

Reflections on Redundancy

Ruben van den Brink, CTO

Maxx Cherevko, Network Architect



About AMS-IX



AMS-IX Amsterdam in numbers



Connected
Networks



Peak Internet
Traffic



Total traffic
volume in 2023



Total port
Capacity



Global presence



- AMS-IX locations
- Powered by AMS-IX

17 internet Exchanges

1150+ connected networks



Our Mission

**A better society
through better
Internet**



Be the exchange you want to see in this world



- Neutrality and Independence
- Online Safety
- Sustainability

We care about neutrality and independence

Neutrality is the cornerstone of our success. The internet is a network of networks, which are owned and operated by organizations and companies that have different, often conflicting, interests. They do have one shared concern: they need a network to carry their business. Each requires excellent connectivity, low latency, high availability and a rich ecosystem of networks to peer with.



We care about online safety

AMS-IX believes responsible use of the internet should not break any laws. Naturally, the virtues of an open and neutral internet have serious, undesirable side effects. As the impact of the online parallel universe continues to grow, cybercrime is threatening people as well as organizations and, ultimately, entire states. These threats are a result of the technical design of internet protocols, which date back decades. AMS-IX supports technical innovations that aim to resolve these fears.

As a neutral and independent operator of interconnection platforms, we explicitly do not police internet traffic and refrain from censoring. However, we do care about the information that flows through our systems. If the content itself is proven to break the law, we collaborate with legal authorities and take responsibility.

“ WE DO CARE ABOUT THE INFORMATION THAT FLOWS THROUGH OUR SYSTEMS ”



We care about sustainability

Digitalization was initially considered part of the solution in the quest to build a more sustainable society. With every major step in port speeds, advancements in network technology have reduced the energy needed to transport each byte of information, from Mbps to Gbps. At the same time, the connectivity sector has become a serious consumer of energy, as witnessed by the ever-increasing power consumption of the data center sector. Data centers are essential to supporting the digital services we all rely on and depend on in our daily lives. AMS-IX believes that a responsible internet is a more sustainable internet.

It is our ambition to reduce the impact of the connectivity sector across the energy chain.



Innovation is key

For Internet Exchanges, we can identify two directions for innovation

- **Horizontal**
Broaden your service portfolio with existing network services, like cloud access, transport, security (anti DDoS): might piss off customers
- **Vertical**
develop new services that leverage the trusted position of an internet exchange, like time services, data exchange, quantum exchange: we offer neutrality and independence




Innovation can be unexpected

Next, Maxx will tell us the AMS-IX story of photonic cross connects and the importance of redundancy in our platform.

It is about the historical pursuit of service availability, about a company that went out of business, about mirrors, and a raspberry pi.

In the unlikely event that you don't care about any of these things, please remember this one takeaway:

Don't underestimate the power of Friday afternoon projects...



The background is a solid orange color. It features several abstract geometric elements: a large, faint circle on the left side, a large, faint circle on the right side, and a faint line connecting the top of the right circle to the top edge of the frame. The text is centered on the left side of the image.

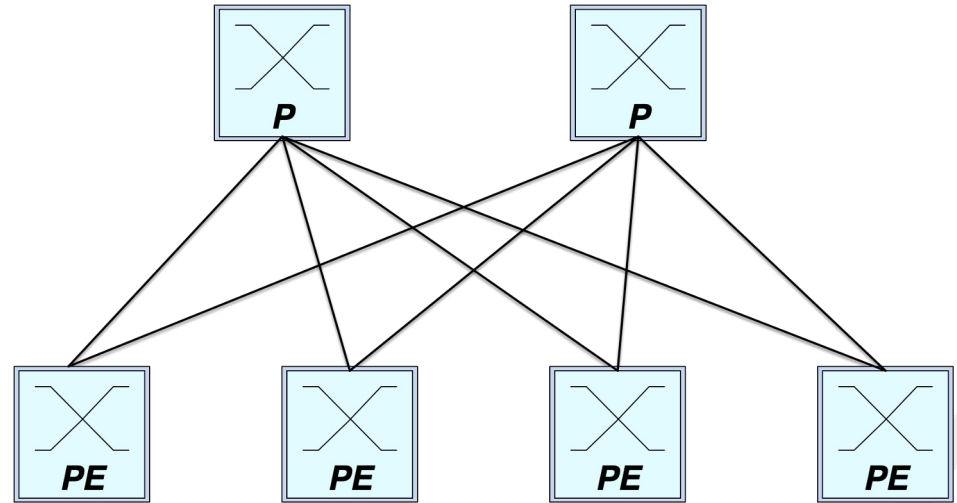
Internet Exchange Designed for Redundancy

