



# AMS-IX

Ariën Vijn  
[arien.vijn@ams-ix.net](mailto:arien.vijn@ams-ix.net)

~ 1993

```

Stockholm      CERN
!              !
!512           !512
!              !
!      !-----!-----!
!      !   IDNX   !
!      !-!-----!-!
!              !   !
!              !448 !32
!              !IP  !CLNP
!              !   !
!-----!-----!-----!-----!   64
!      (ibr-router) ----- RedIRIS
!              !
!      Amsterdam. ----- Leuven
!      ebone.net      !   64      !-----! (UCD)
!              ----- PTT Telecom ----- IXI (ULB)
!-----!-----!              !-----! (YUNAC)
!
=====!=====!=====!===== ibr-lan
!              !
!-----!-----!      !-----!-----!
!Amsterdam1.!      !Amsterdam. !
!router.      !      !nl.eu.net !
!surfnnet.nl !      !              !
!-----!-----!      !-----!-----!
!              !
SURFnet      EUnet

```




~ 1993

```


Stockholm      CERN
!              !
!512           !512
!              !
!      !-----!-----!
!      !   IDNX   !
!      !-!-----!-!
!              !      !
!              !448  !32
!              !IP   !CLNP
!              !      !
!-----!-----!-----!-----! 64
!      (ibr-router) ----- RedIRIS
!              !
!      Amsterdam. ----- Leuven
!      ebone.net      ! 64          !-----! (UCD)
!              ----- PTT Telecom ----- IXI (ULB)
!-----!-----!          !-----! (YUNAC)
!
=====!=====!=====!===== ibr-lan
!              !
!-----!-----!      !-----!-----!
!Amsterdam1.!      !Amsterdam. !
!router.      !      !nl.eu.net !
!surfnet.nl !      !          !
!-----!-----!      !-----!-----!
!              !          !
SURFnet      EUnet

```



Vampire tap


Thicknet



Transeiver

Drop cable

AUI



~ 1993

```

Stockholm      CERN
!              !
!512           !512
!              !
!      !-----!-----!
!      !   IDNX   !
!      !-!-----!-!
!              !      !
!              !448  !32
!              !IP   !CLNP
!              !      !
!-----!-----!-----!-----! 64
!      (ibr-router) ----- RedIRIS
!              !
!      Amsterdam. ----- Leuven
!      ebone.net      ! 64          !-----!
!              ----- PTT Telecom ----- IXI (UCD)
!              !-----!
!-----!-----!          !-----!          (YUNAC)

```

```

!
=====!=====!=====!===== ibr-lan
!              !
!-----!-----!      !-----!-----!
!Amsterdam1.!      !Amsterdam. !
!router.      !      !nl.eu.net !
!surfnnet.nl !      !          !
!-----!-----!      !-----!-----!
!              !          !
SURFnet      EUnet

```



# Agenda

- Introduction.
- AMS-IX Organization.
- Technical Details.
- Operational aspects.
- Chicago setup.
  - Scalability.
- Questions / discussion.


# Agenda

- Introduction.
- AMS-IX Organization.
- Technical Details.
- Operational aspects.
- Chicago setup.
  - Scalability.
- Questions / discussion.



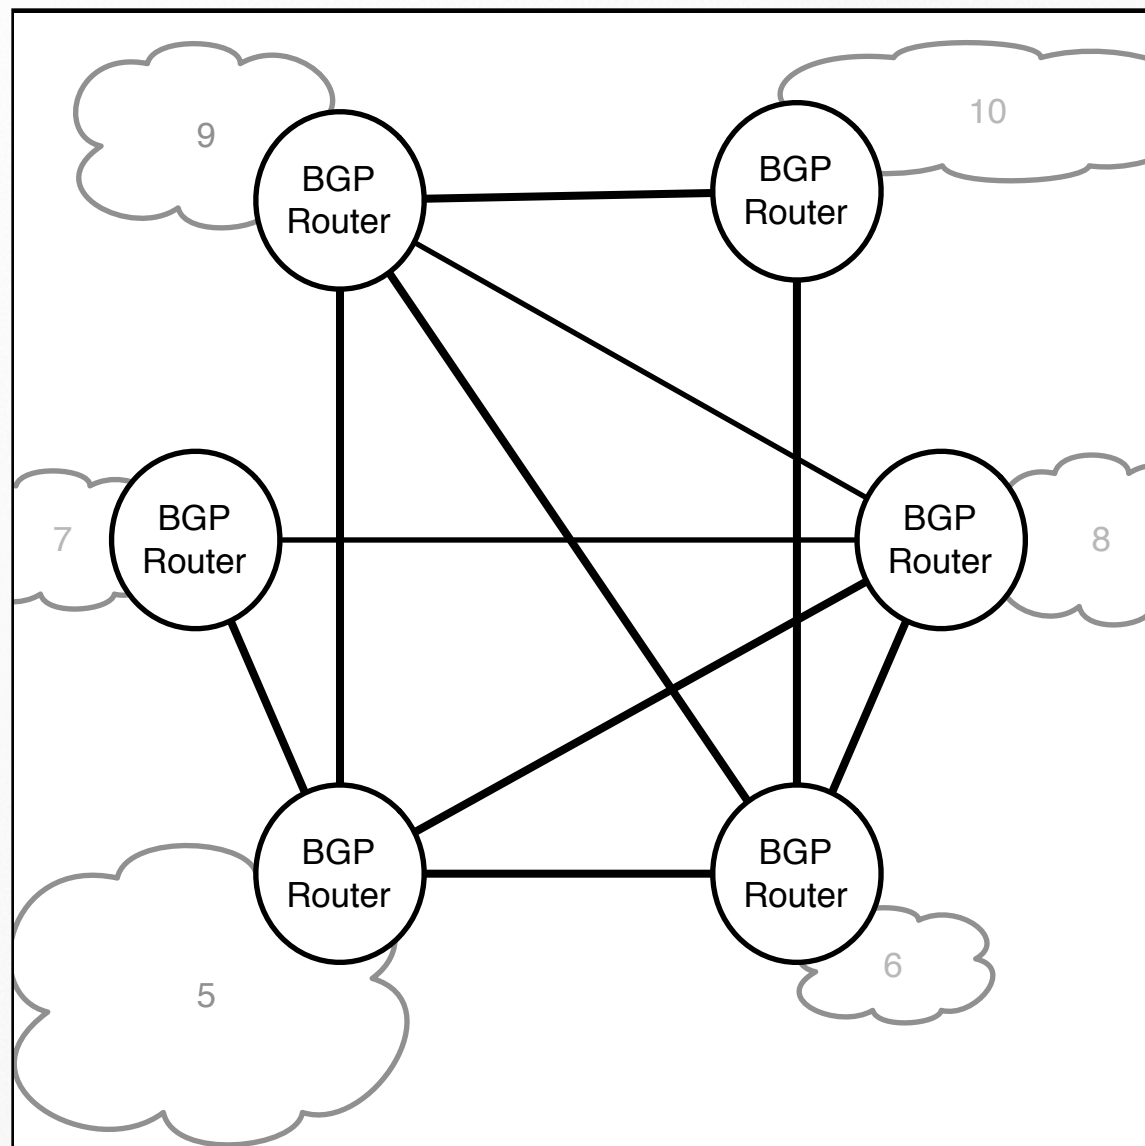
# Introduction

- How to peer?



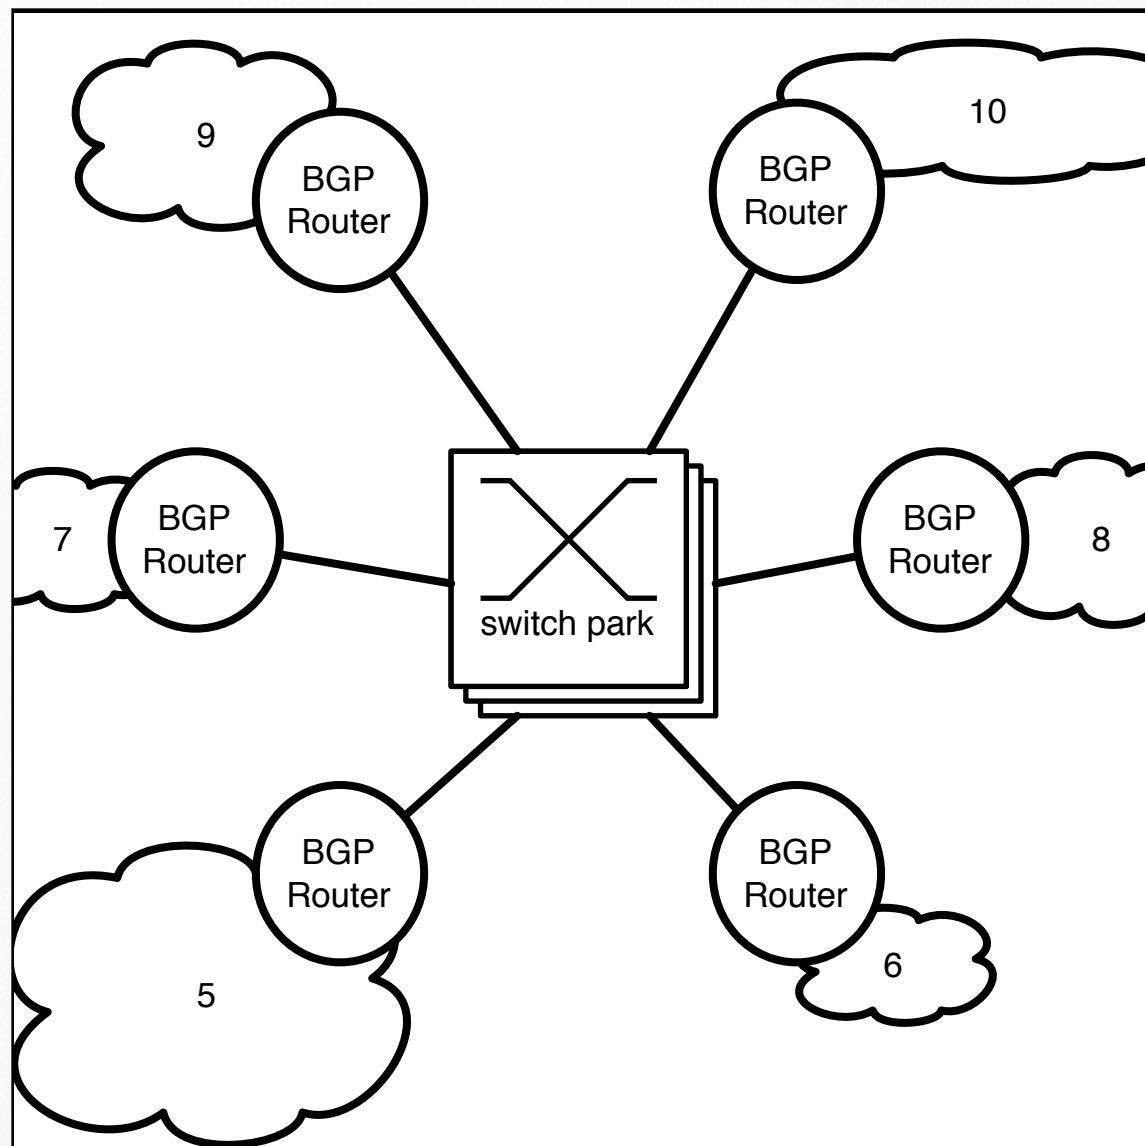


# Introduction



- How to peer?
- **Border Gateway Protocol (BGP).**
  - Interconnect BGP routers.
  - Max. n-1 interfaces per router.

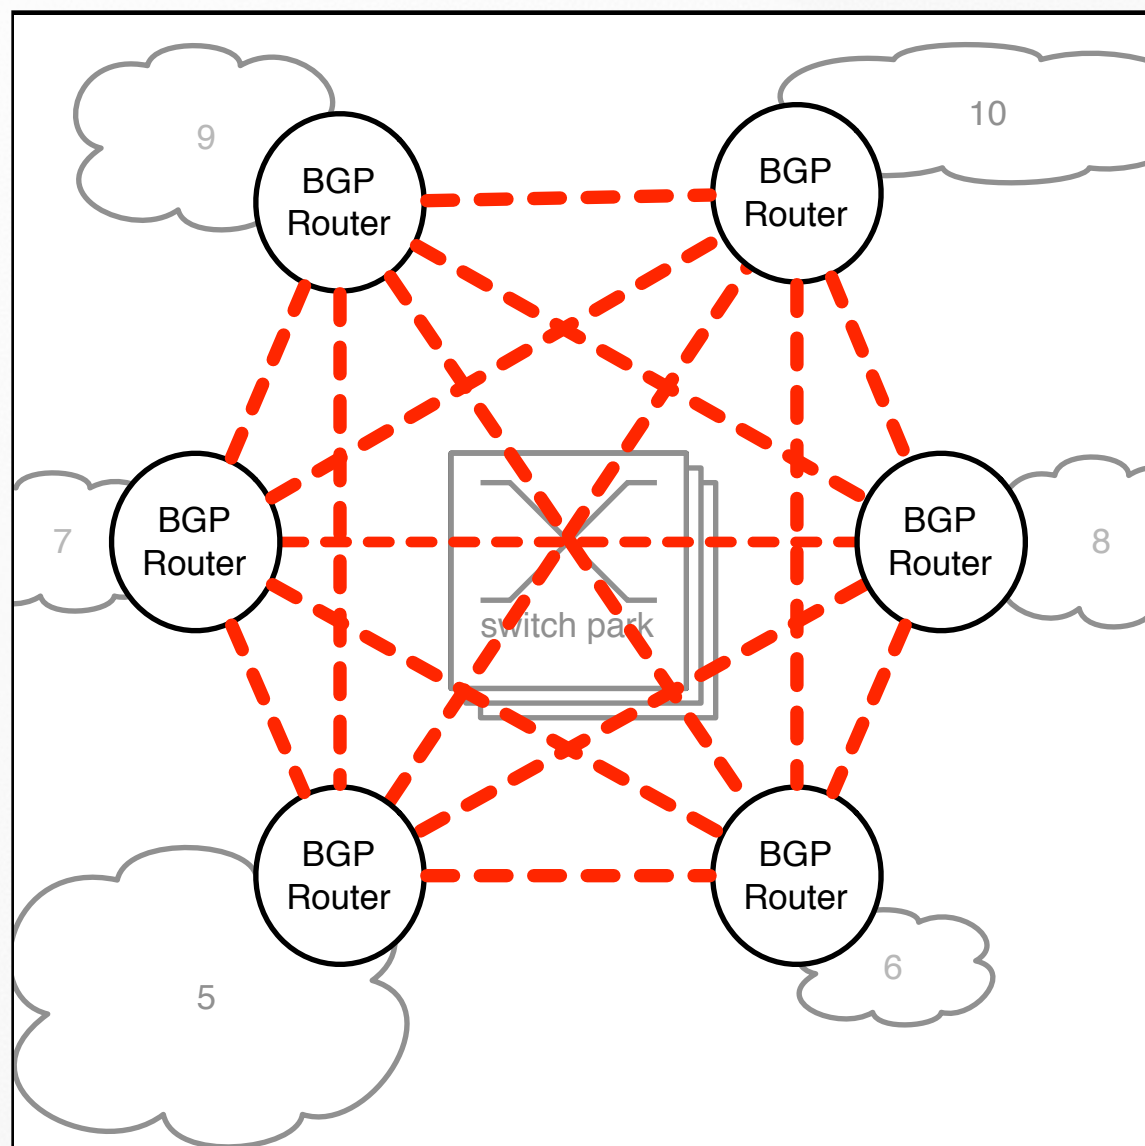
# Introduction



- How to peer?
- Border Gateway Protocol (BGP).
  - Interconnect BGP routers.
  - Max. n-1 interfaces per router.
- Using a shared infrastructure.
  - 1 connection per router.




# Introduction



- How to peer?
- **Border Gateway Protocol (BGP).**
  - Interconnect BGP routers.
  - Max. n-1 interfaces per router.
- Using a shared infrastructure.
  - 1 connection per router.
  - Everybody can peer with everybody else.

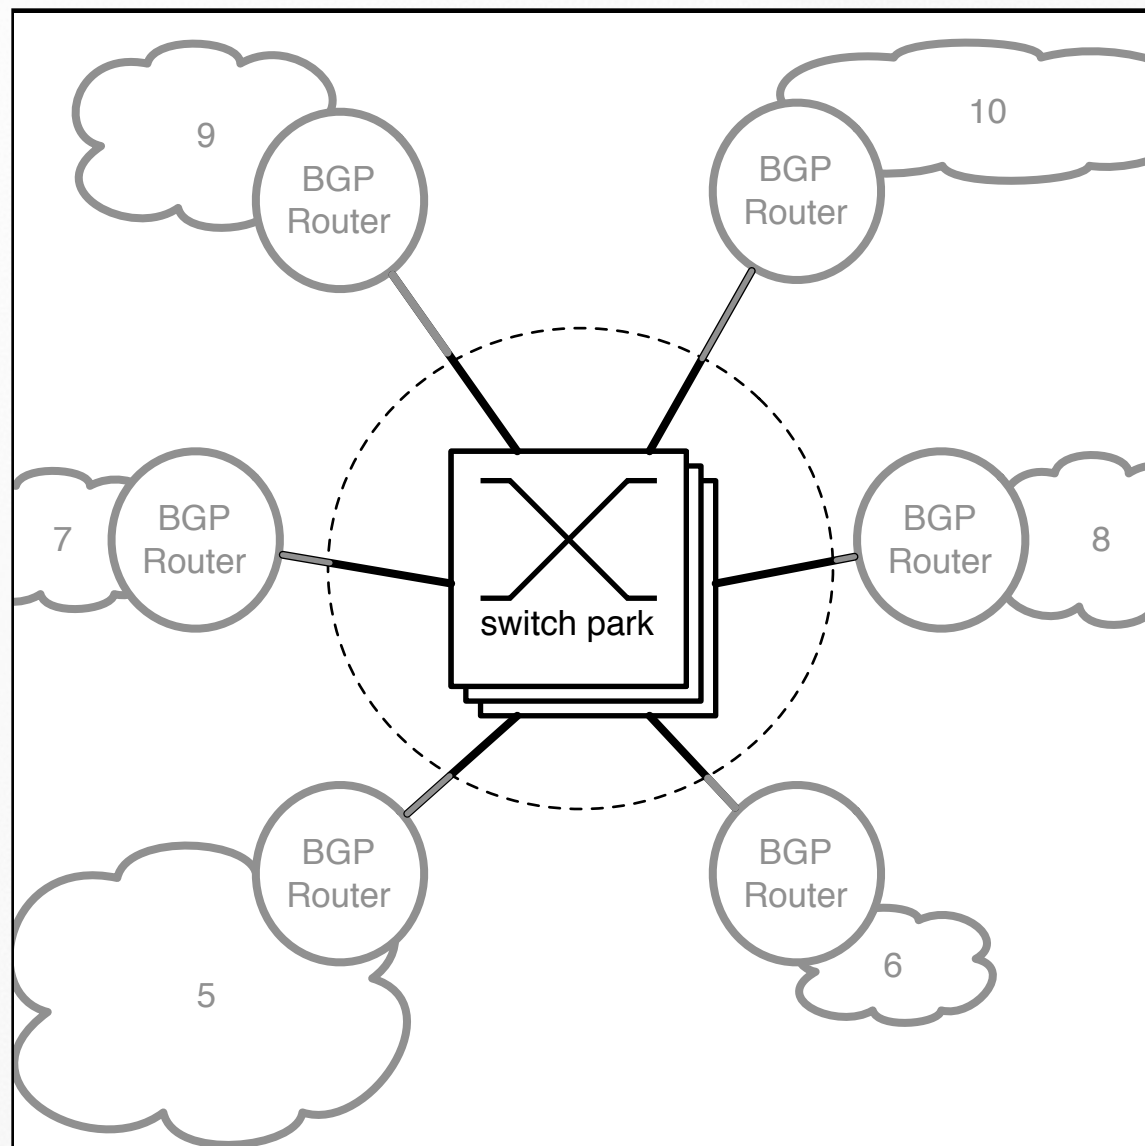
# Introduction



- How to peer?
- **Border Gateway Protocol (BGP).**
  - Interconnect BGP routers.
  - Max. n-1 interfaces per router.
- Using a shared infrastructure.
  - 1 connection per router.
  - Everybody can peer with everybody else.
- Not everybody peers with everybody else.
  - It is a bilateral deal.



