Then more work to move the rack post backwards. You need to flip those blue tabs to move the posts.



but the DWDM node is in the way.



Then it's just to push the router into the rack




with your knees



d pull the rest

and when it's half in, change side



an



Two boxes side by side



Bos-r2 was inserted between bos-r1 and jkg-r1.

```
bergroth@bos-r2-re0> show isis adjacency
Interface          System          L State      Hold (secs) SNPA
et-0/0/0.1        jkg-r1-re0     2 Up         24
et-11/0/0.1       bos-r1-re0     2 Up         21
```

2020-05-16 20:50

Left Borås to go home.

2020-05-16 22:20

Short stop in Jönköping to fix a breaker issue.

2020-05-17 00:30

Stopped at Linköping to get some sleep at the site, they had plenty of power to connect the heating in the Caravan.



2020-05-17 08:40

Back home

2020-05-18 10:00

Dropped off all the gear at the Stockholm Office. The DWDM cards will be used to upgrade other parts of the network. Do not know if that also going to be a caravan tour or not.



All in all 2600km by car.

<<Magnus