Topology after ring > 2003

double star with VSRP > 2006



Spine-Leaf single PE



Increased scalability:

- Easy expansion/Flexibility

Enhanced redundancy:

- Multiple path
- Fast failover (ECMP, secondary LSP)

Other benefits:

- Cost-effective (pizzas vs chassis)
- Simpler management/troubleshooting
- High-port spines (should be dense)
- Cabling infrastructure (all PE to all P)
- **Potential vendor lock-in**

Issues with single PE design

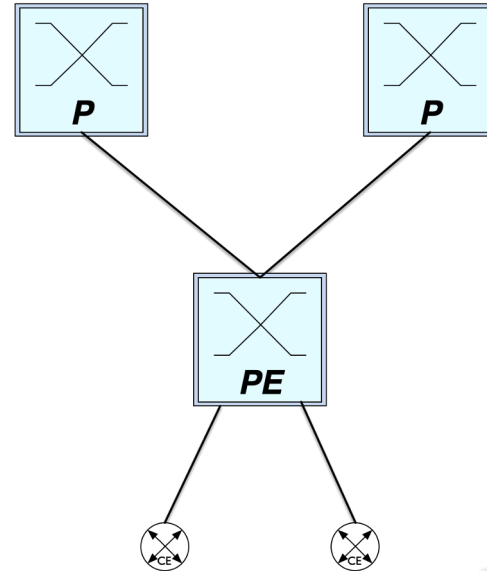


Different operational incidents (time of resolution):

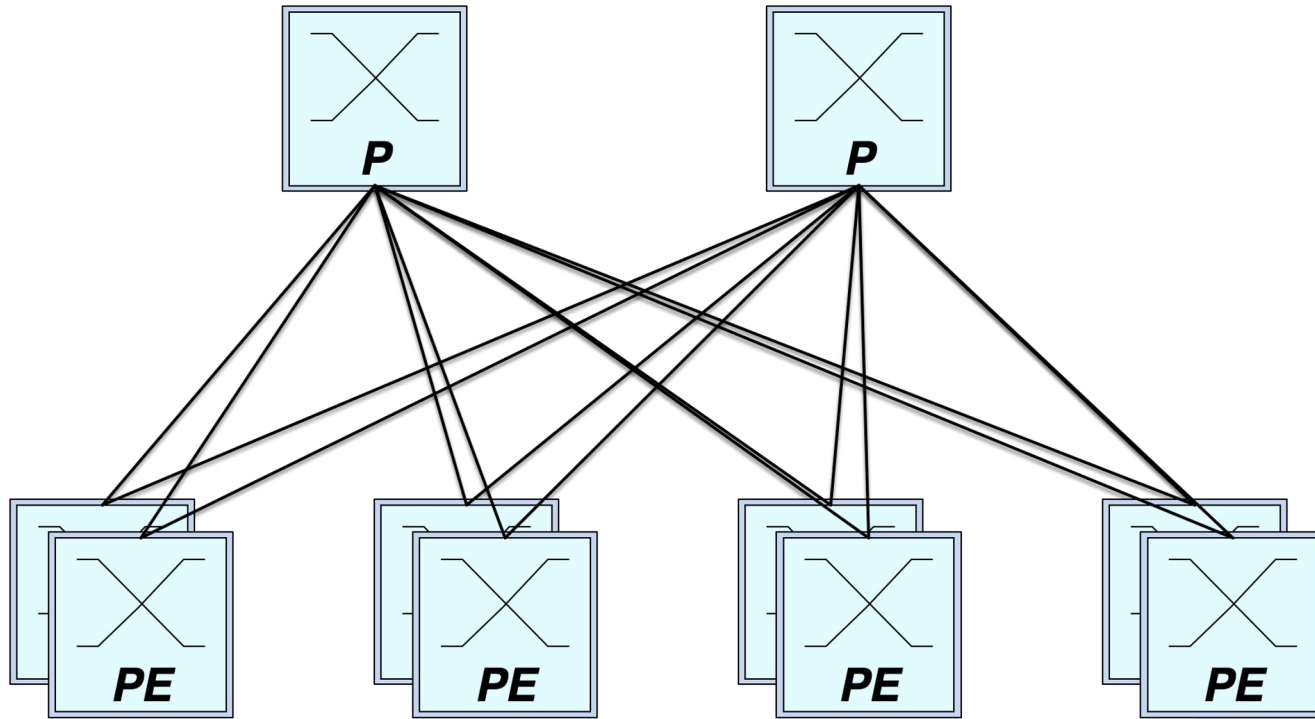
- CE facing optics died
- Line-card on PE HW issue
- Line-card/MM unexpected reboot (SW issue)
- Both backbone fiber-cut(unlikely but possible)
- MM CPU is overloaded with something!?!?
- Whole chassis "just died"
- and many others....