# Introduction



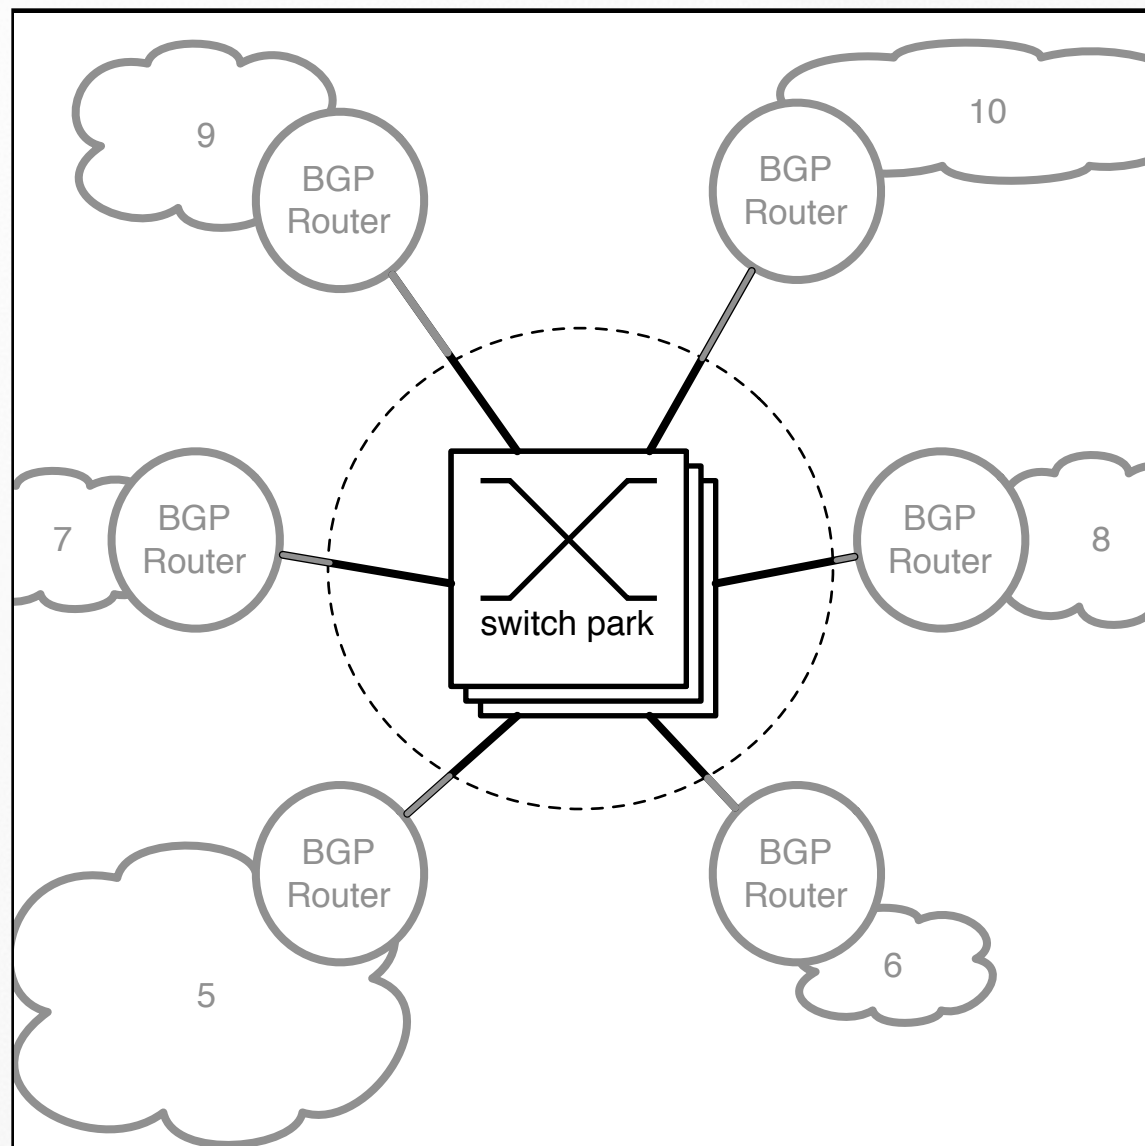
- An internet exchange.
- Facilitates exchange of traffic.
  - IPv4 and IPv6.
- Preferably neutral and open for new parties to join.
- **AMS-IX** is an internet exchange point in Amsterdam
  - **AMsterdam Internet eXchange**

# Agenda

- Introduction.
- **AMS-IX Organization.**
- Technical Details.
- Operational aspects.
- Chicago setup.
  - Scalability.
- Questions / discussion.



# AMS-IX organisation



- **AMS-IX is a professional, non-profit, independent and neutral Internet Exchange.**

# AMS-IX organisation

- AMS-IX is a professional, non-profit, independent and neutral Internet Exchange.
- Early 90's
  - Thick ethernet cable provided shared medium that became popular for peering.
  - No legal entity.

```
Stockholm      CERN
!              !
!512           !512
!              !
!              !-----!
!              ! IDNX  !
!              !-----!
!              !
!              !448  !32
!              !IP   !CLNP
!              !
!-----!-----!-----! 64
! (ibr-router) ----- RedIRIS
!              !
! Amsterdam. ----- Leuven
! ebone.net    ! 64      !-----! (UCD)
!              !-----! PTT Telecom ----- IXI (ULB)
!              !-----! (YUNAC)
!
=====!=====!=====!===== ibr-lan
!              !
!-----!-----!-----!-----!
!Amsterdam!.. !Amsterdam. !
!router.   !   !nl.eu.net !
!surfnet.nl !   !
!-----!-----!-----!-----!
!              !
! SURFnet    ! EUnet
```



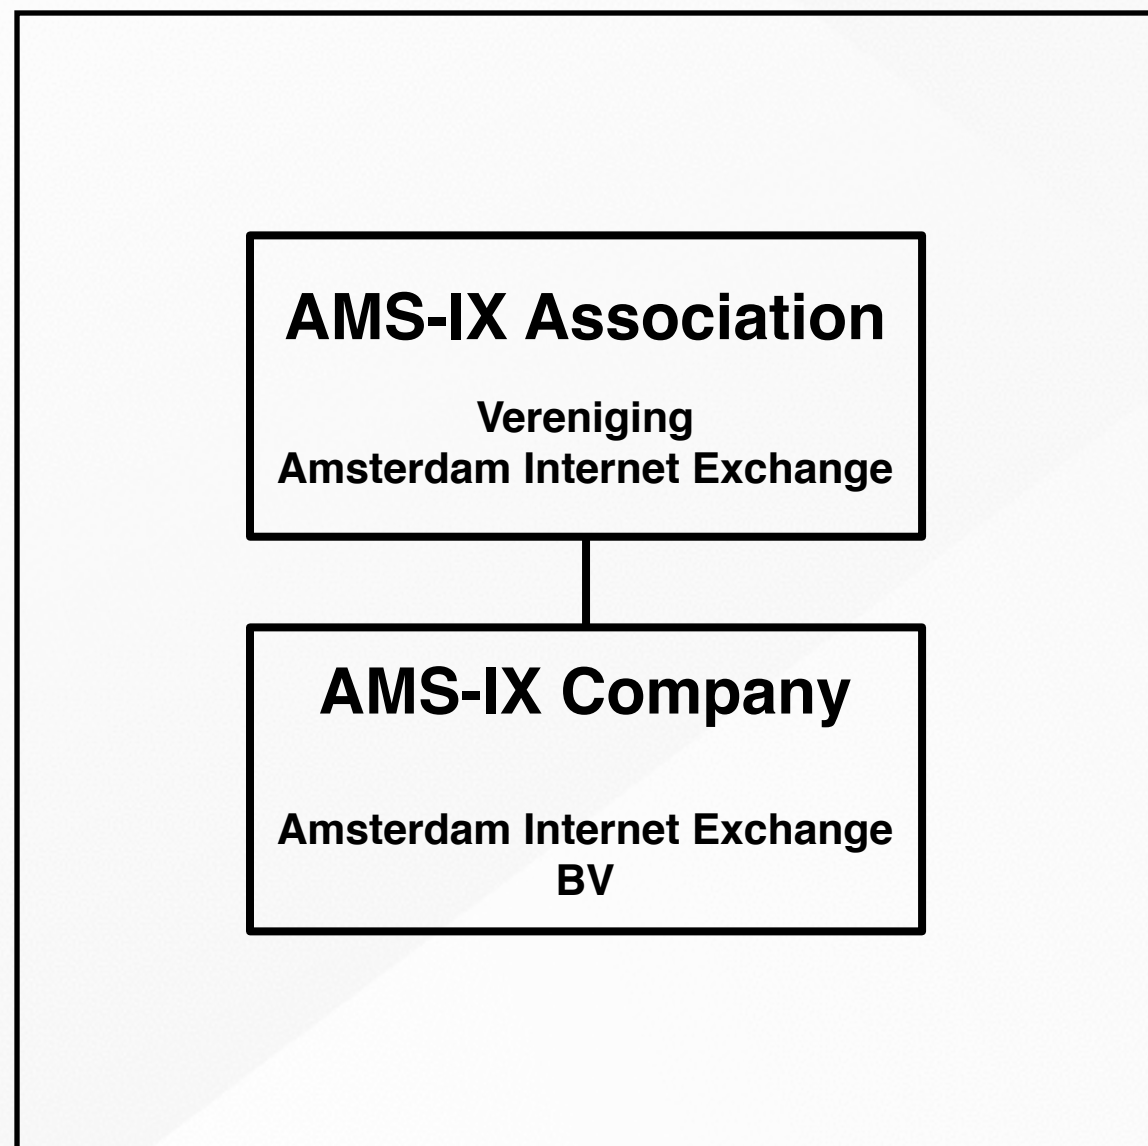
# AMS-IX organisation

A2000	IXE/PSI
AT&T EMEA	KPN Quest
BT	NLnet
Belnet	RIPE NCC
Demon	SURFnet
EUnet	Telecom Finland
EuroNet	UUnet/MCI
GTS Europe (Ebony)	Unisource
Global One	Wirehub
IBM Global Network	XS4ALL

- **AMS-IX is a professional, non-profit, independent and neutral Internet Exchange.**
- **Early 90's**
  - **Thick ethernet cable provided shared medium that became popular for peering.**
  - **No legal entity.**
- **Formalised 29 December 1997**
  - **Association in the Netherlands, operating under Dutch Law.**
  - **Over 675 Networks (October 2014).**



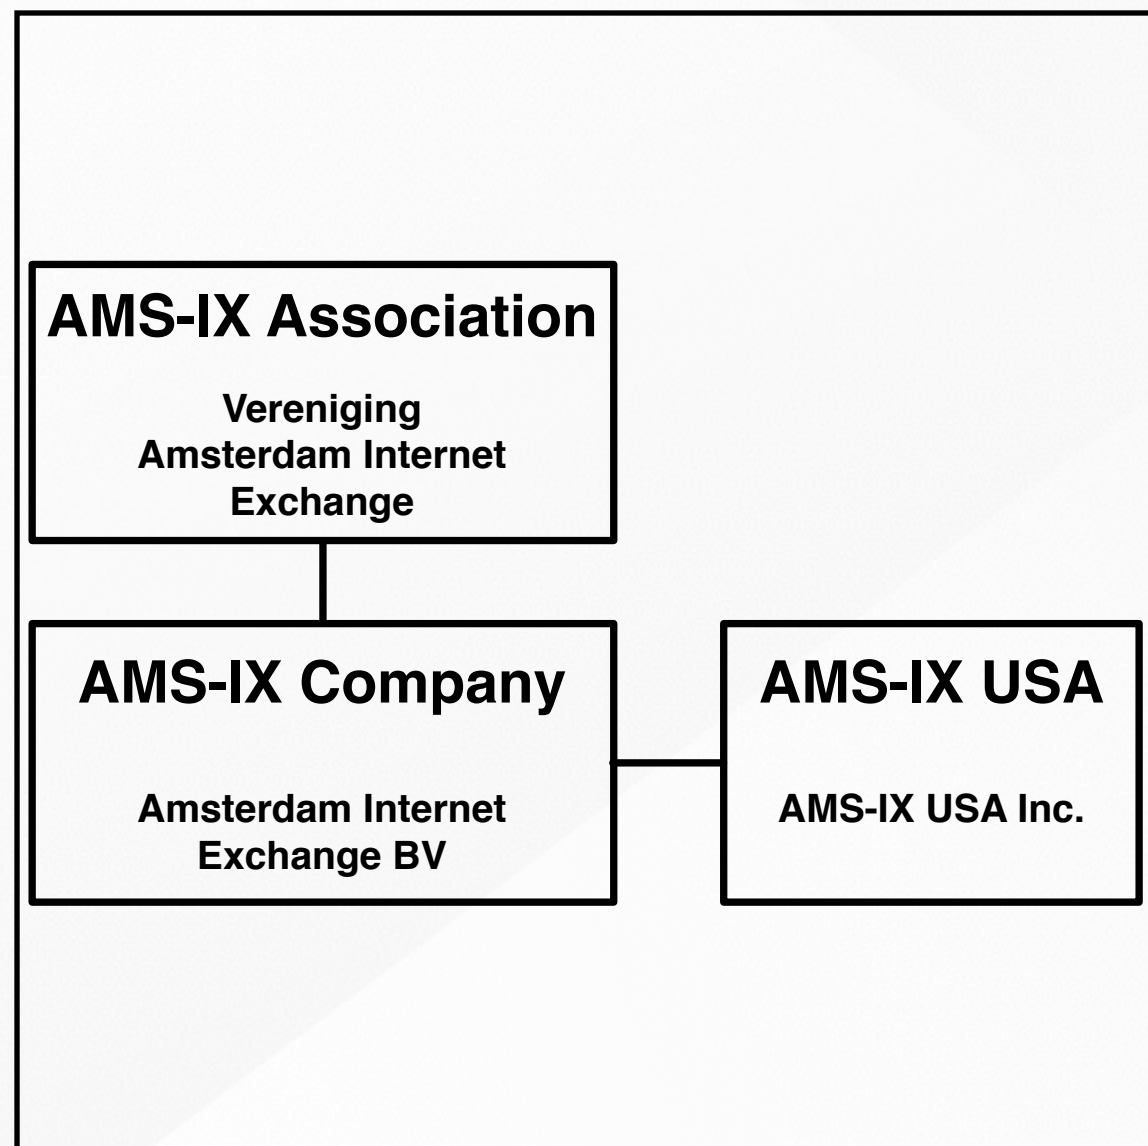
# AMS-IX organisation



- **Amsterdam Internet Exchange B.V.**
  - **Limited Liability Company.**
  - **The AMS-IX Association is the single shareholder.**
  - **Founded in 2000 as the administrative body.**
  - **Took over operations in December 2001.**
  - **AMS-IX NOC.**



# AMS-IX organisation



- **Amsterdam Internet Exchange B.V.**
  - **Limited Liability Company.**
  - **The AMS-IX Association is the single shareholder.**
- **Founded in 2000 as the administrative body.**
- **Took over operations in December 2001.**
  - **AMS-IX NOC.**
  - **AMS-IX USA Inc.**
    - **Founded in 2013.**

# AMS-IX organisation


AMS-IX bodies	AMS-IX association	AMS-IX B.V.
Meeting	General Assembly (ALV)	Shareholders Meeting
Board	Executive Board	Supervisory Board
Management		Management



# Agenda

- Introduction.
- AMS-IX Organization.
- **Technical Details.**
- Operational aspects.
- Chicago setup.
  - Scalability.
- Questions / discussion.

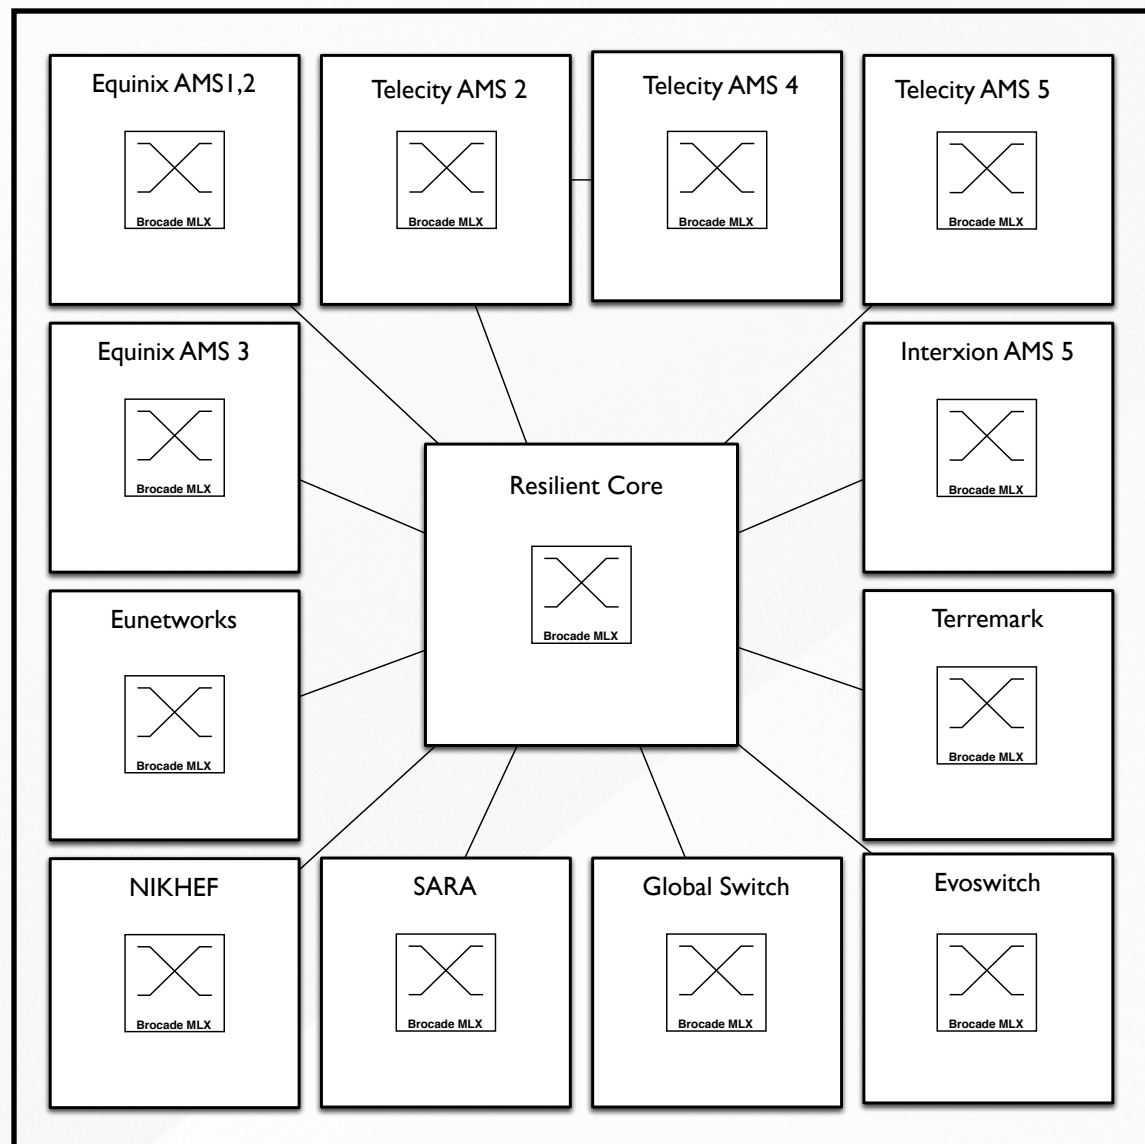
# Introduction



- **Modern AMS-IX infrastructure**



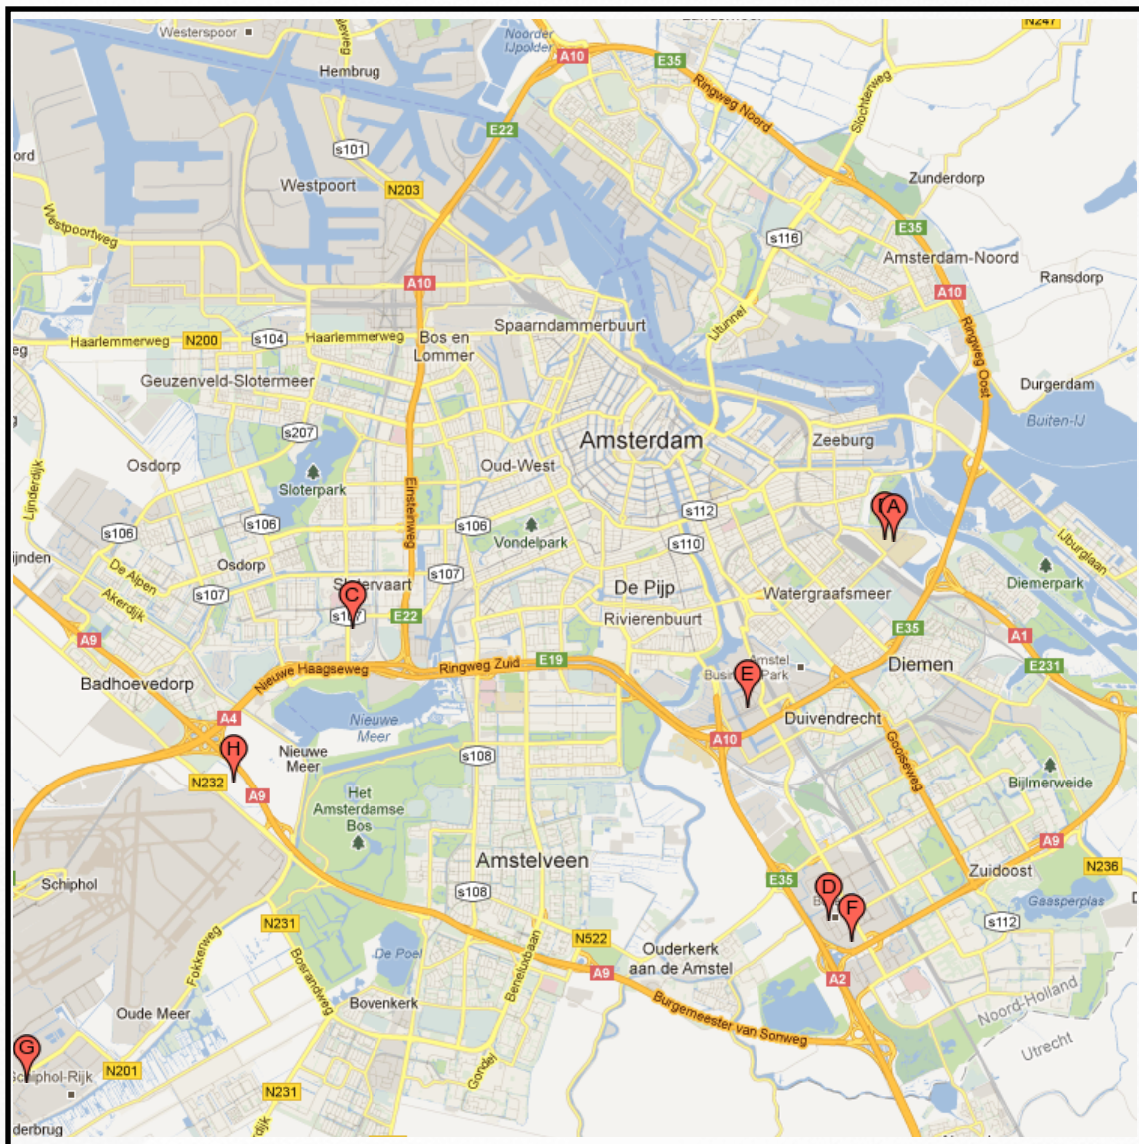
# Introduction



- **AMS-IX infrastructure**
- **12 locations**

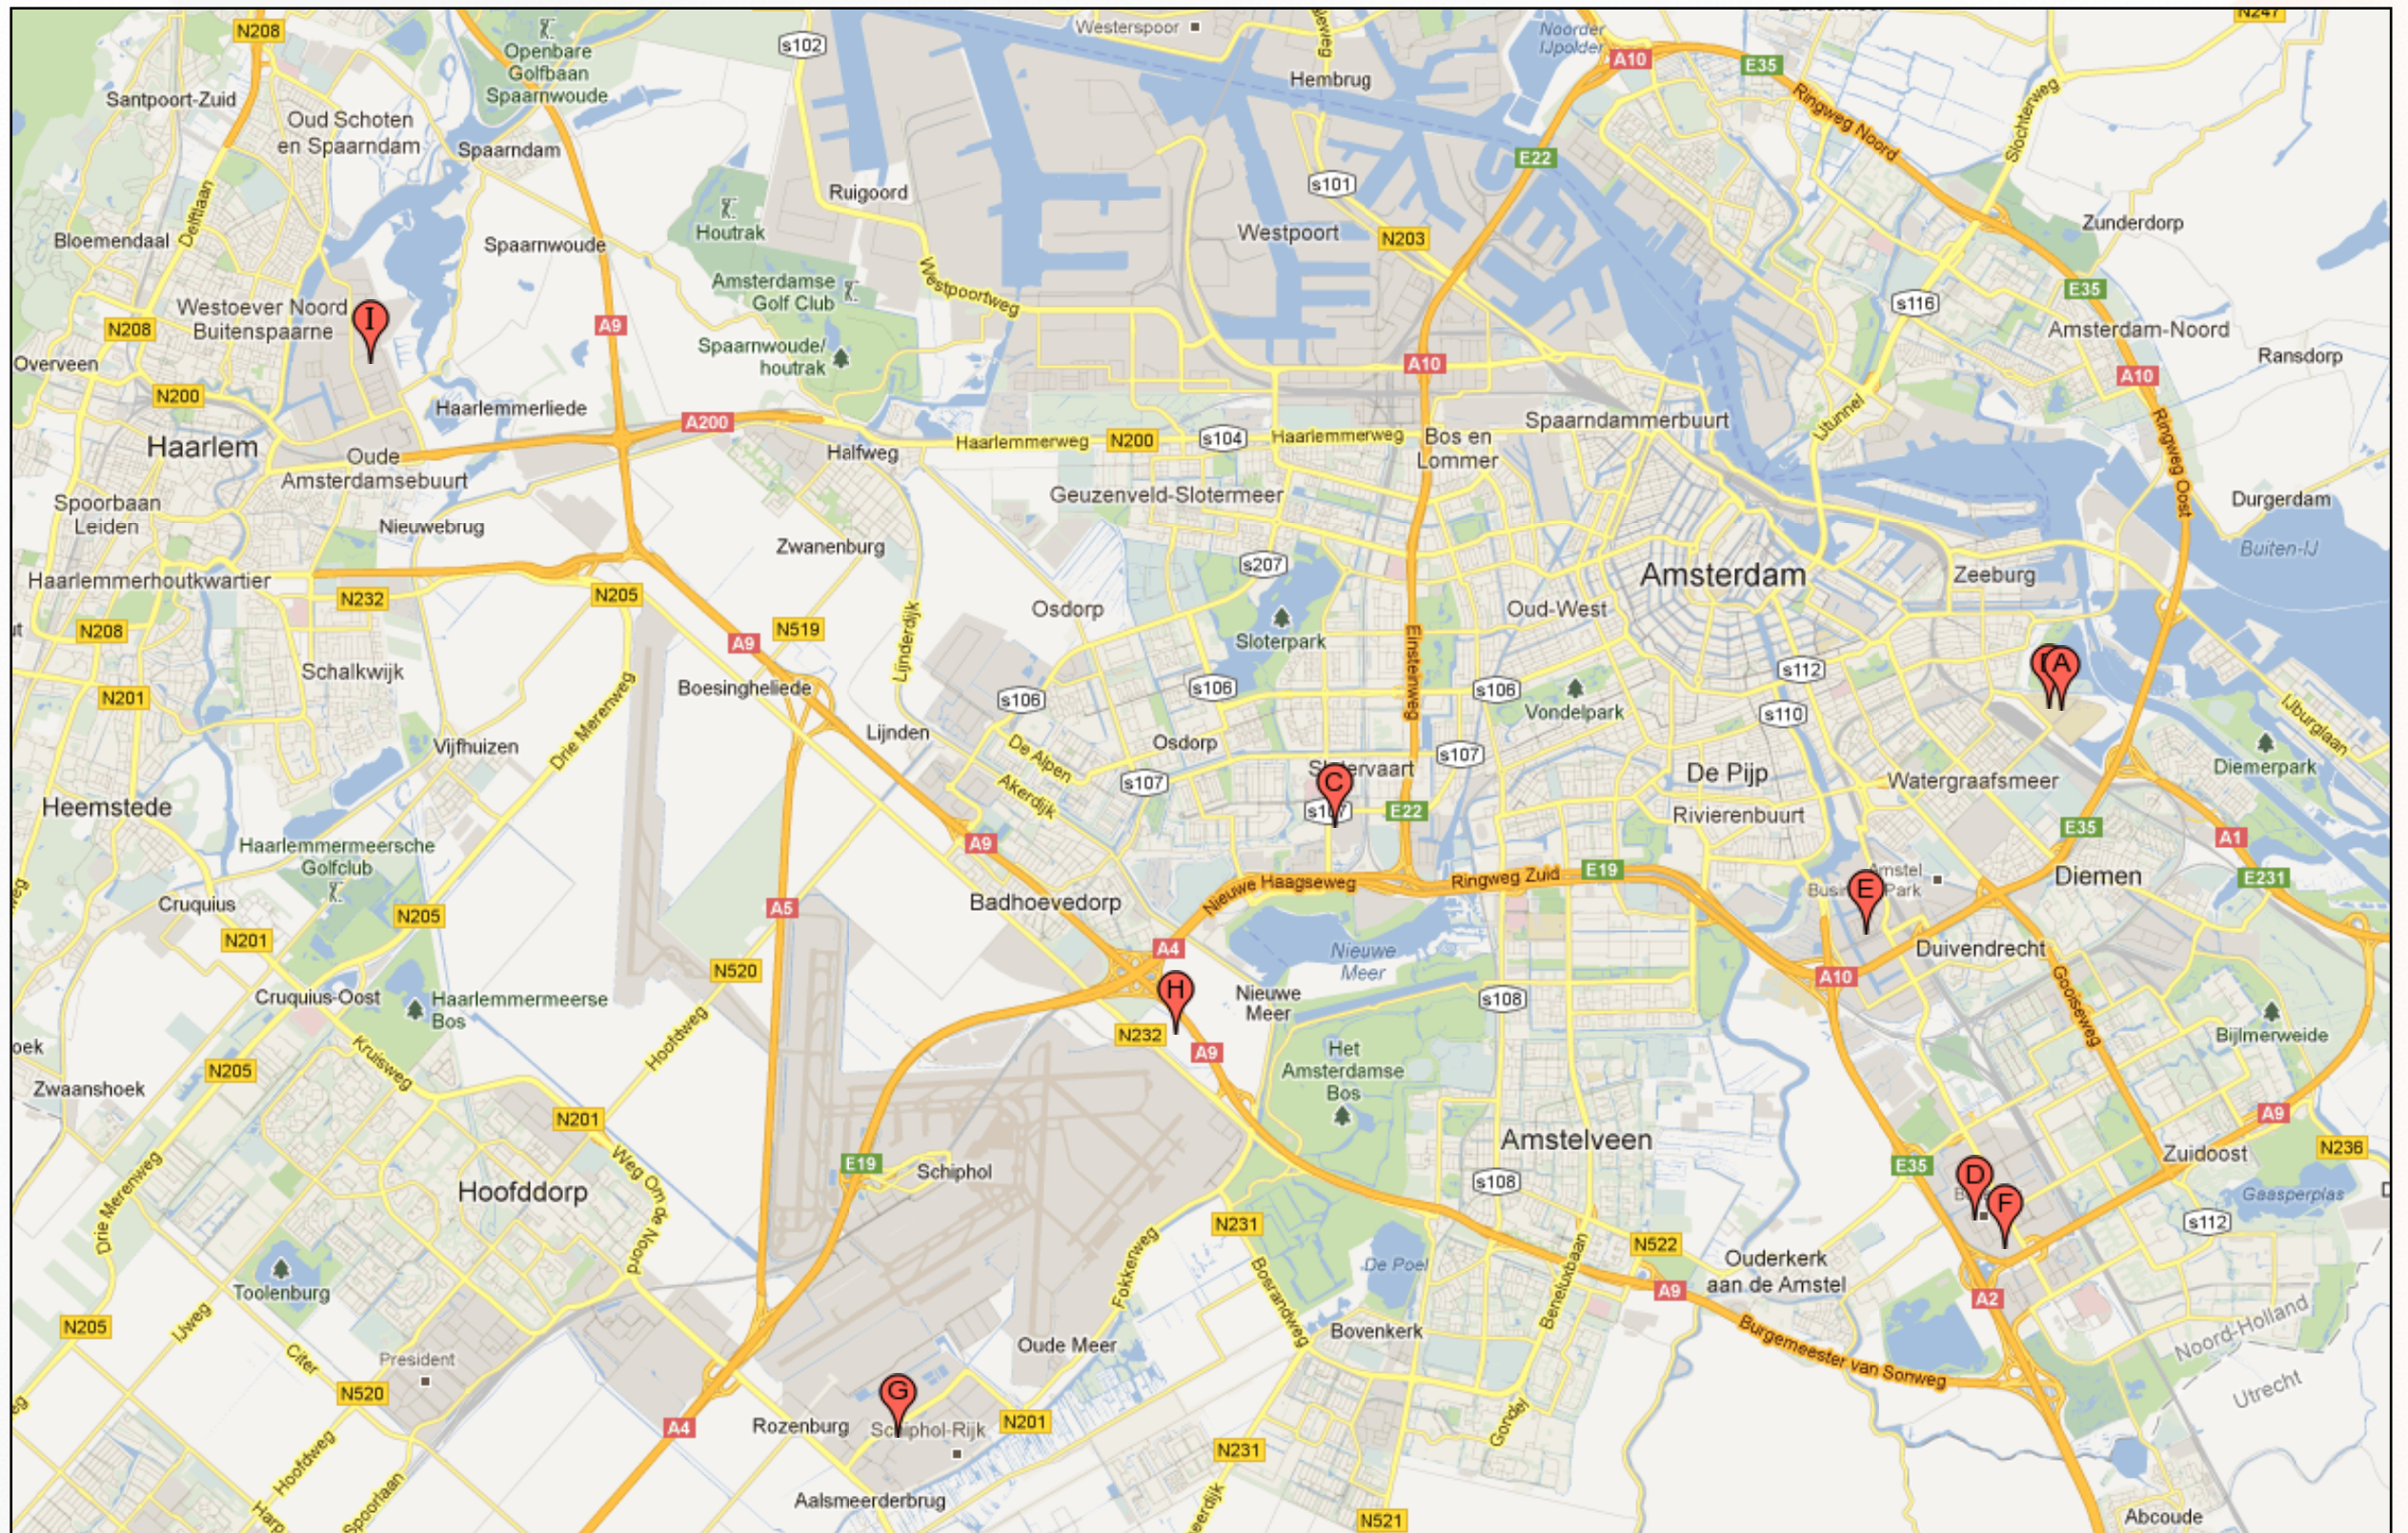


# Introduction



- **AMS-IX infrastructure**
- **12 locations**
- **7 geographical areas**







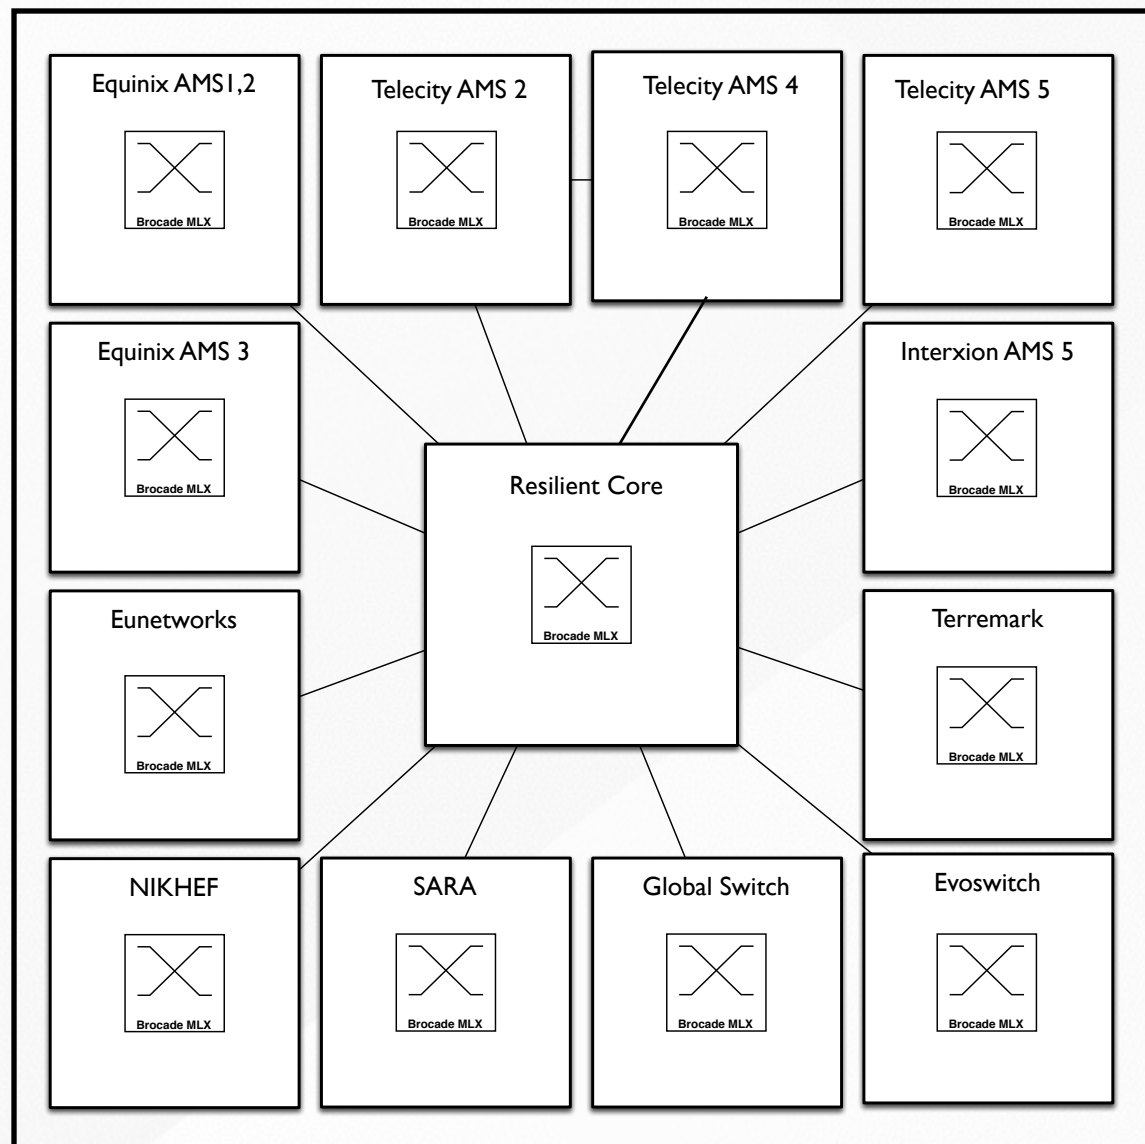
# Introduction



- **AMS-IX infrastructure**
  - 12 data centre
  - 7 geographical areas
  - 9 different organisations




# Introduction




- **AMS-IX infrastructure**
  - 12 data centre
  - 7 geographical areas
  - 9 different organisations
- 2 dark fibre suppliers
  - Eurofiber
  - euNetworks

# Technical overview



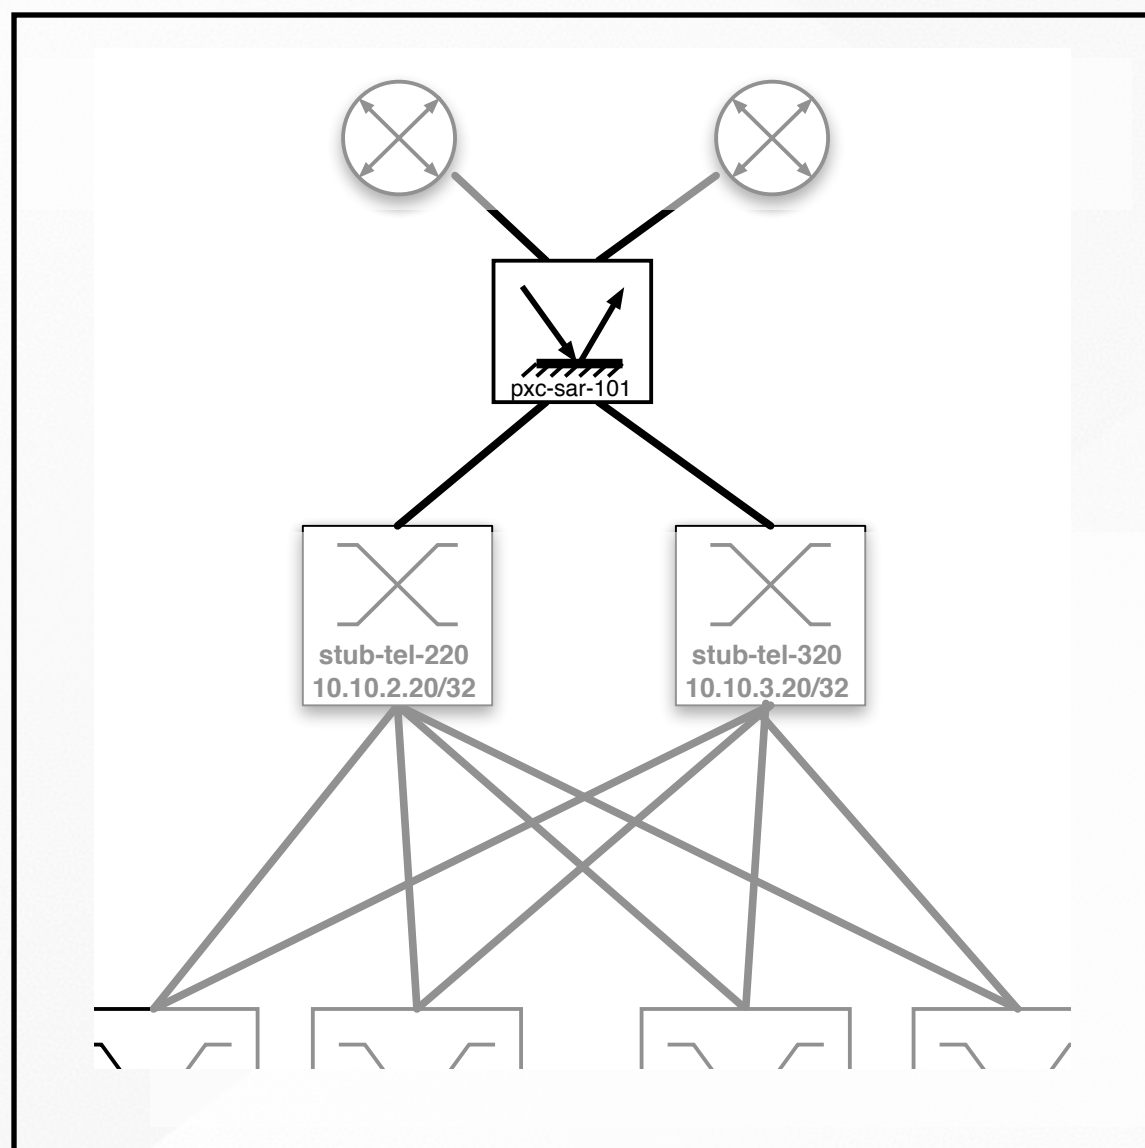


# Technical overview



- Border routers of member (CE)



# Technical overview





- Border routers of members
- Photonic Cross Connect (PXC)
  - Layer 1
  - Low energy consumption
  - Reliable



# Technical overview





# Technical overview







# Technical overview





# Technical overview






# Technical overview




# Technical overview



- Signal (light) comes in
- Photonic Cross Connect (PXC)
  - Layer I
  - Low energy consumption
  - Very Reliable




# Technical overview



- PE routers (stub switches)
  - Puts a label on the frame
  - From now on it is label switched
  - Distributes traffic over core switches (P routers).


# Technical overview



- PE routers (stub switches)
  - Puts a label on the frame
  - From now on it is label switched
  - Distributes traffic over core switches (P routers).




# Technical overview



- PE routers (stub switches)
  - Puts a label on the frame
  - From now on it is label switched
  - Distributes traffic over core switches (P routers).