How to avoid long
downtime?

Spine-Leaf single PE



Double PE topology

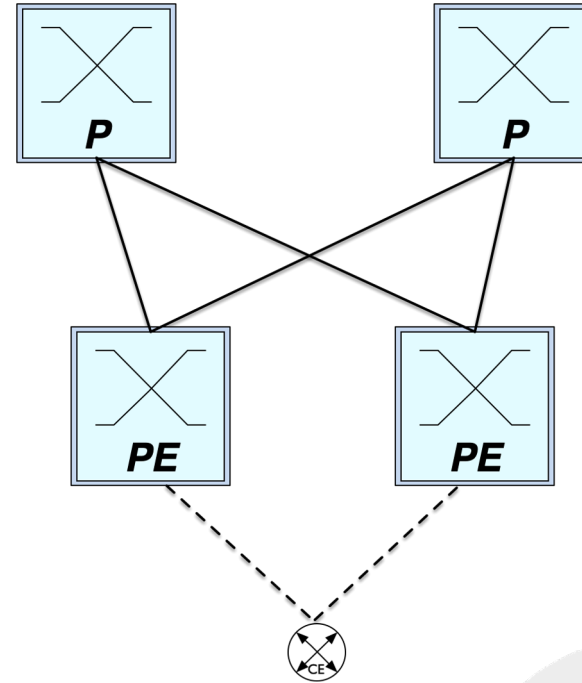


Problem solved?



Different operational incidents:


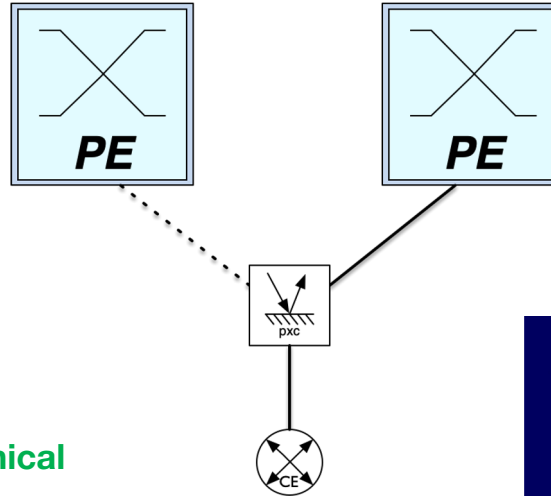
- CE facing optics died
- Line-card on PE HW issue
- Line-card/MM unexpected reboot (SW issue)
- Both backbone fiber-cut
- MM CPU is overloaded with something!?!?
- Whole chassis "just died"
- and many others....