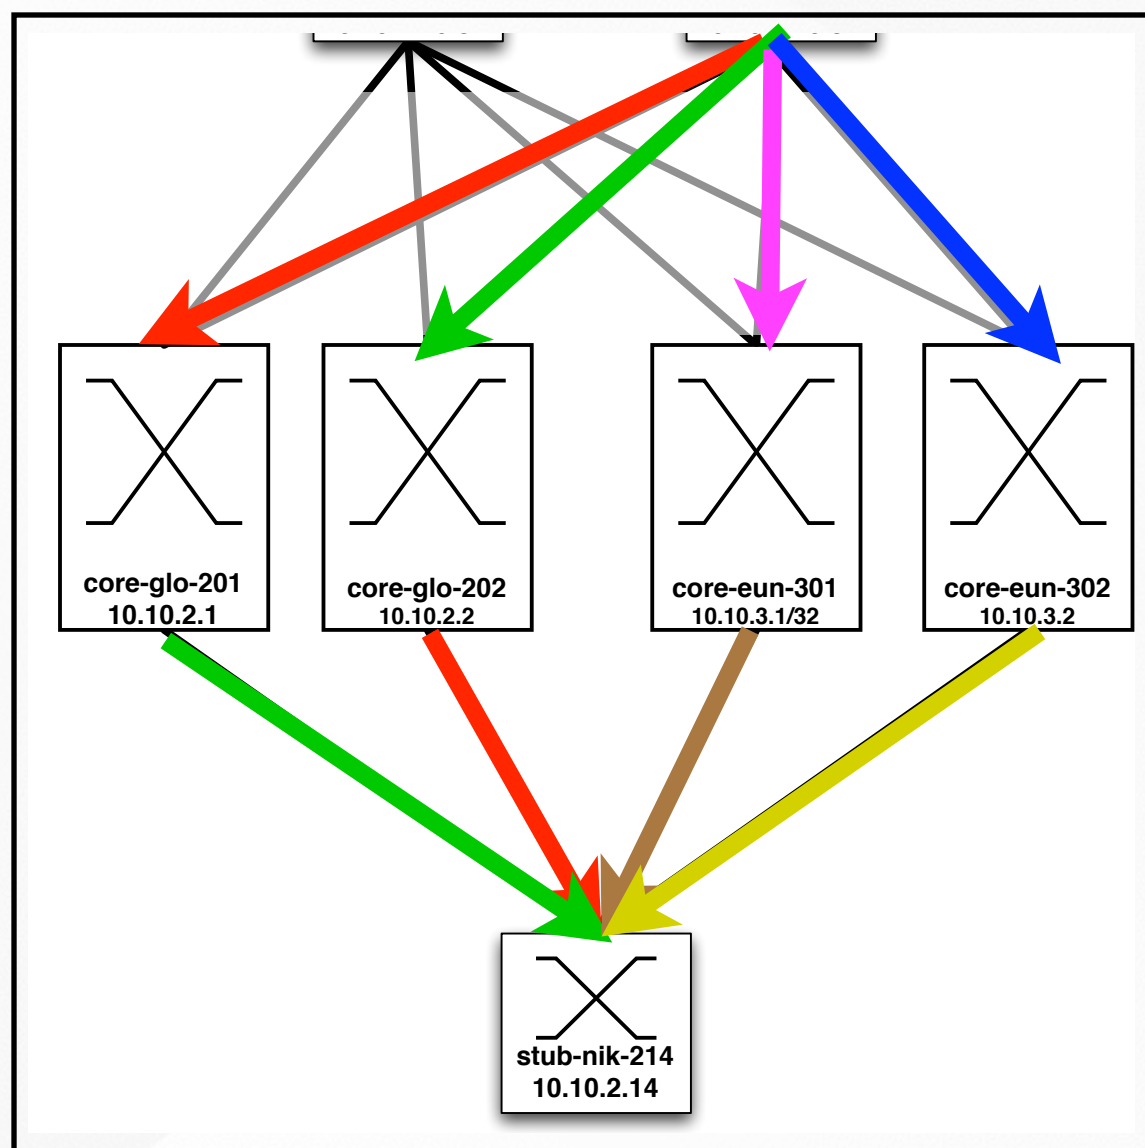
# Technical overview



- PE routers (stub switches)
  - Puts a label on the frame
  - From now on it is label switched
  - Distributes traffic over core switches (P routers).

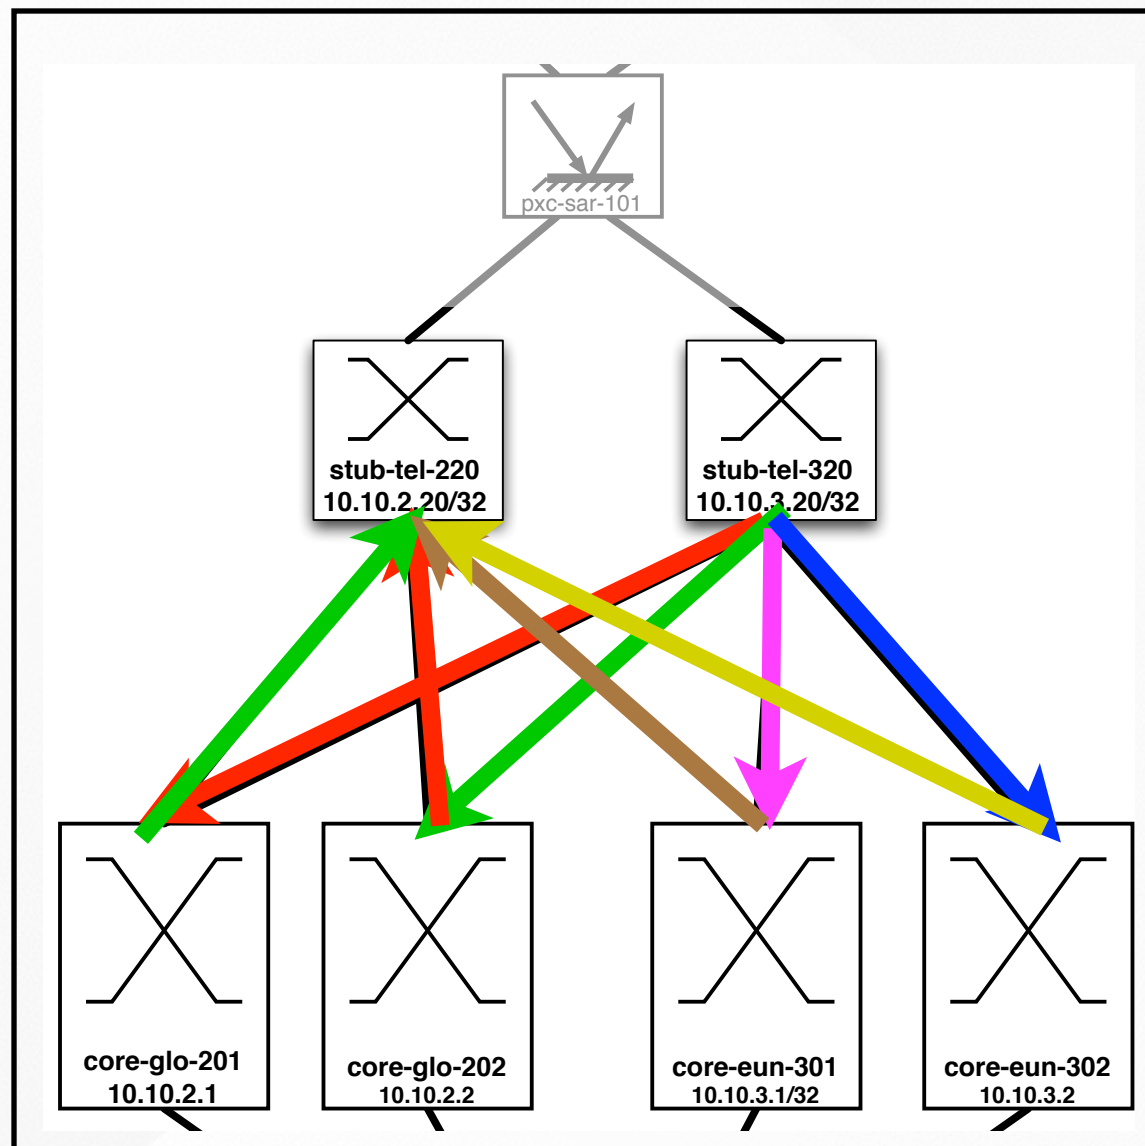


# Technical overview



- PE routers (stub switches)
  - Puts a label on the frame
  - From now on it is label switched
  - Distributes traffic over core switches (P routers).
- P-routers
  - Switch traffic through to the destination PE.


# Technical overview



- PE routers (stub switches)
  - Puts a label on the frame
  - From now on it is label switched
  - Distributes traffic over core switches (P routers).
- P-routers
  - Switch traffic through to the destination PE.




# Technical overview



- Full mesh IP network
- OSPF as routing protocol

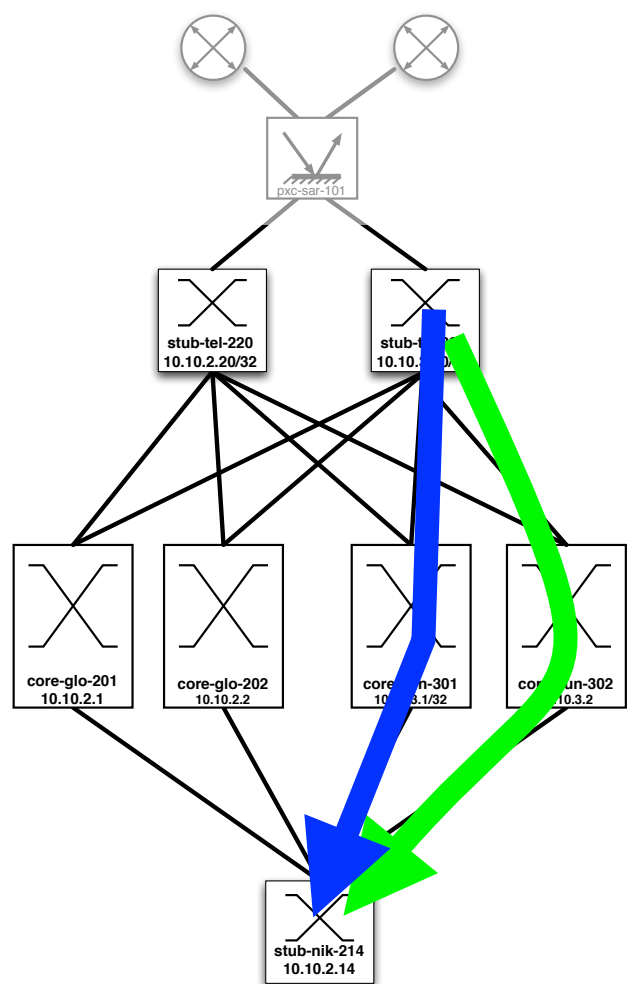
# Technical overview



- Full mesh IP network
- OSPF as routing protocol
- BFD for fast OSPF shutdown
- Replacing that by Ethernet OAM

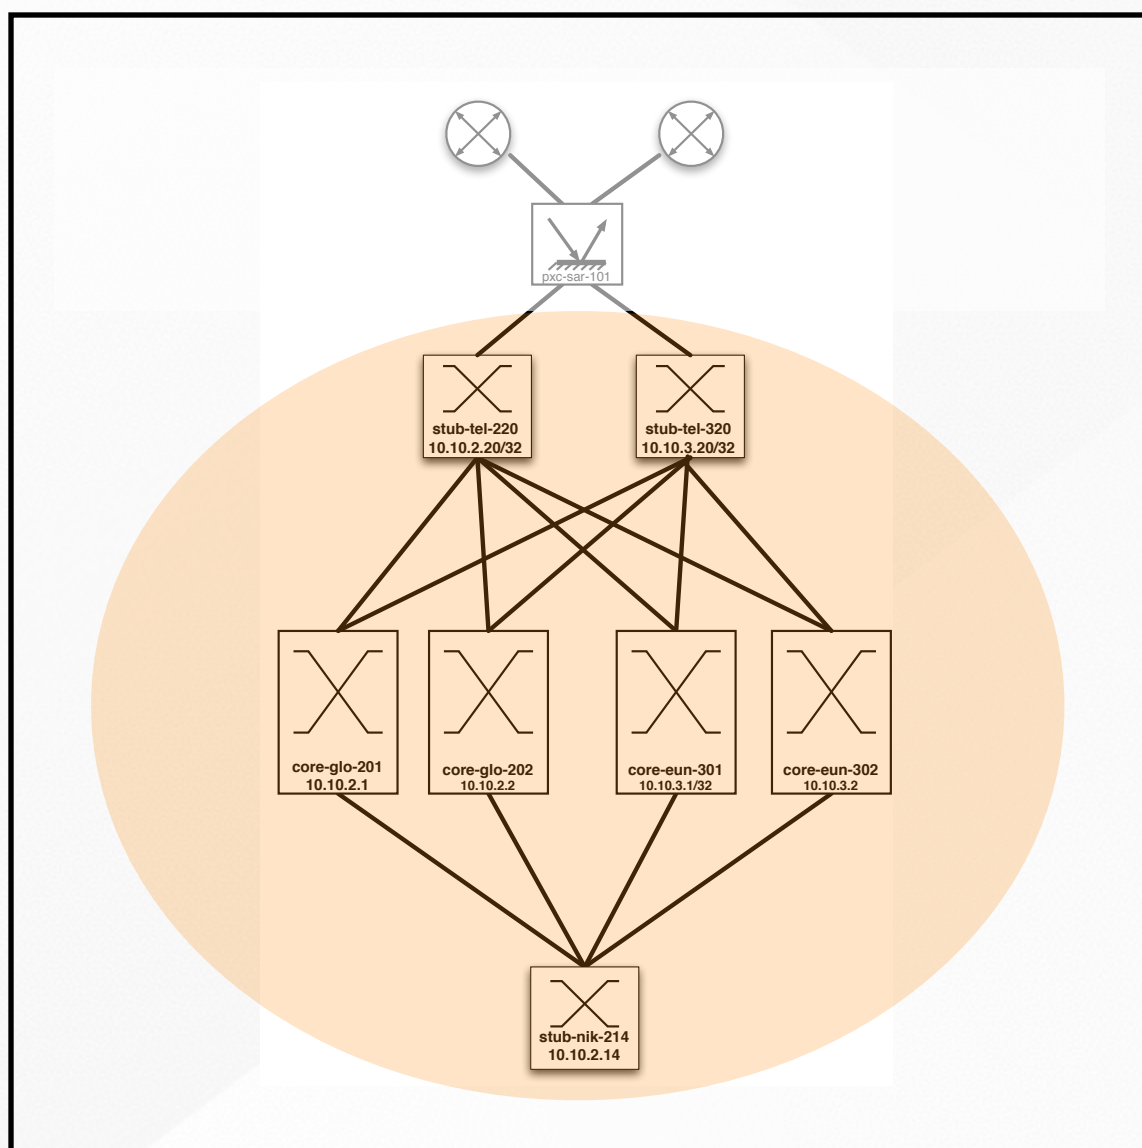


# Technical overview



- BFD for fast OSPF shutdown
- Replacing that by Ethernet OAM.
- MPLS to switch packets.
- RSVP-TE signalled LSPs over predefined paths.
- Primary and secondary (backup) paths are defined.

# Technical overview



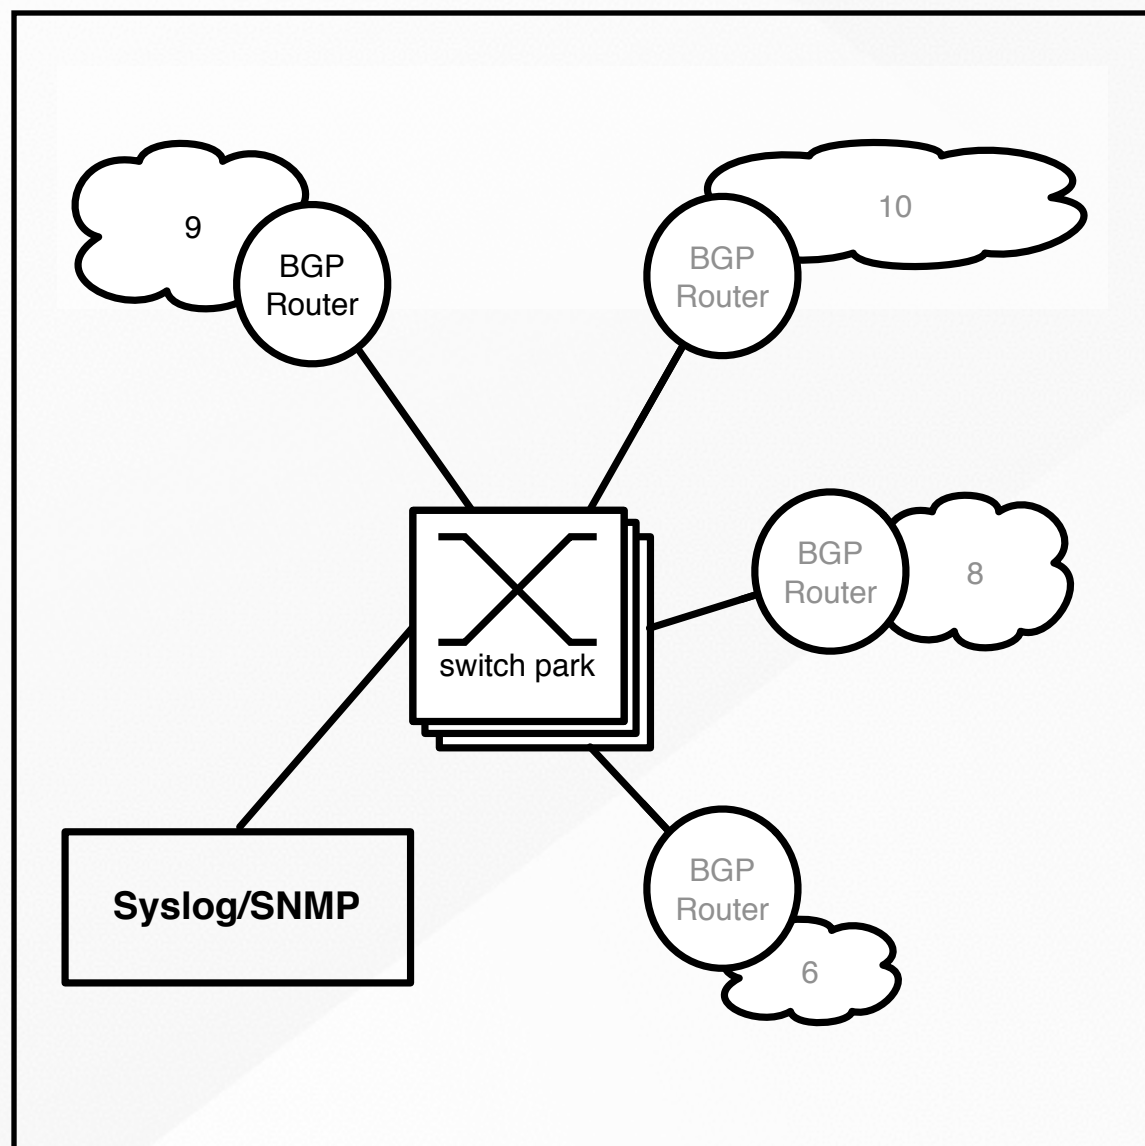
- **MPLS to switch packets.**
- **RSVP-TE signalled LSPs over predefined paths.**
- **Primary and secondary (backup) paths are defined.**
- **VPLS instance per peering LAN**
- **Static defined VPLS peers (LDP signalled)**
- **Load balanced over parallel LSPs over multiple core routers**



# Agenda

- Introduction.
- AMS-IX Organization.
- Technical Details.
- Operational aspects.
- Chicago setup.
  - Scalability.
- Questions / discussion.

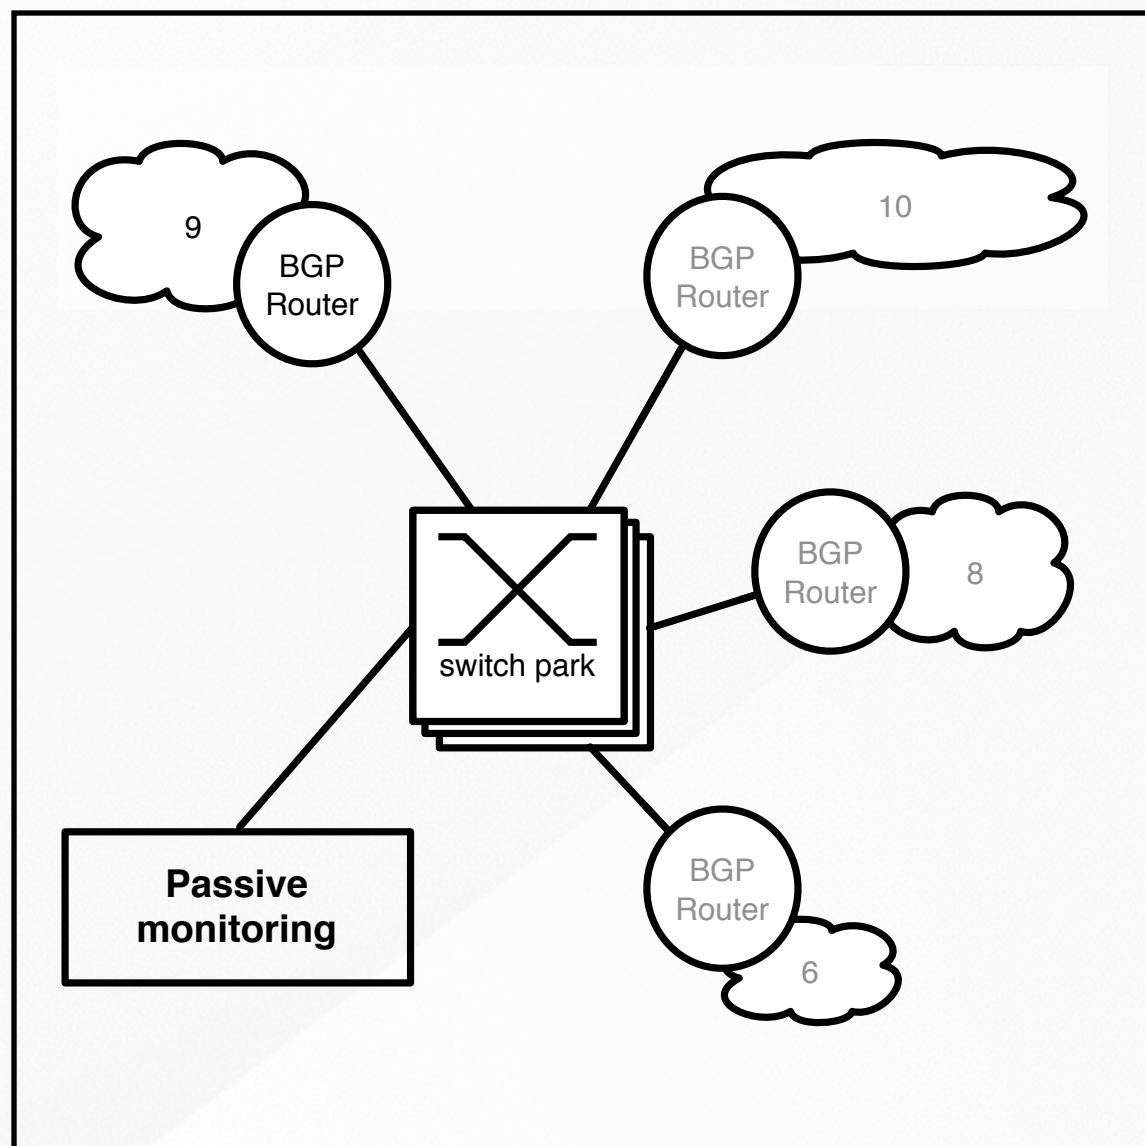
# Operational



- One MAC address per port
- No exception
  - One MAC address per 802.1q tag and physical port tuple for partner ports.
- Port shut down when loop is detected.
- Source MAC addresses that are allowed on other ports.

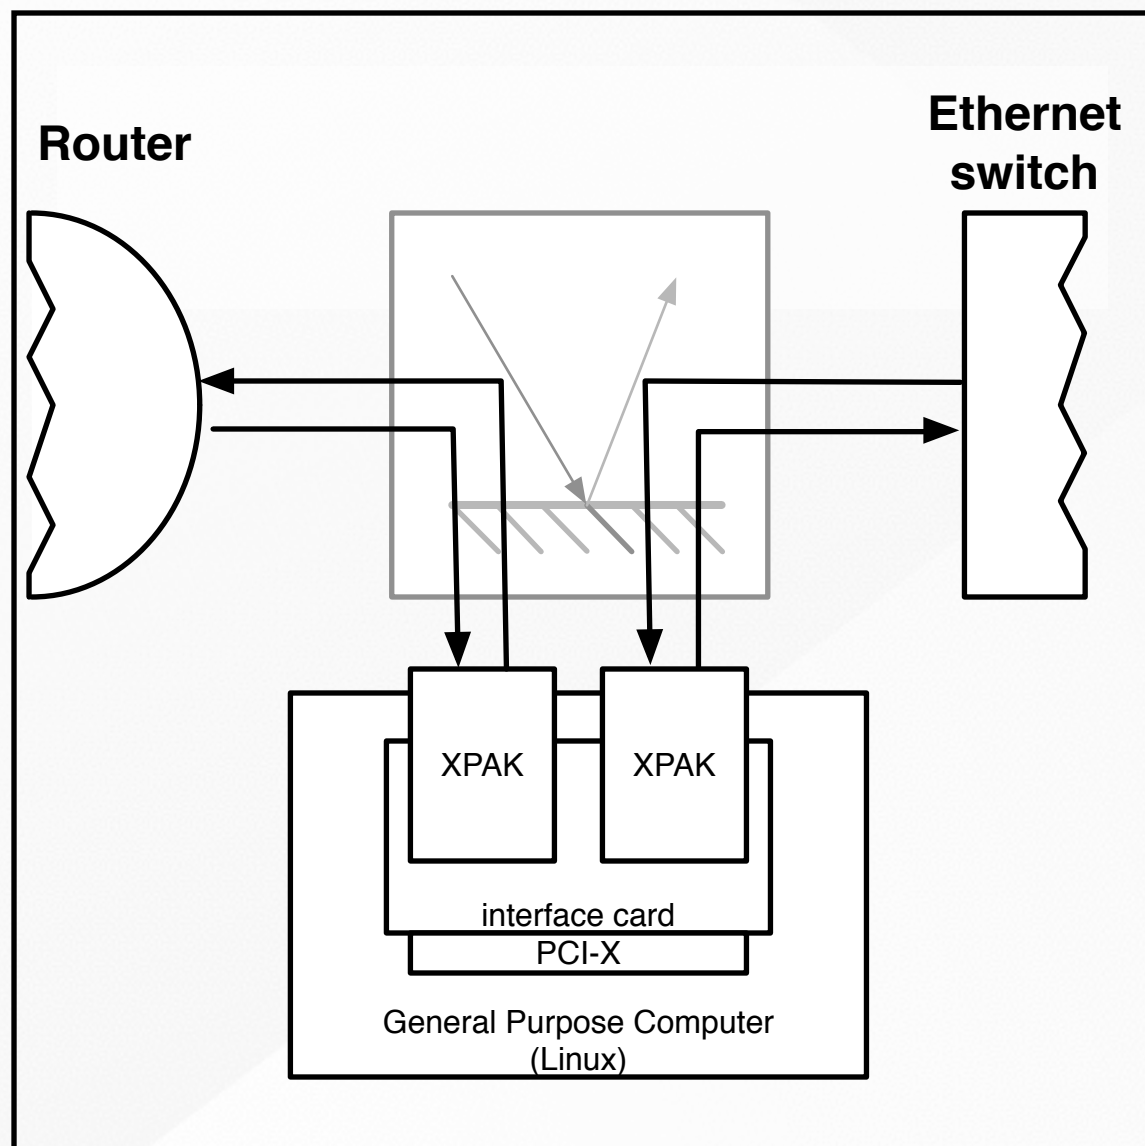


# Operational



- **Housekeeping**
  - **Broadcast traffic (ARP)**
  - **Multicast traffic (ICMPv6 ND)**
  - **Unwanted traffic**
    - **CDP, EDP, FDP, etc.**
    - **DEC MOP**
    - **DHCP**
    - **ARP for other subnets.**
    - **etc.**

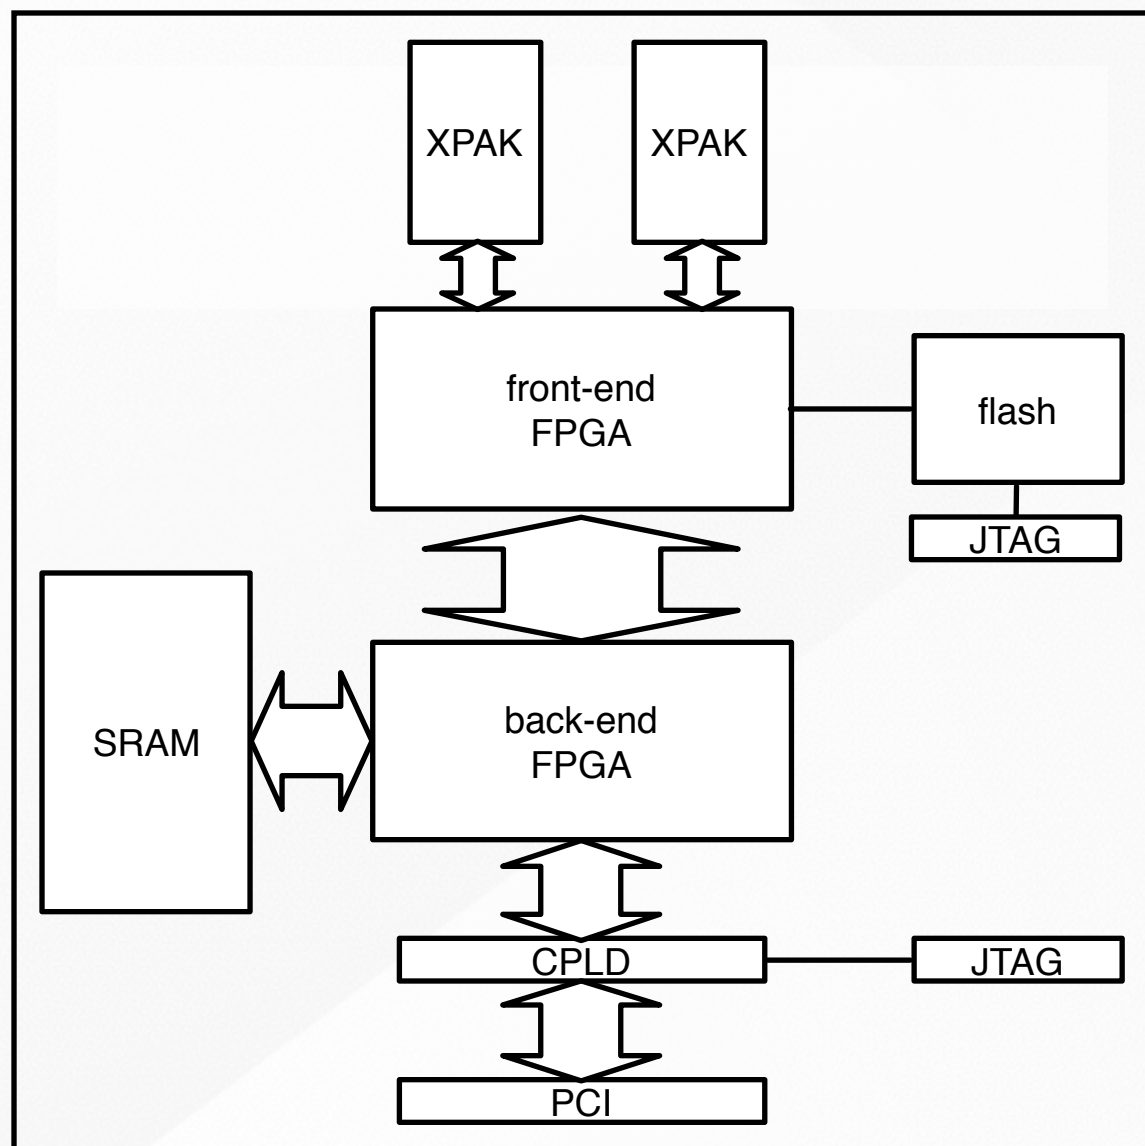
# Operational



- Targeted passive monitoring.
- For troubleshooting only.

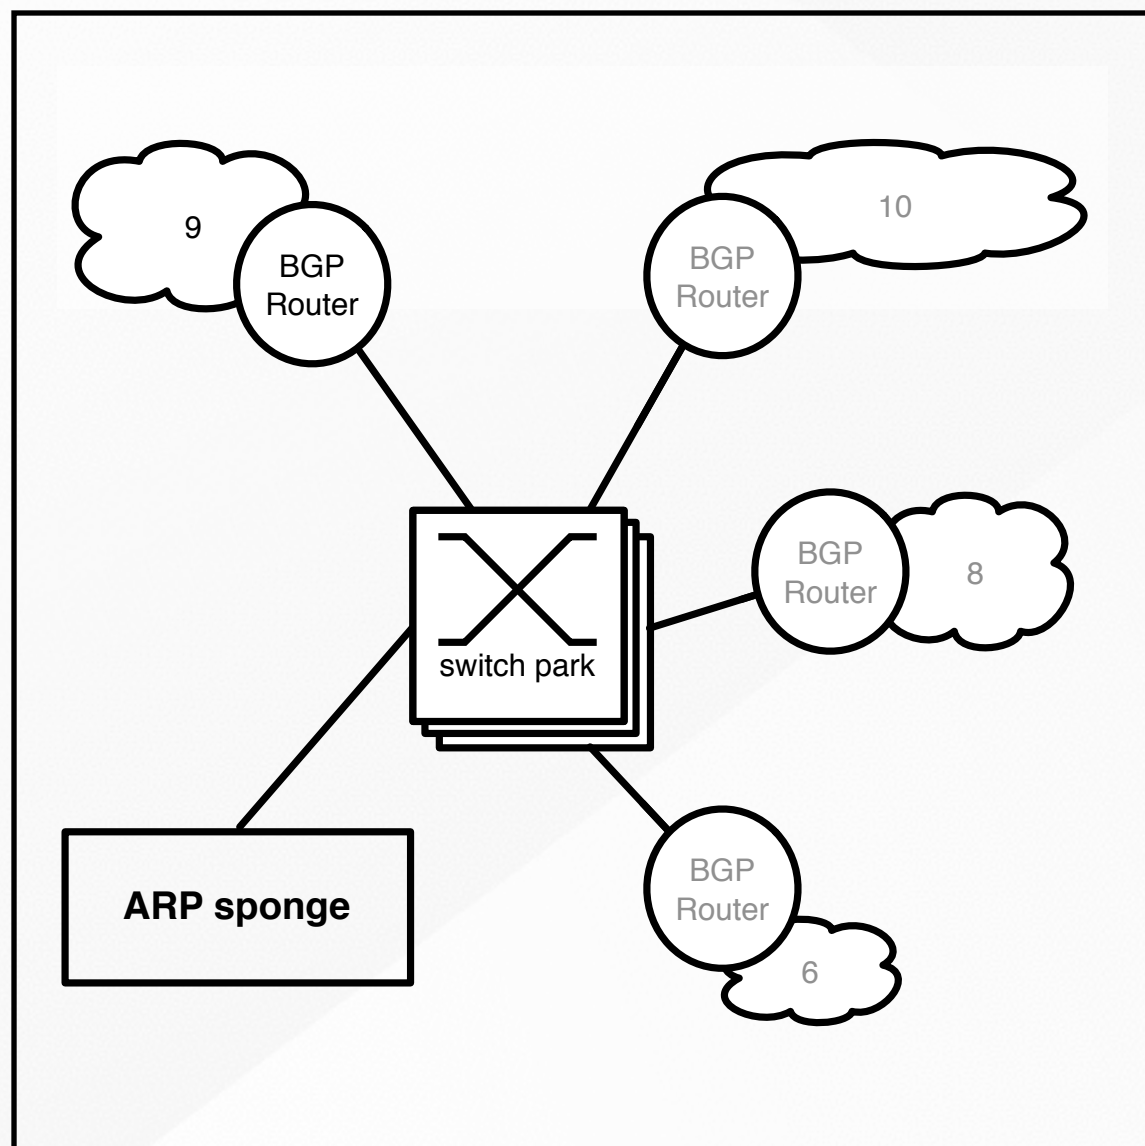


# Operational



- Targeted passive monitoring.
- For troubleshooting only.
- FPGA based solution for 10G
  - Maximum frame rate monitoring.
- Integrated in libpcap.
  - tcpdump, wireshark

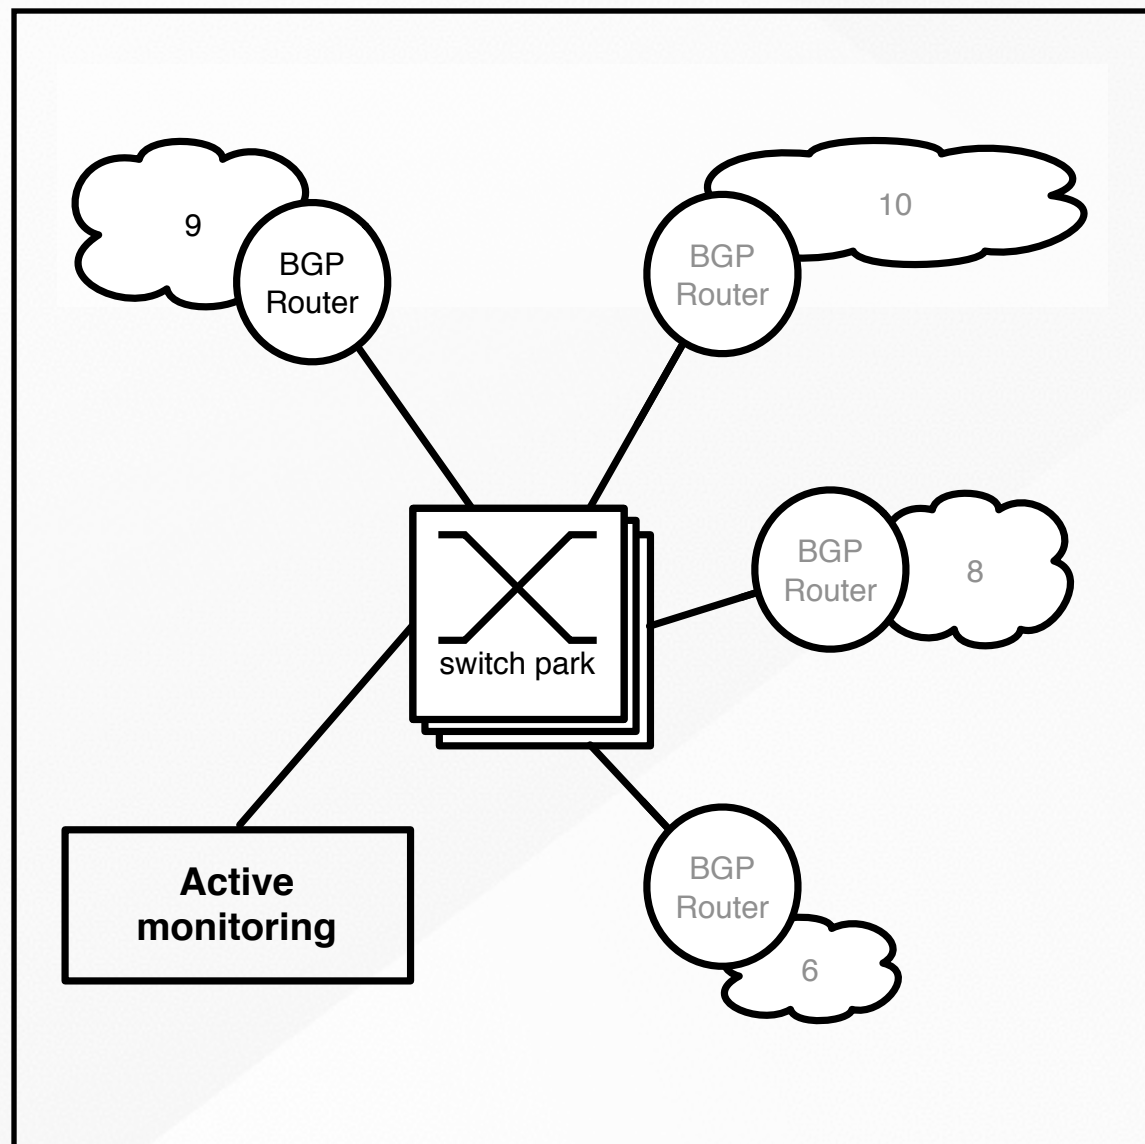
# Operational



- **ARP sponge**
  - **Friendly ARP spoofer to mitigate ARP broadcasts level.**
  - **Answers ARP queries for unused and non-operational IPv4 addresses.**
  - **Stops doing so when it notices that the address is used.**