Active-passive L1 switch setup



- L1 Active-passive setup
- Photonic X-Connect (PXC)
- reliable devices
- able to map any-to-any
- measure light level
- detects Tx/Rx issues
- fast swap (<25ms)
- 3D-MEMS (Microelectromechanical systems)
- BGP timers 90s/30s (180/60)



Autoswap and operational aspects

- LSP state monitoring
- 10x last year we had autoswaps (some due to HW/SW issues and some due to interop issues with multivendor network)
- Even L1 support line can solve 90% of issues with customers
- Field engineers have a lot of time to fix issues like optics, LC, BB etc.
- Replacing of unit is not service affecting
- Upgrade of the unit/reboot neither

```
# unexpected      reboot of LC1 on stub-eq3-344
2023-02-06T15:33:38+01:00 stub-eq3-344

# unexpected reboot of chassis
2023-05-09T11:37:05+02:00 stub-nik-341

# 20230516 maintenance load-balance customers niche
2023-05-16T00:13:59+02:00 stub-ix9-241

#?????????
2023-07-19T09:50:42+02:00 stub-eq5-247

# unscheduled reboot for stub-nik-341
2023-10-01T08:56:49+02:00 stub-nik-341

#20231113 - Replacement of core-glo-205 with Juniper unit
2023-11-14T00:14:35+01:00 stub-nik-341/stub-ix9-241

#lACP issue
2023-11-22T19:06:25+01:00 stub-nik-341

#Unexpected reboot of MLX switches in AMS-IX network
2023-11-27T08:05:58+01:00 12x MLXe

#swaps before reboot due to memory leak
2024-01-10T03:48:35+01:00 1x switch

#Unscheduled swaps on multiple PEs at DR1, DG1, EQX1, EQX3, EQX7 OSPF
2024-01-09T00:58:32+01:00 5x switches
```

Glimmerglass has left the building



- Around 70 units in production @AMS-IX NL network
- Really stable devices but no one lasts forever
- Lack of 1U dense units for IXaaS
- Glimmerglass went out of business since 01-01-2019
- Replacement?



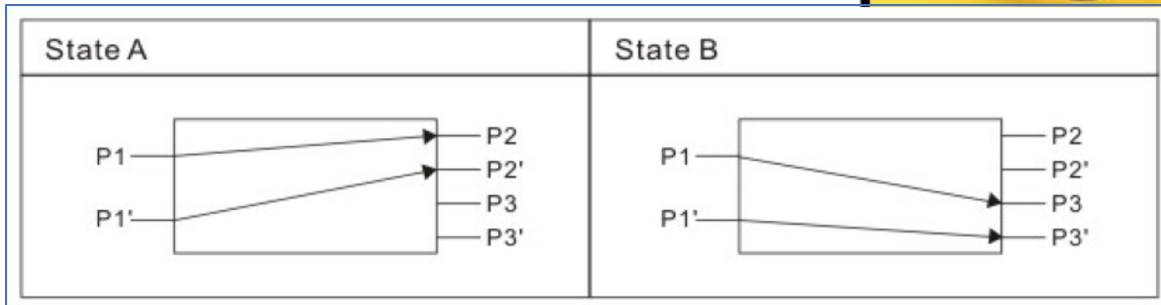

The image features a solid orange background. On the left side, there are faint, semi-transparent white circles and lines that suggest a network or molecular structure. On the right side, there is a larger, semi-transparent white circle with a smaller white circle inside it, also suggesting a network or molecular structure. The text is positioned on the left side of the image.

Investigation
Design
Development
MVP



D1x2 Optical Switch






- Latching
- 1310-1650nm
- insertion loss: 0.6dB
- fast (<8ms)
- current (55mA)
- cheap





Building Proof of Concept (PoC)





Patience
Time




Building Proof of Concept (PoC)



Is rackspace cheap?