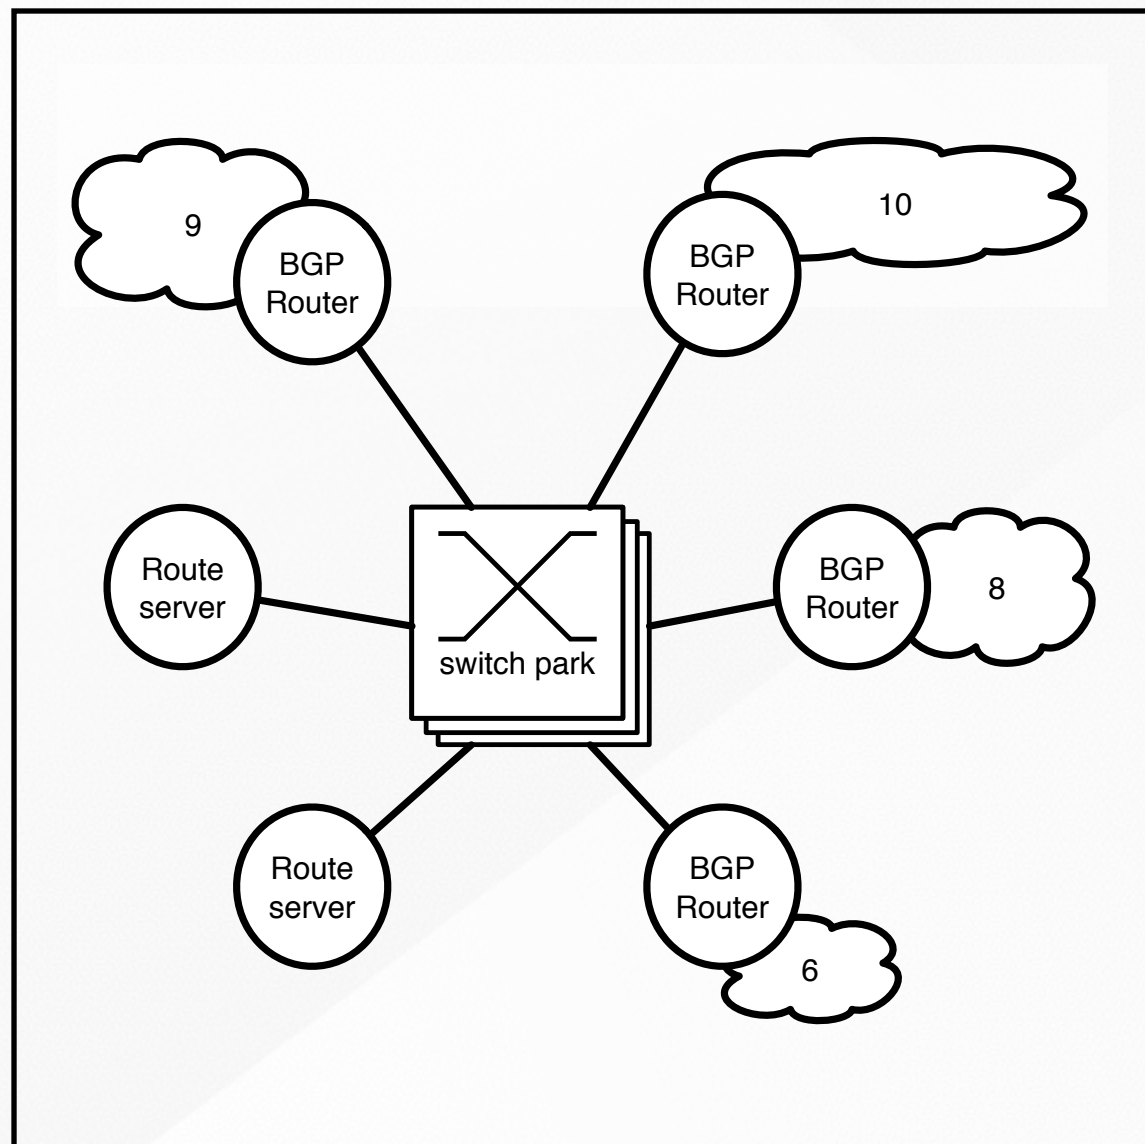


# Operational



- **Active monitoring**
  - **Proxy ARP**
  - **BGP sessions**

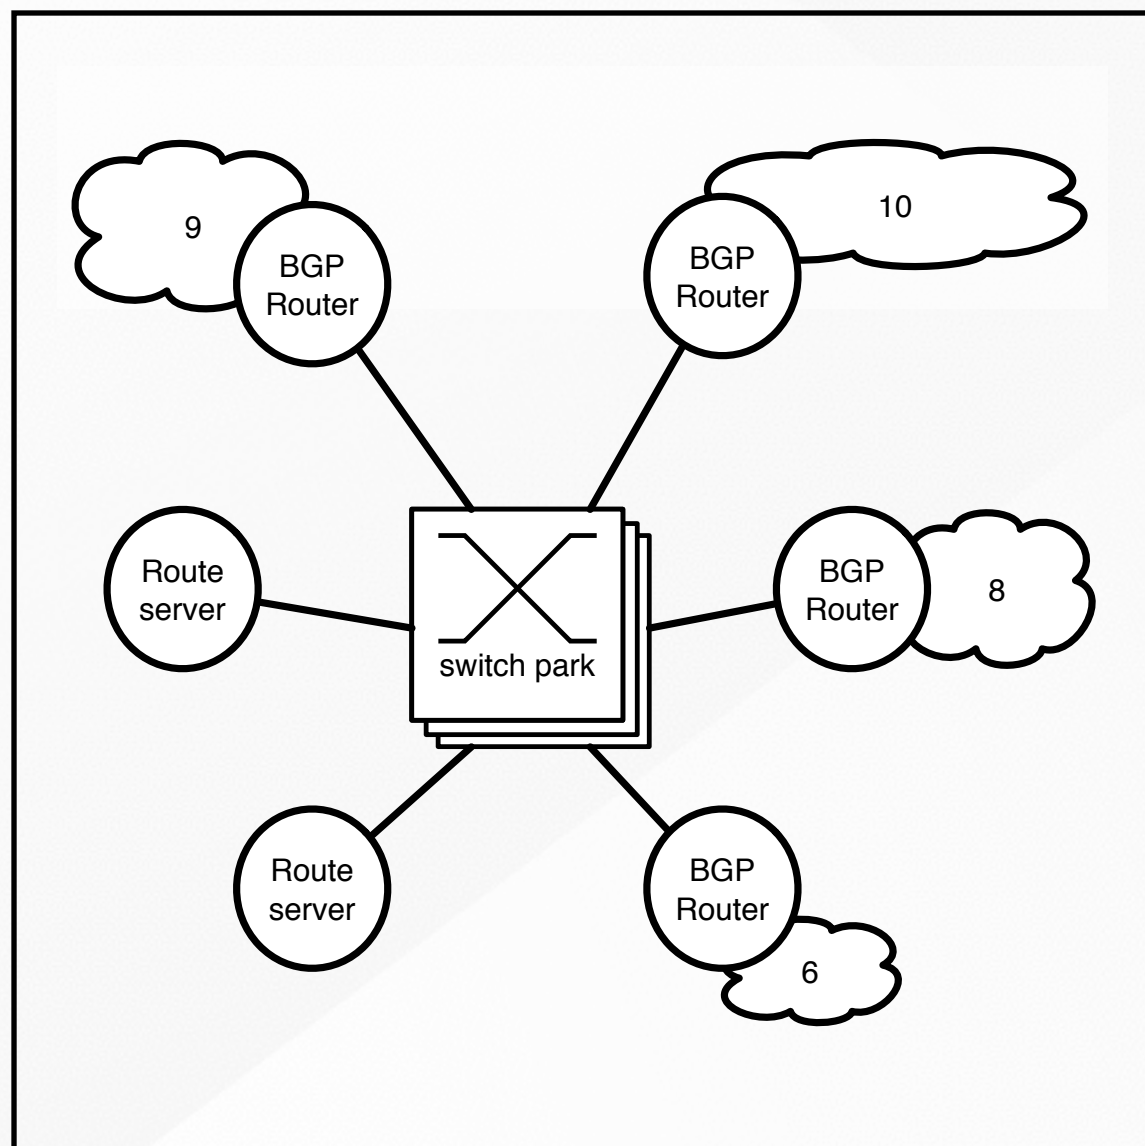
# Operational



- **Active monitoring**
  - **Proxy ARP**
  - **BGP sessions**
    - **Own network ASI200**
    - **Route servers AS6777**



# Operational



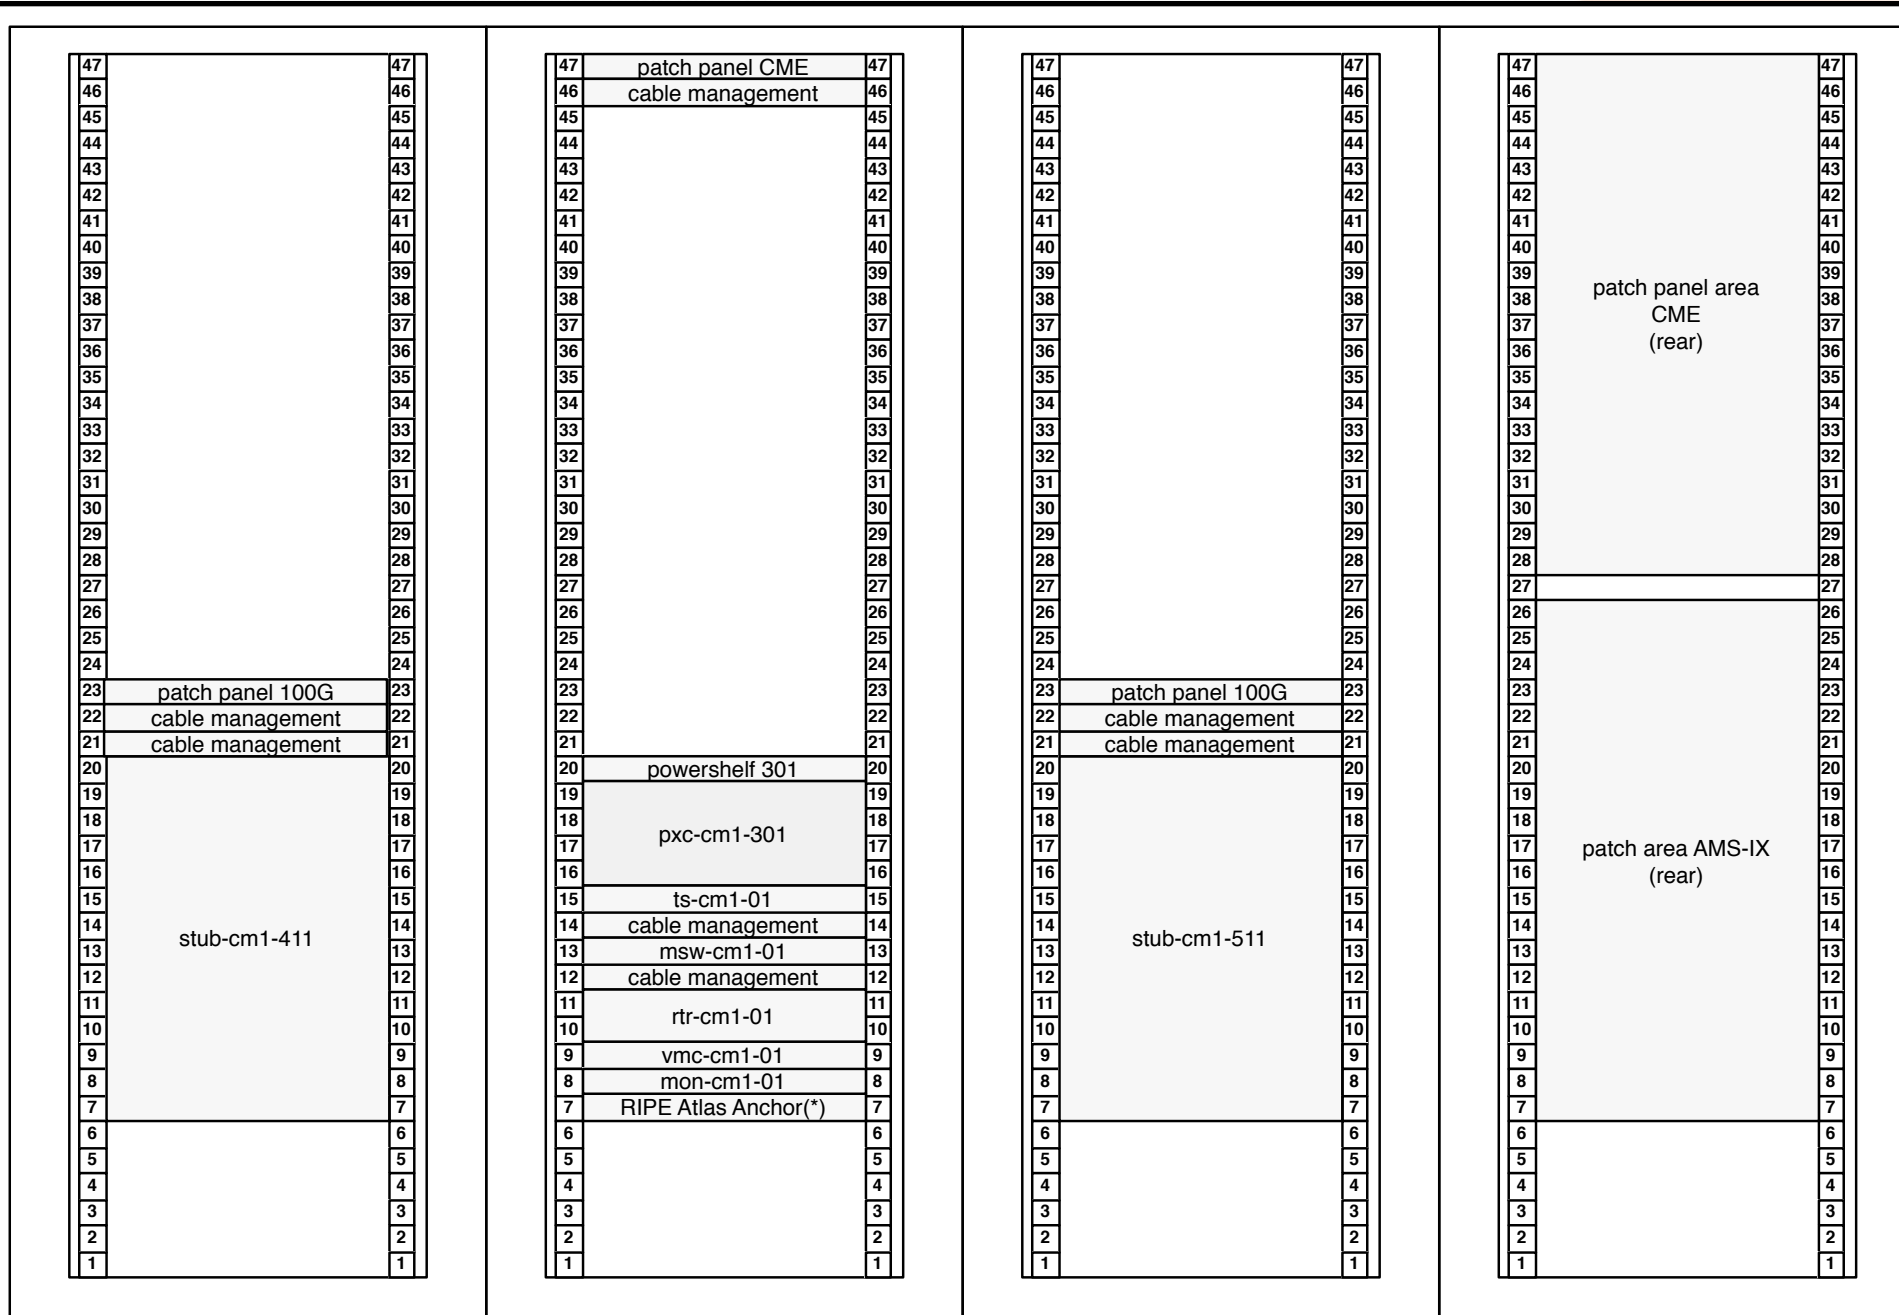
- **Route servers**
  - **Two BGP sessions to get routes from all participants (unless filtered).**
- **BGP next-hop attribute so payload traffic goes directly from peer to peer.**
- **IRRdb based filtering per ASN.**

# Agenda

- Introduction.
- AMS-IX Organization.
- Technical Details.
- Operational aspects.
- **Chicago setup.**
  - **Scalability.**
- Questions / discussion.

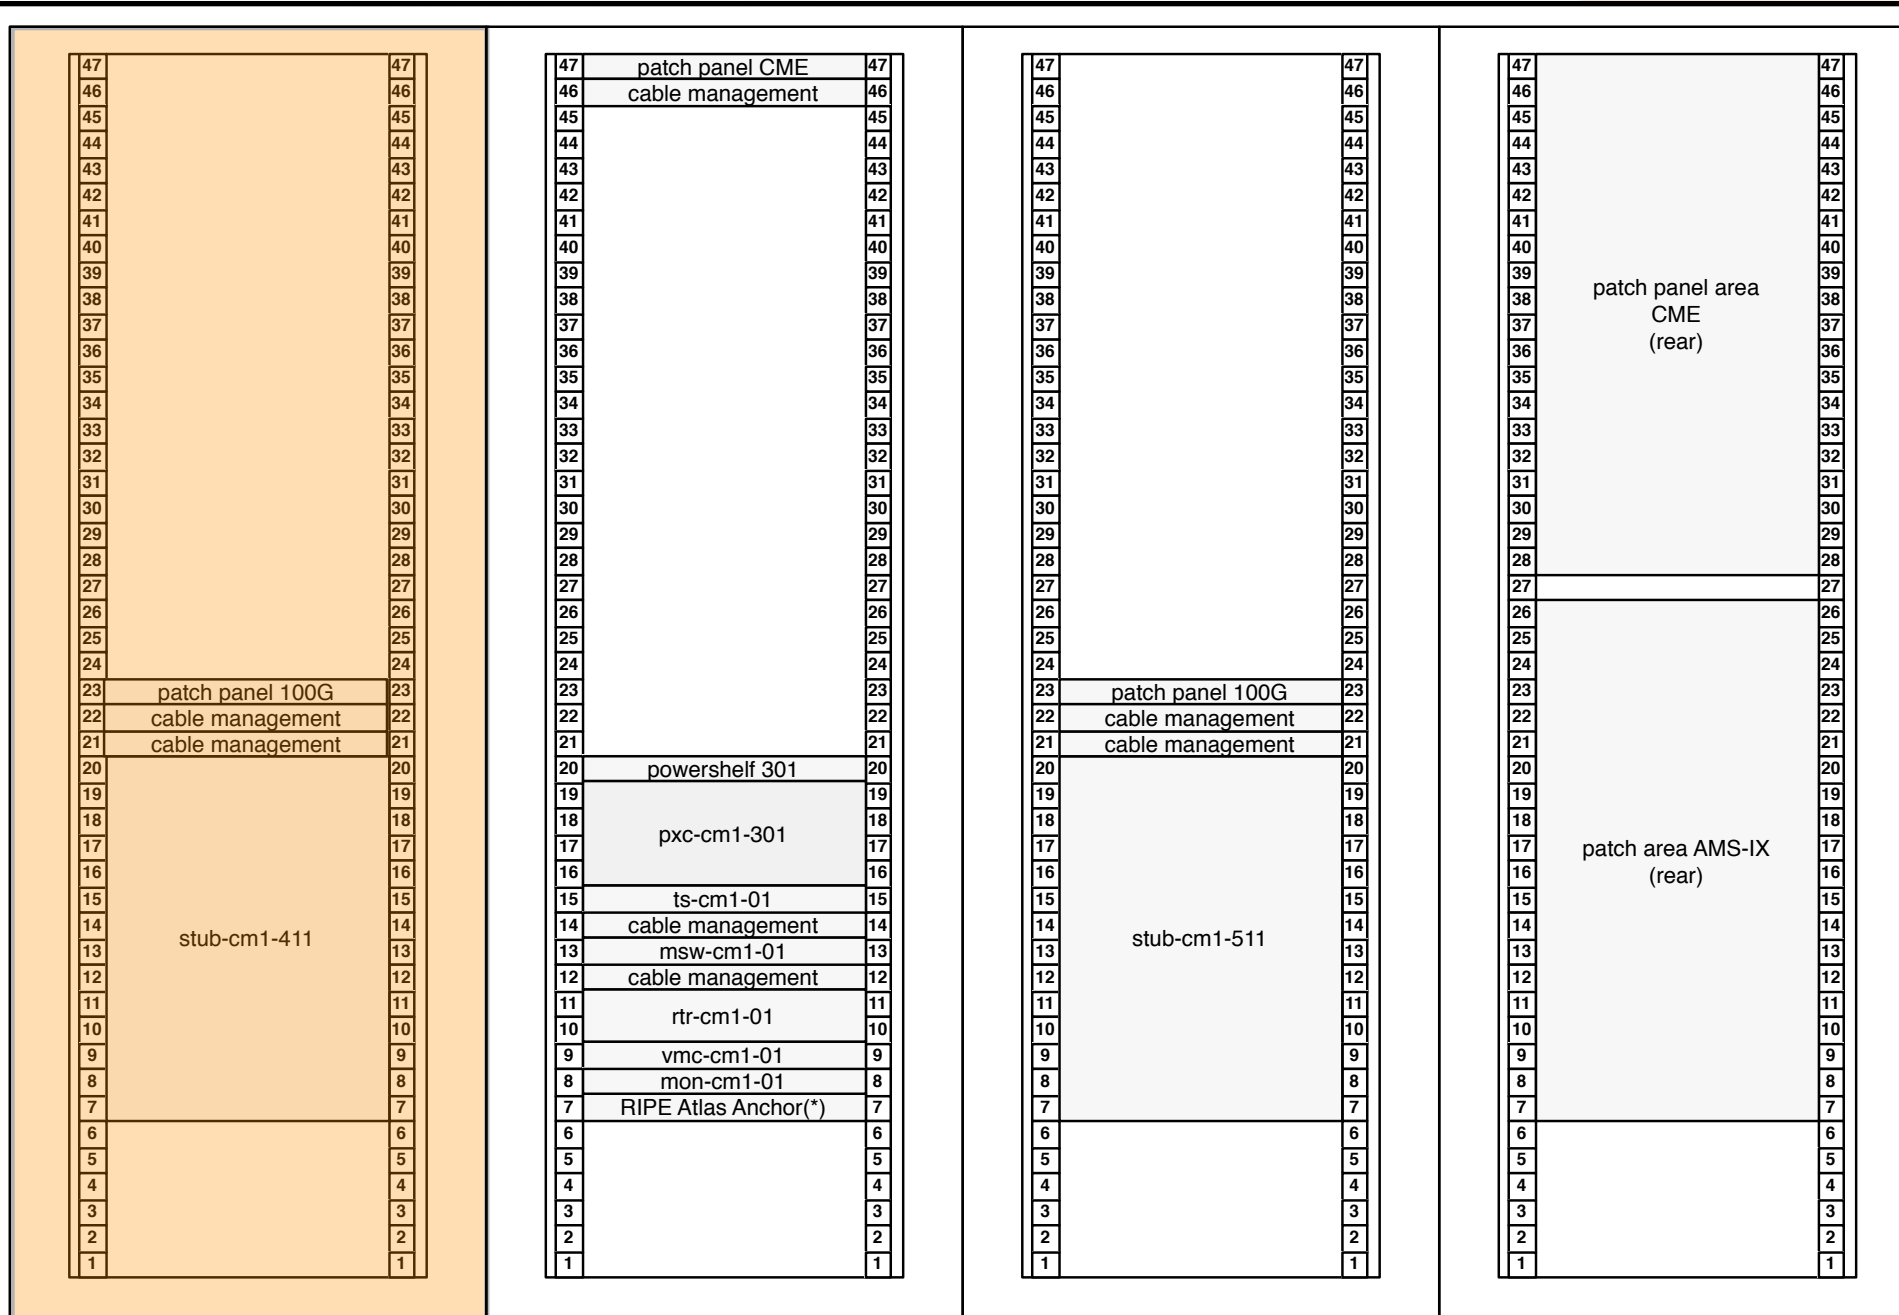


# Scalability



Cold Isle

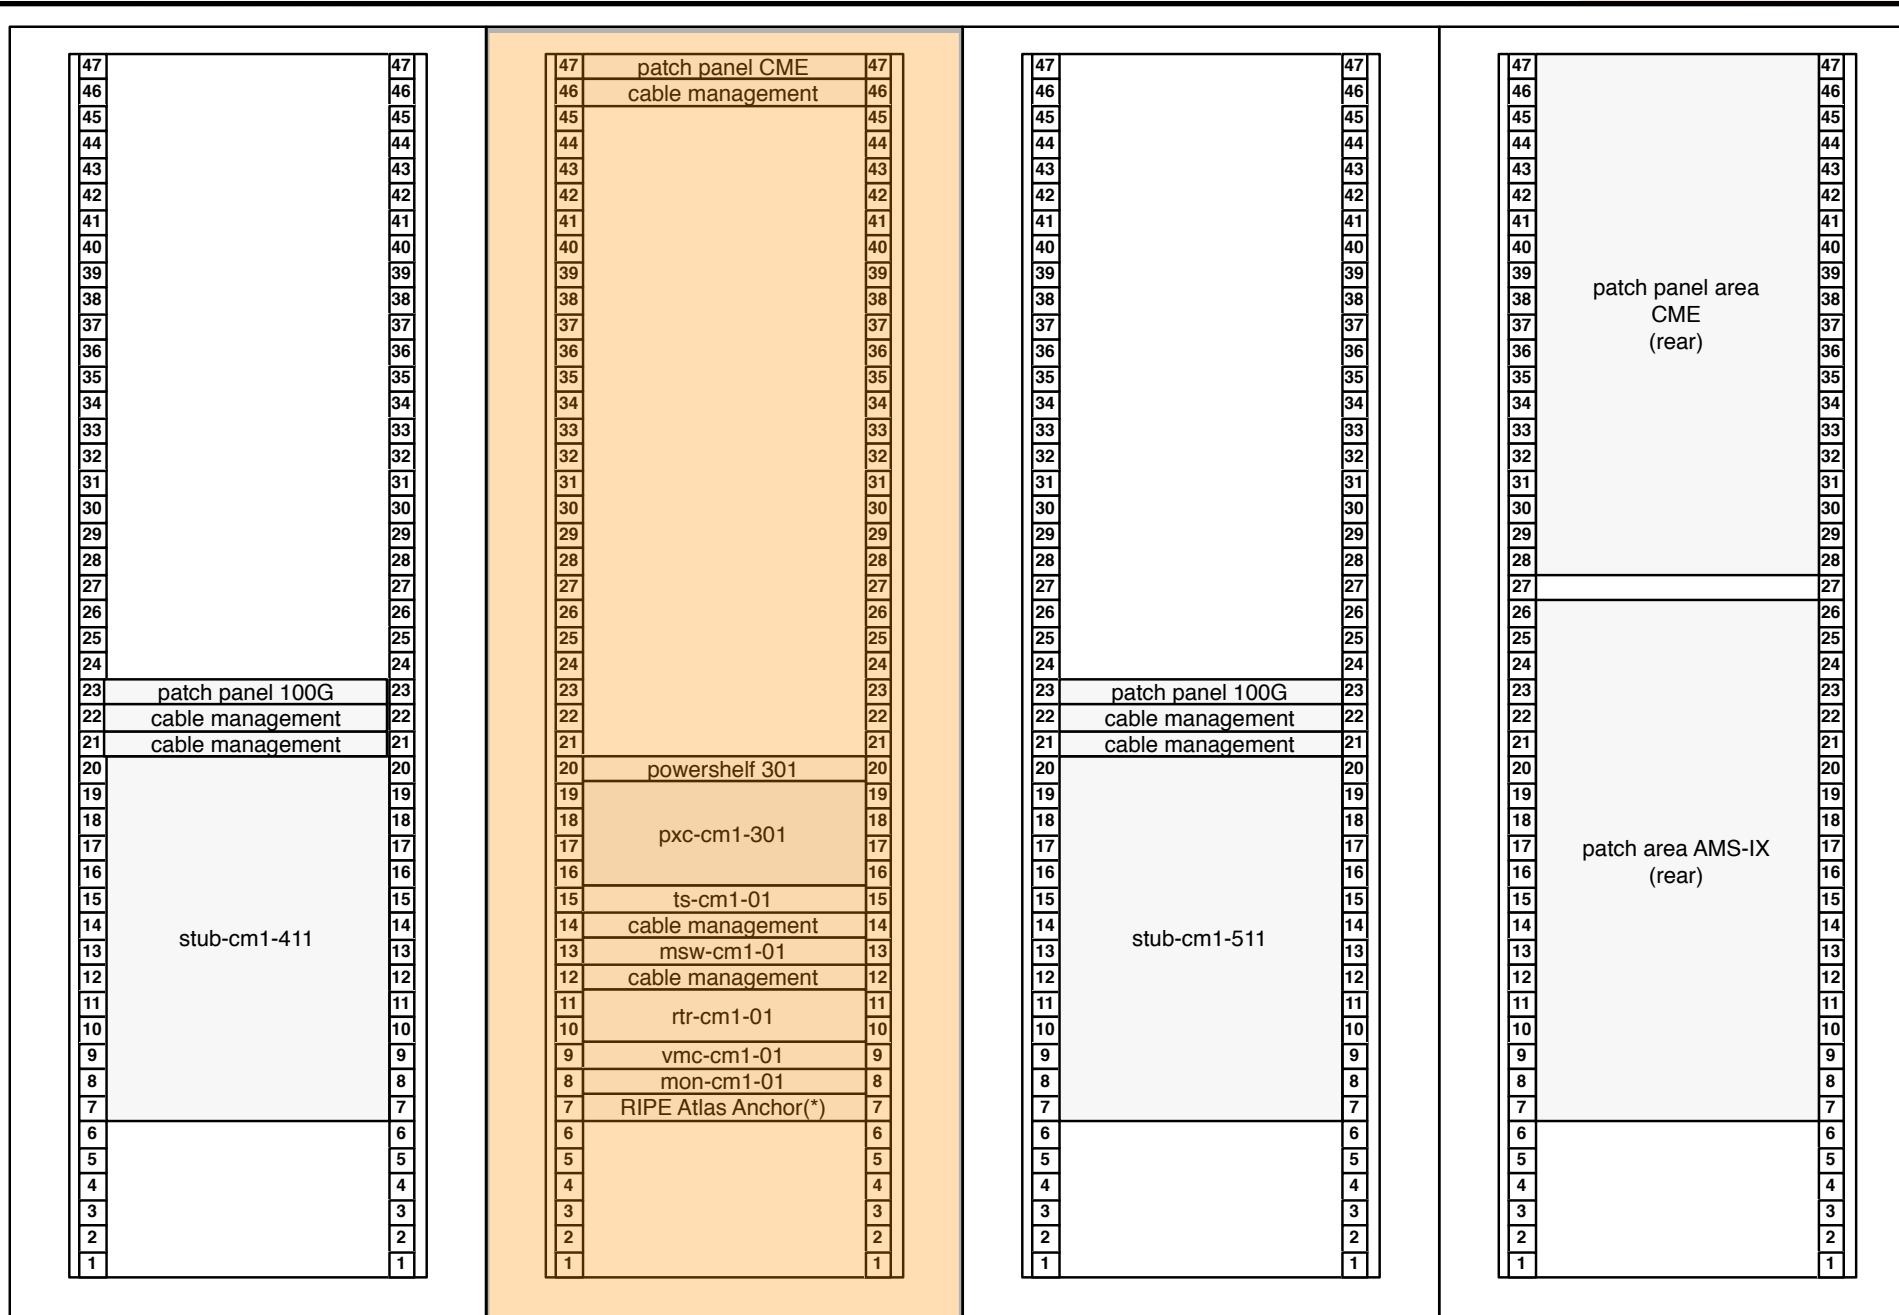
# Scalability



Cold Isle

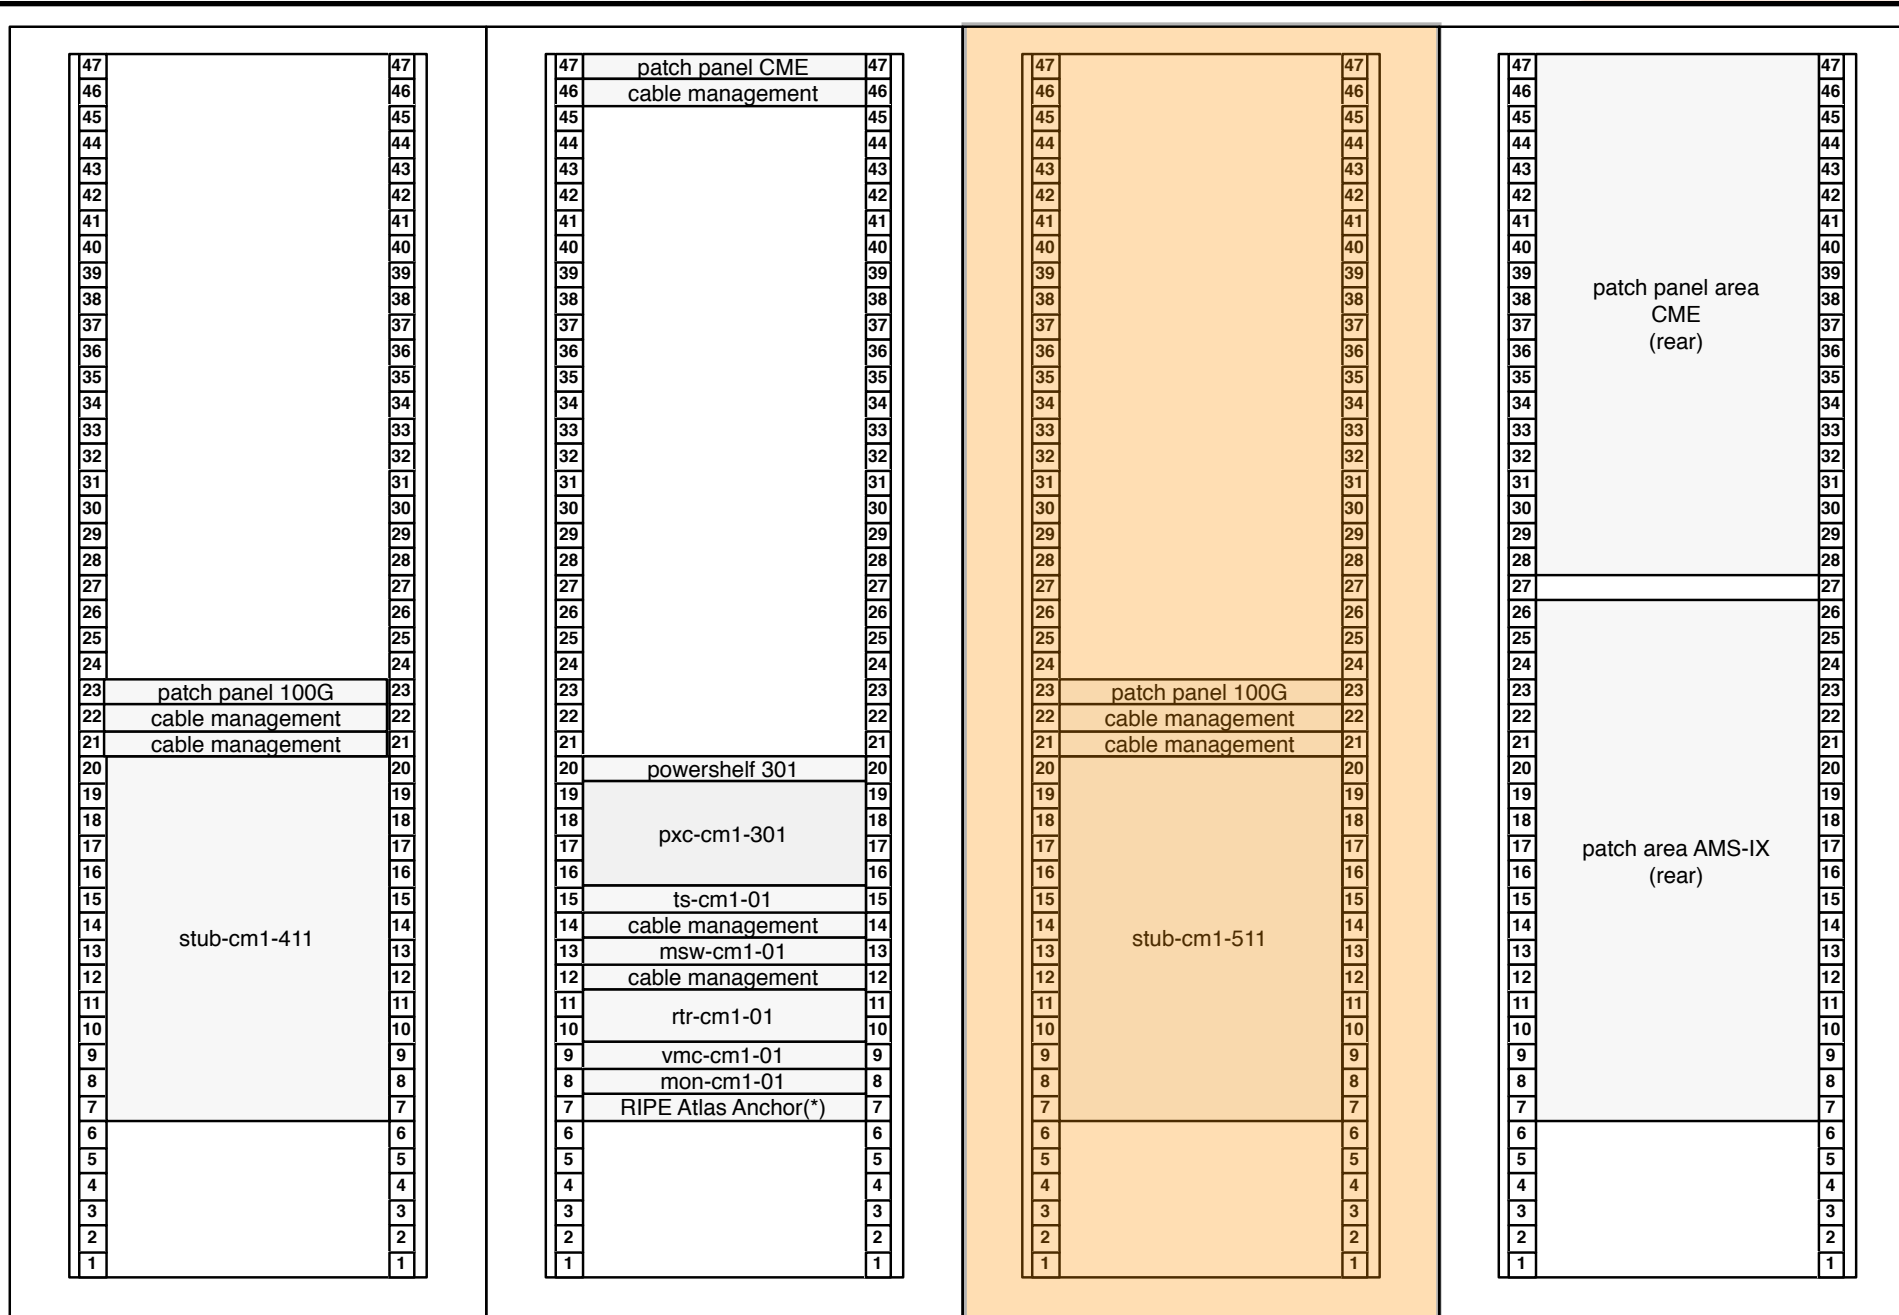


# Scalability



Cold Isle

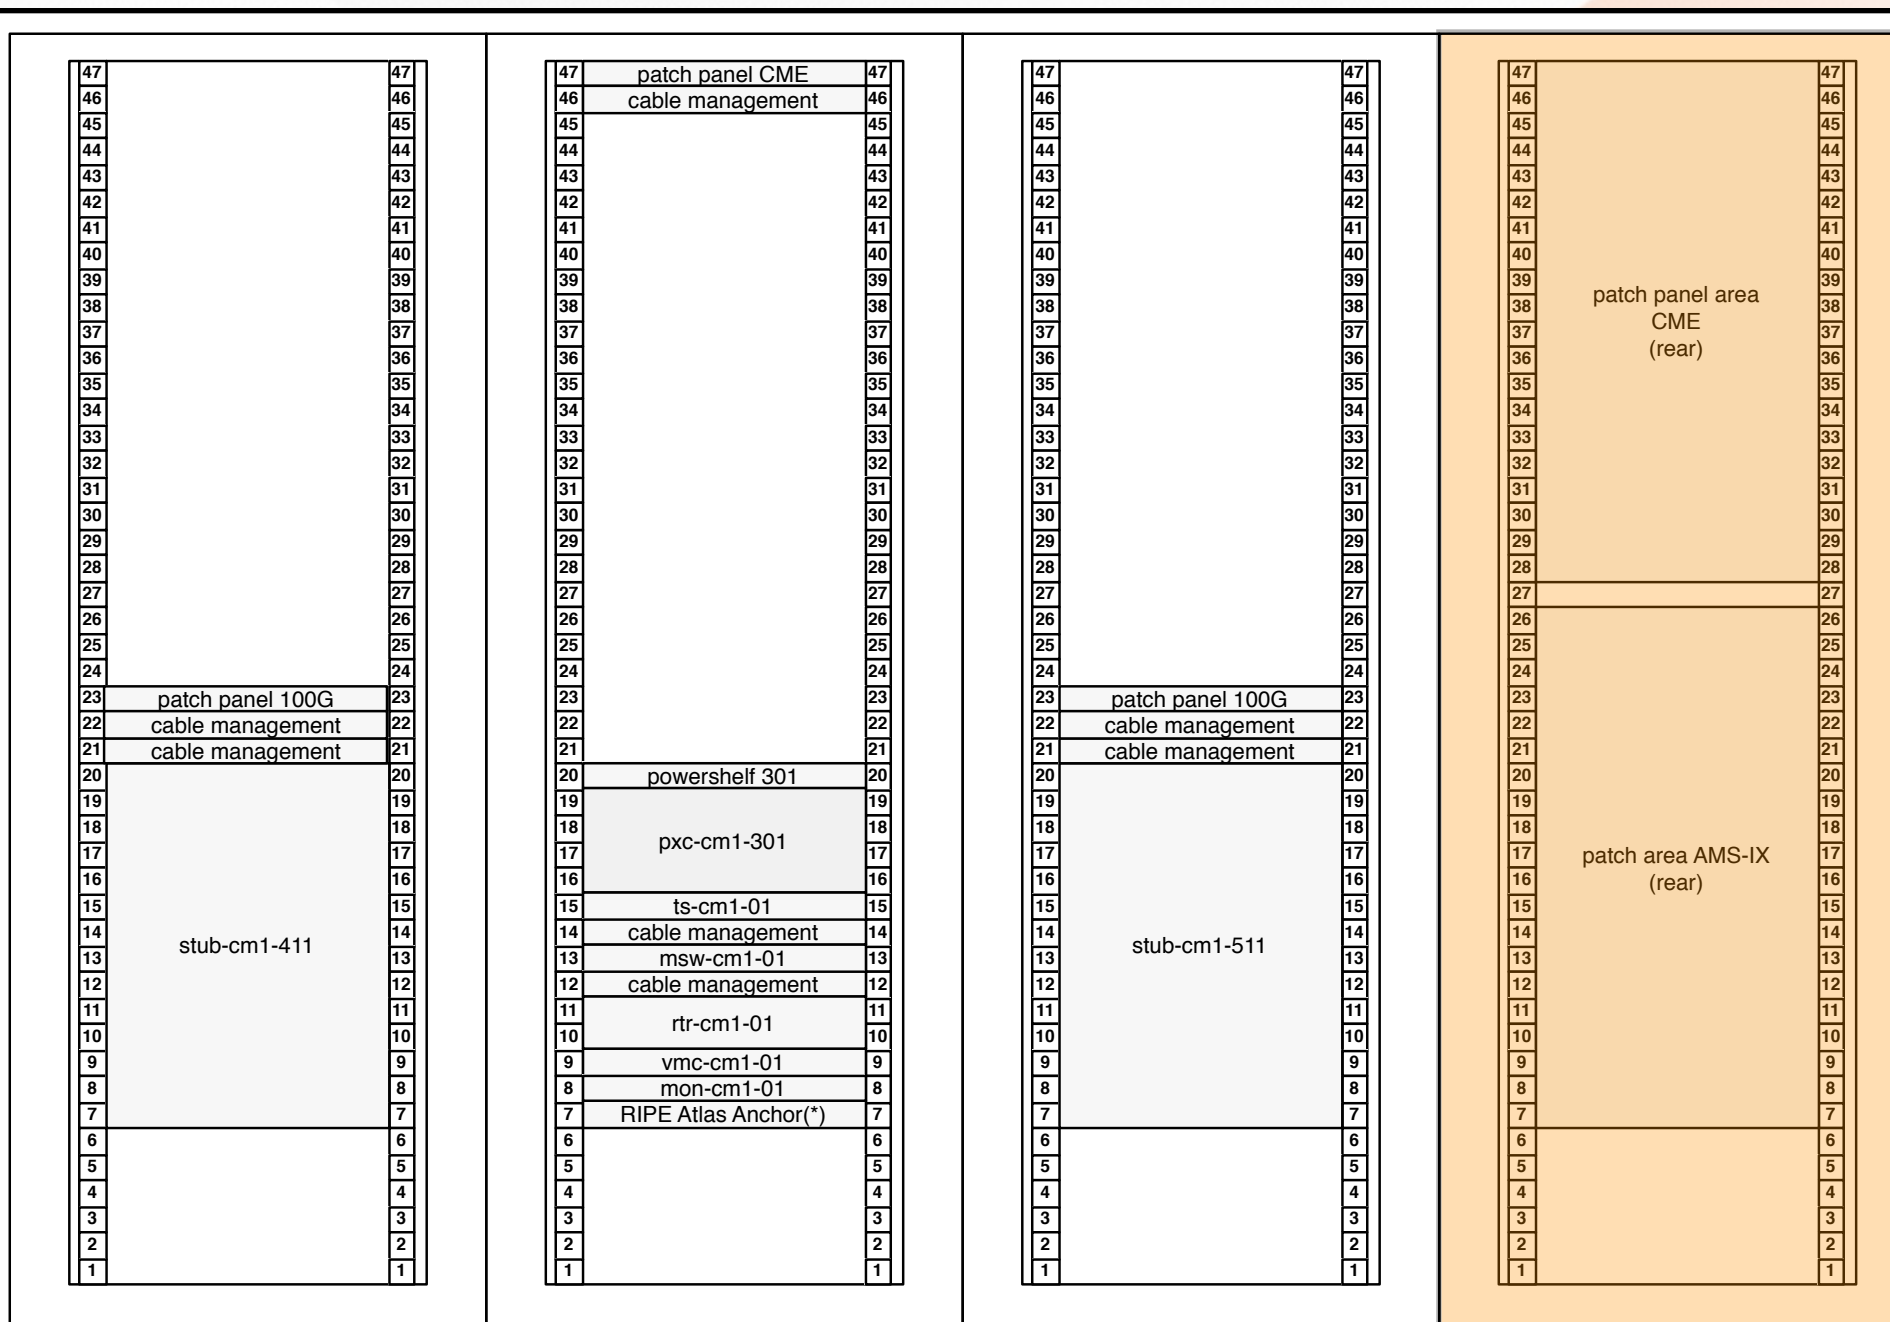