MPO/MTP24 to LC



Minimum viable product (MVP)




- 2 rack units
- 8x optical modules
- 1 module x 12 relays
- latching relays
- 1x redundant PSU
- 1x management module
- ALL modules+PSU+MM are hot swappable
- one chassis can serve $8 \times 12 = 96$ customer connections
- MM is Raspberry Pi 4 :-)



AMS-IX MXC v1



GM-64
&
GM-160



AMS-IX MXC (boring)internals :-)

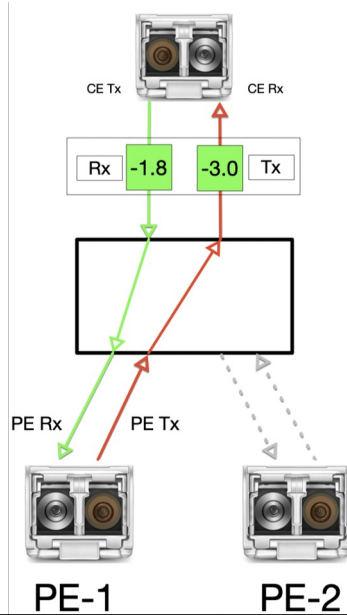


- Two types of chassis 2U for "bigger" locations 1U for small(interchangeable modules)
- Low insertion loss:
 - Pigtail < 0,3db * 2 (Senco)
 - Power reader 0,1db-0,2db (Santec)
 - Relay < 0,7db (Coreray)
 - Splicing 0,1db * 3
 - Total +/- 1.8db
- Latching relays
- NXP i.MX8 heart of MM
- Dual power supply
- Low power consumption(15-16w)
- Fast swaps(<8ms)



Web-UI & REST-API

- REST-API vs TL1
- Tx/Rx issues detecting



GM-64 TEST12

amsix

Control
Settings
Status
Update
Users
Logout

Module 1			Module 2			Module 3			Module 4		
	A	B		A	B		A	B		A	B
#	A/RX	B/TX	#	A/RX	B/TX	#	A/RX	B/TX	#	A/RX	B/TX
01	-57.2	-56.1	01	-57.5	-1.4	01	-57.3	-57.7	01	-56.6	-57.3
02	-57.1	-57.1	02	-57.4	-3.0	02	-57.1	-55.2	02	-56.8	-57.7
03	-56.9	-55.1	03	-57.5	-2.0	03	-57.1	-57.2	03	-57.7	-57.1
04	-56.9	-56.1	04	-57.4	-3.1	04	-56.9	-57.2	04	-57.3	-54.3
05	-56.1	-55.7	05	-57.1	-2.1	05	-56.6	-57.0	05	-57.0	-56.3
06	-57.2	-57.2	06	-57.4	-2.7	06	-57.7	-57.5	06	-56.7	-57.3
07	-56.4	-56.4	07	-56.7	-1.9	07	-53.8	-53.8	07	-56.4	-55.0
08	-56.7	-56.7	08	-9.2	-7.4	08	-56.3	-57.4	08	-57.7	-54.3
09	-5.0	-2.4	09	-54.0	-56.7	09	-55.9	-55.8	09	-57.3	-56.8
10	-3.0	-2.3	10	-56.8	-56.3	10	-57.3	-56.3	10	-55.5	-57.5
11	-5.6	-3.9	11	-53.3	-57.2	11	-56.5	-57.2	11	-56.6	-52.5
12	-3.2	-3.3	12	-56.7	-56.0	12	-57.0	-56.0	12	-56.9	-57.5
13	-2.4	-3.1	13	-56.6	-56.9	13	-56.7	-56.6	13	-57.5	-54.9
14	-3.0	-3.4	14	-57.3	-55.1	14	-56.0	-54.5	14	-55.6	-56.0
15	-3.3	-3.9	15	-57.1	-57.3	15	-56.8	-56.5	15	-54.0	-56.5
16	-3.6	-2.4	16	-52.6	-55.2	16	-55.8	-57.1	16	-57.2	-55.2

```
url='https://192.168.11.2:5000/api/mappings'  
query={"mappings": [{"id": "1/1", "mapping": "b"}, {"id": "1/2", "mapping": "b"}]}  
r = requests.patch(url, json=query, headers={'x-access-tokens':token}, verify=False)  
print(r.text)  
print(r.content)
```


MXC in (Virtual) Reality

MXC in Reality



- UAT testing in AMS-IX lab is ongoing
- Aim to place production order for the first batch before summer
- Initial deployment in 2024
- Several IX's have shown interest in this technology




If the MXC is sufficiently tested in the field, we consider making it commercially available.



MXC in Virtual Reality



MXC in Virtual Reality



Questions?