# Scalability



Cold Isle

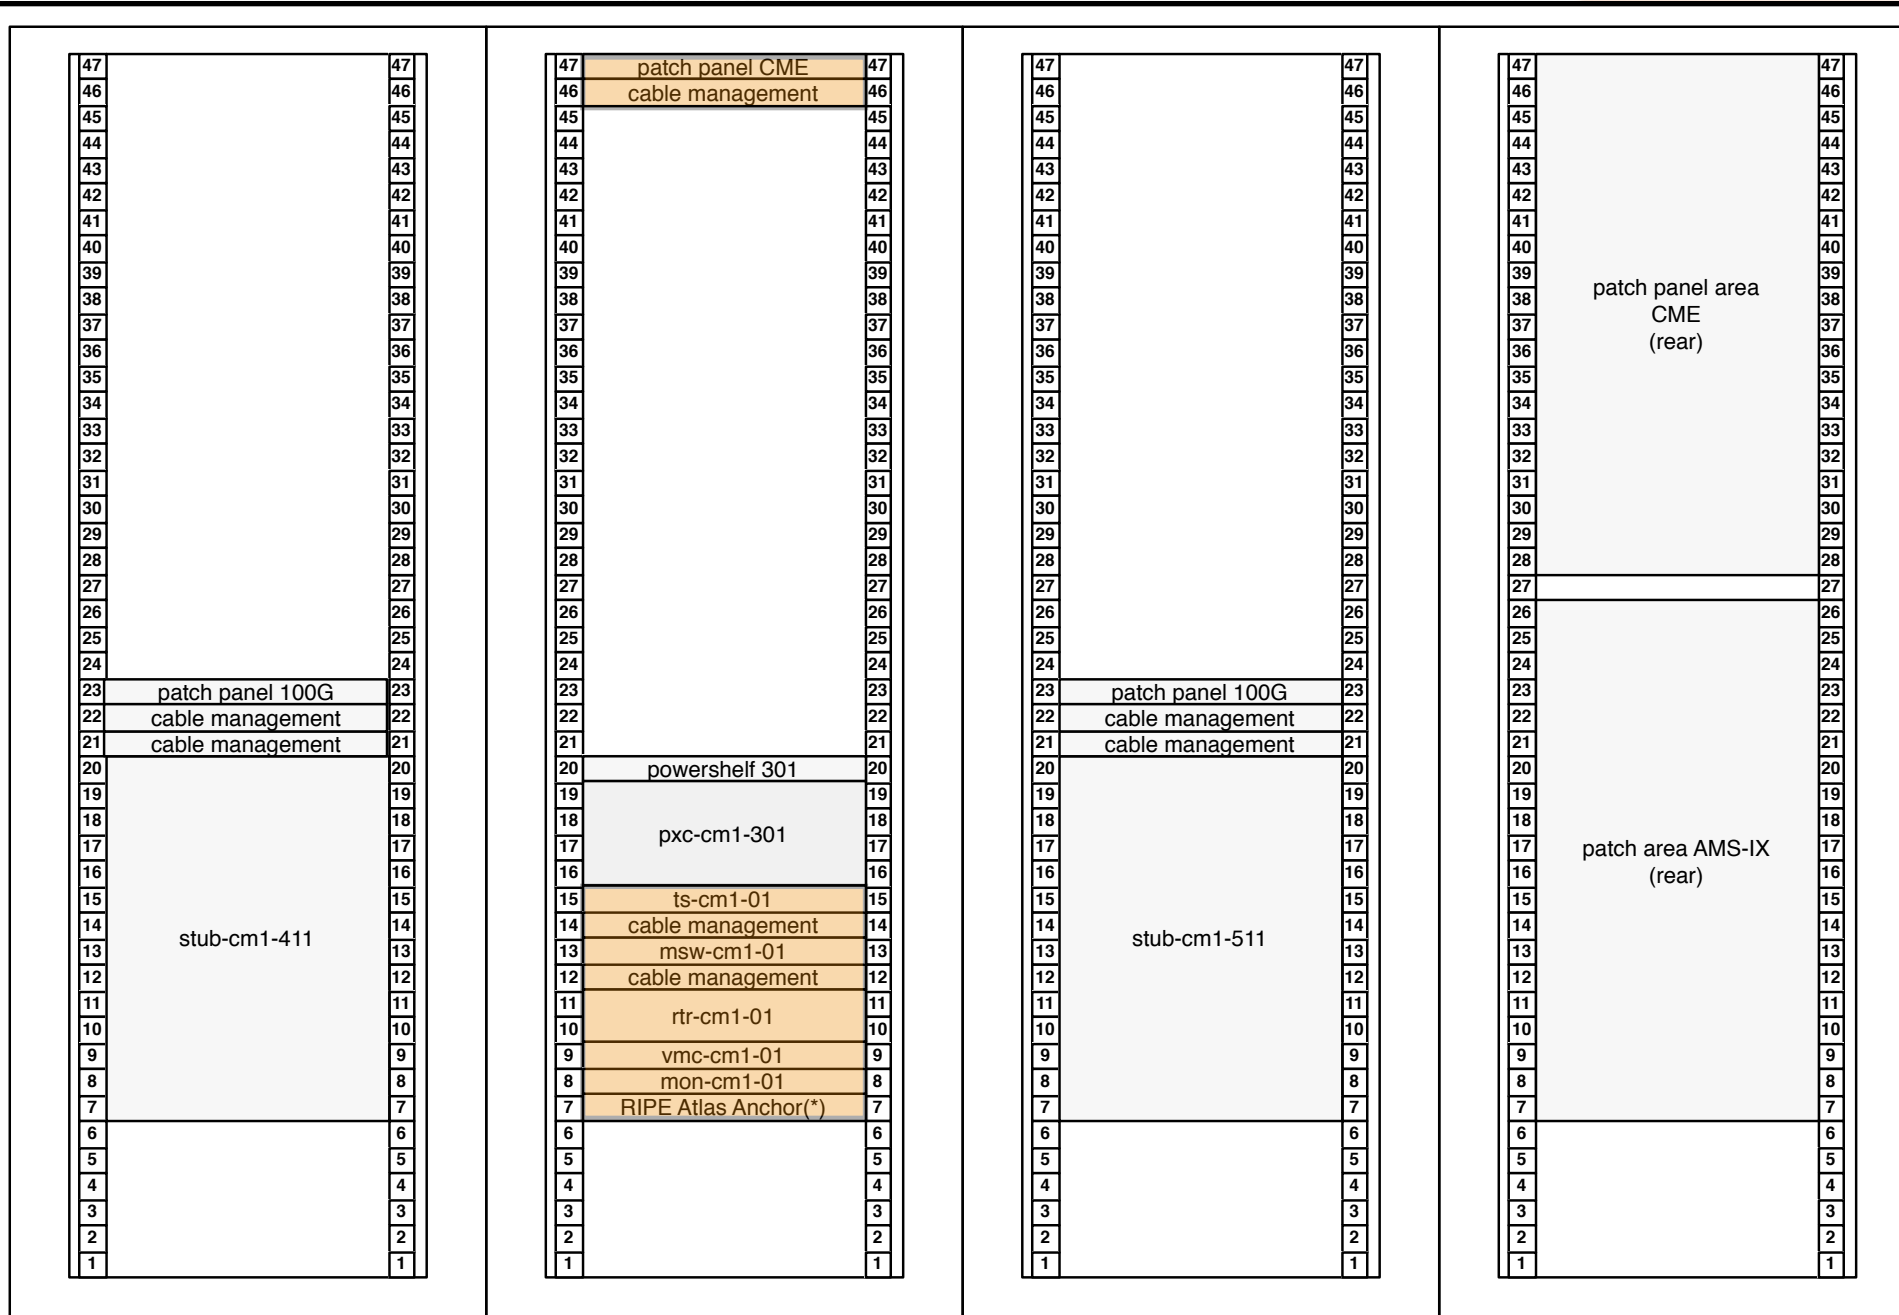


# Scalability



Cold Isle

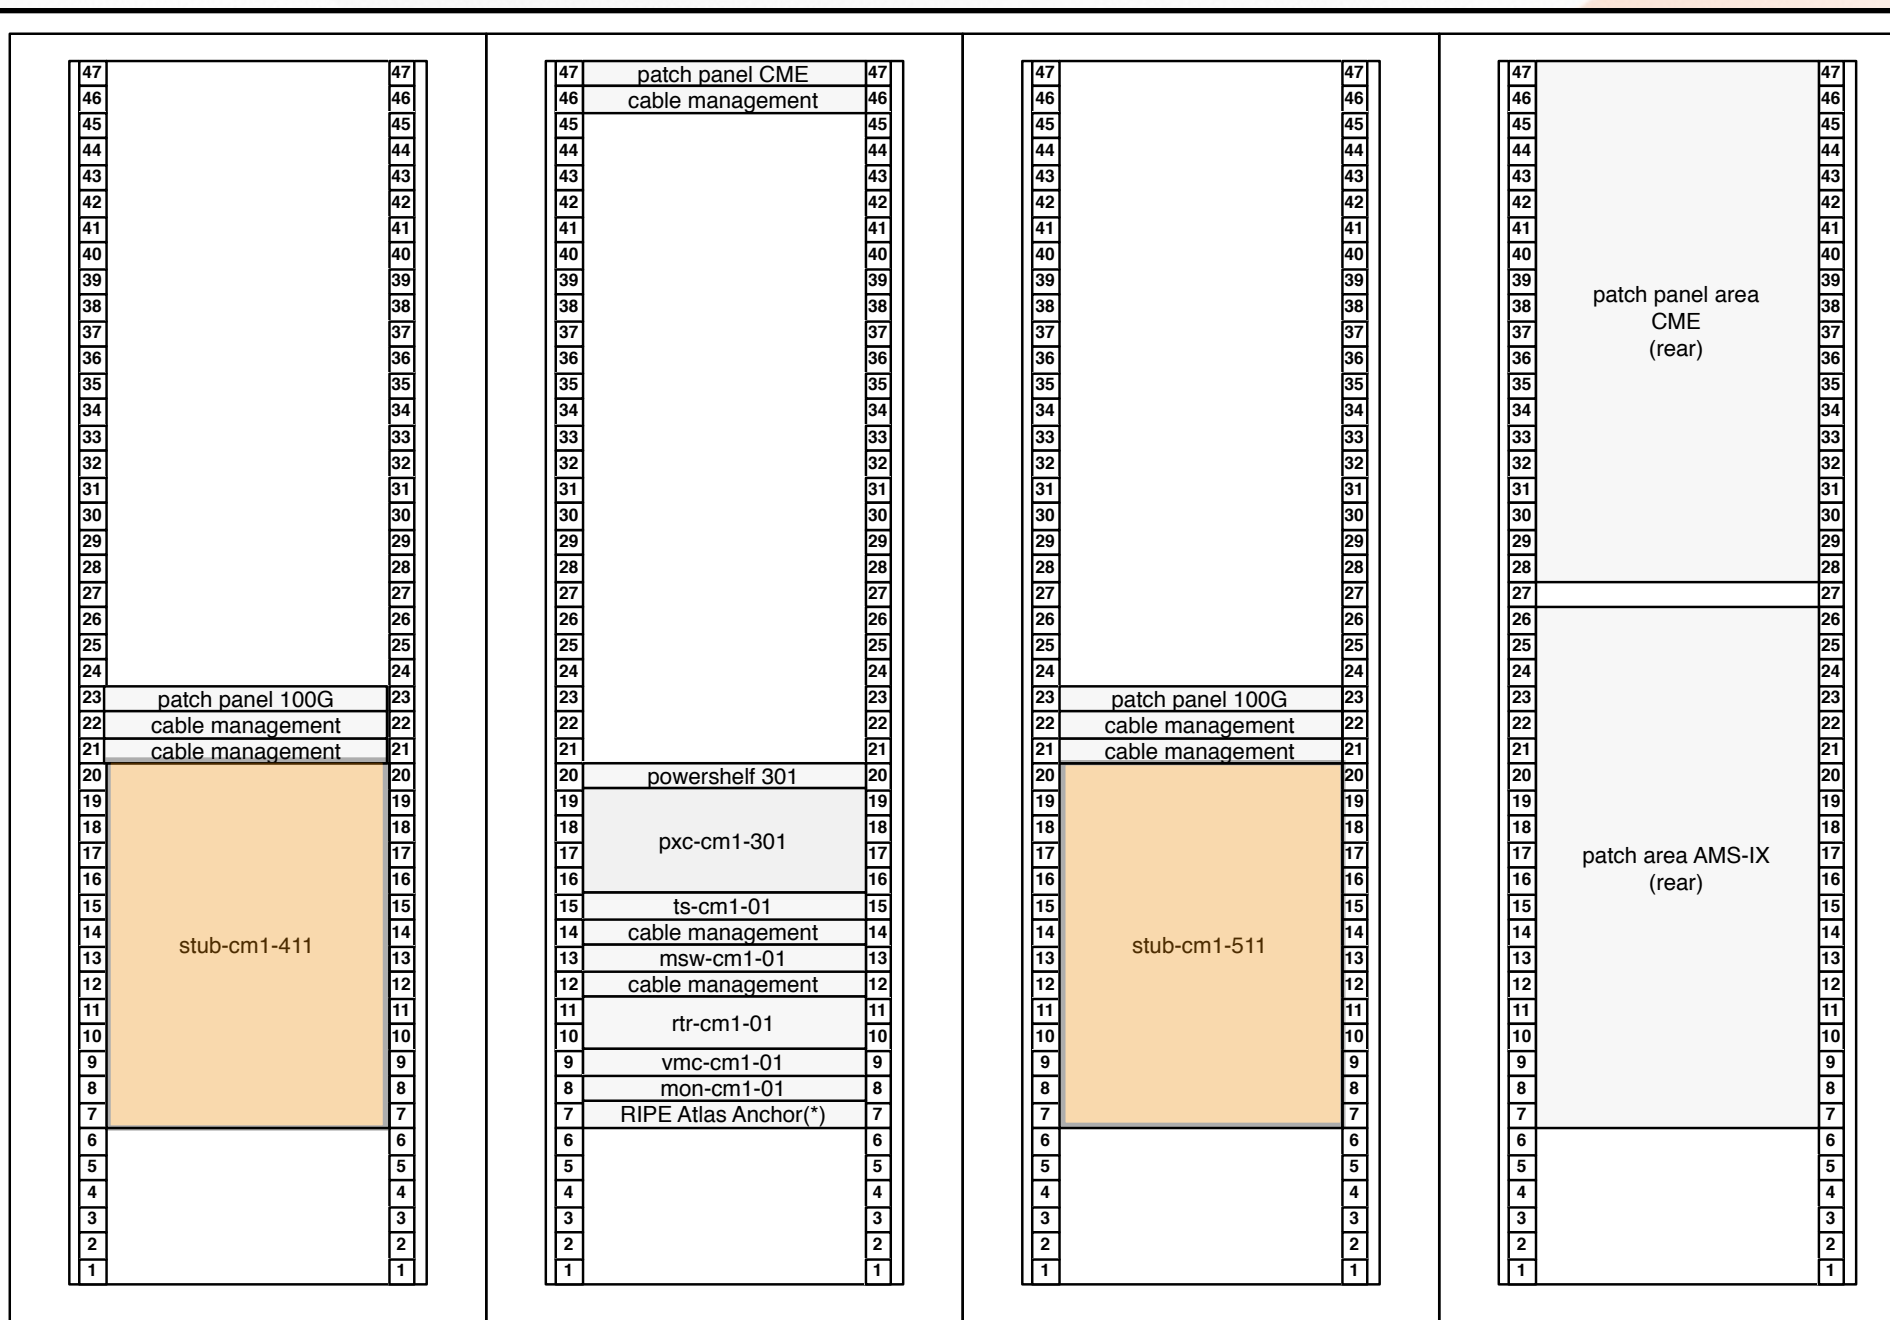
# Scalability



Cold Isle

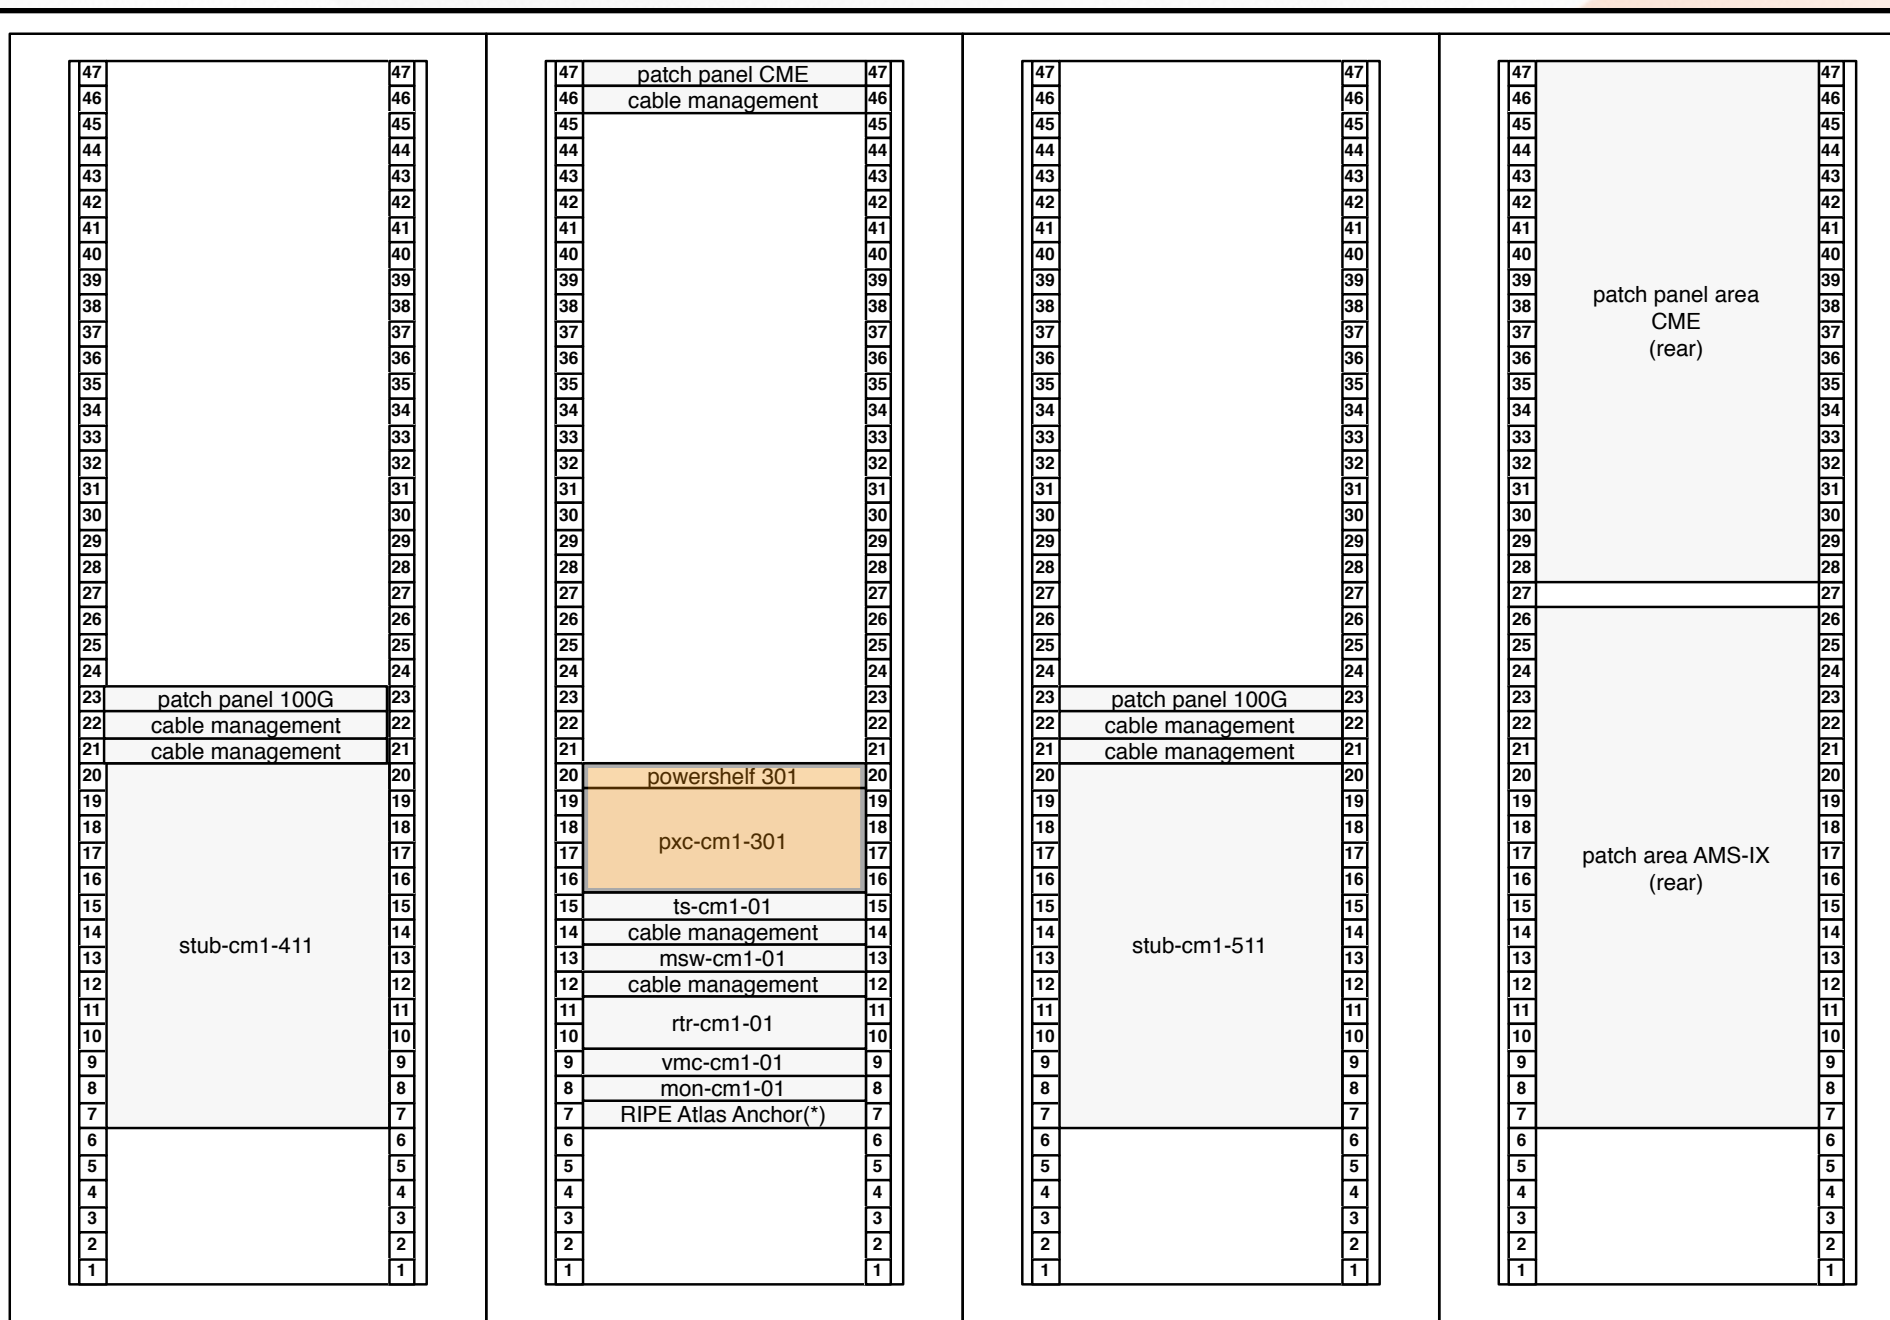


# Scalability



Cold Isle

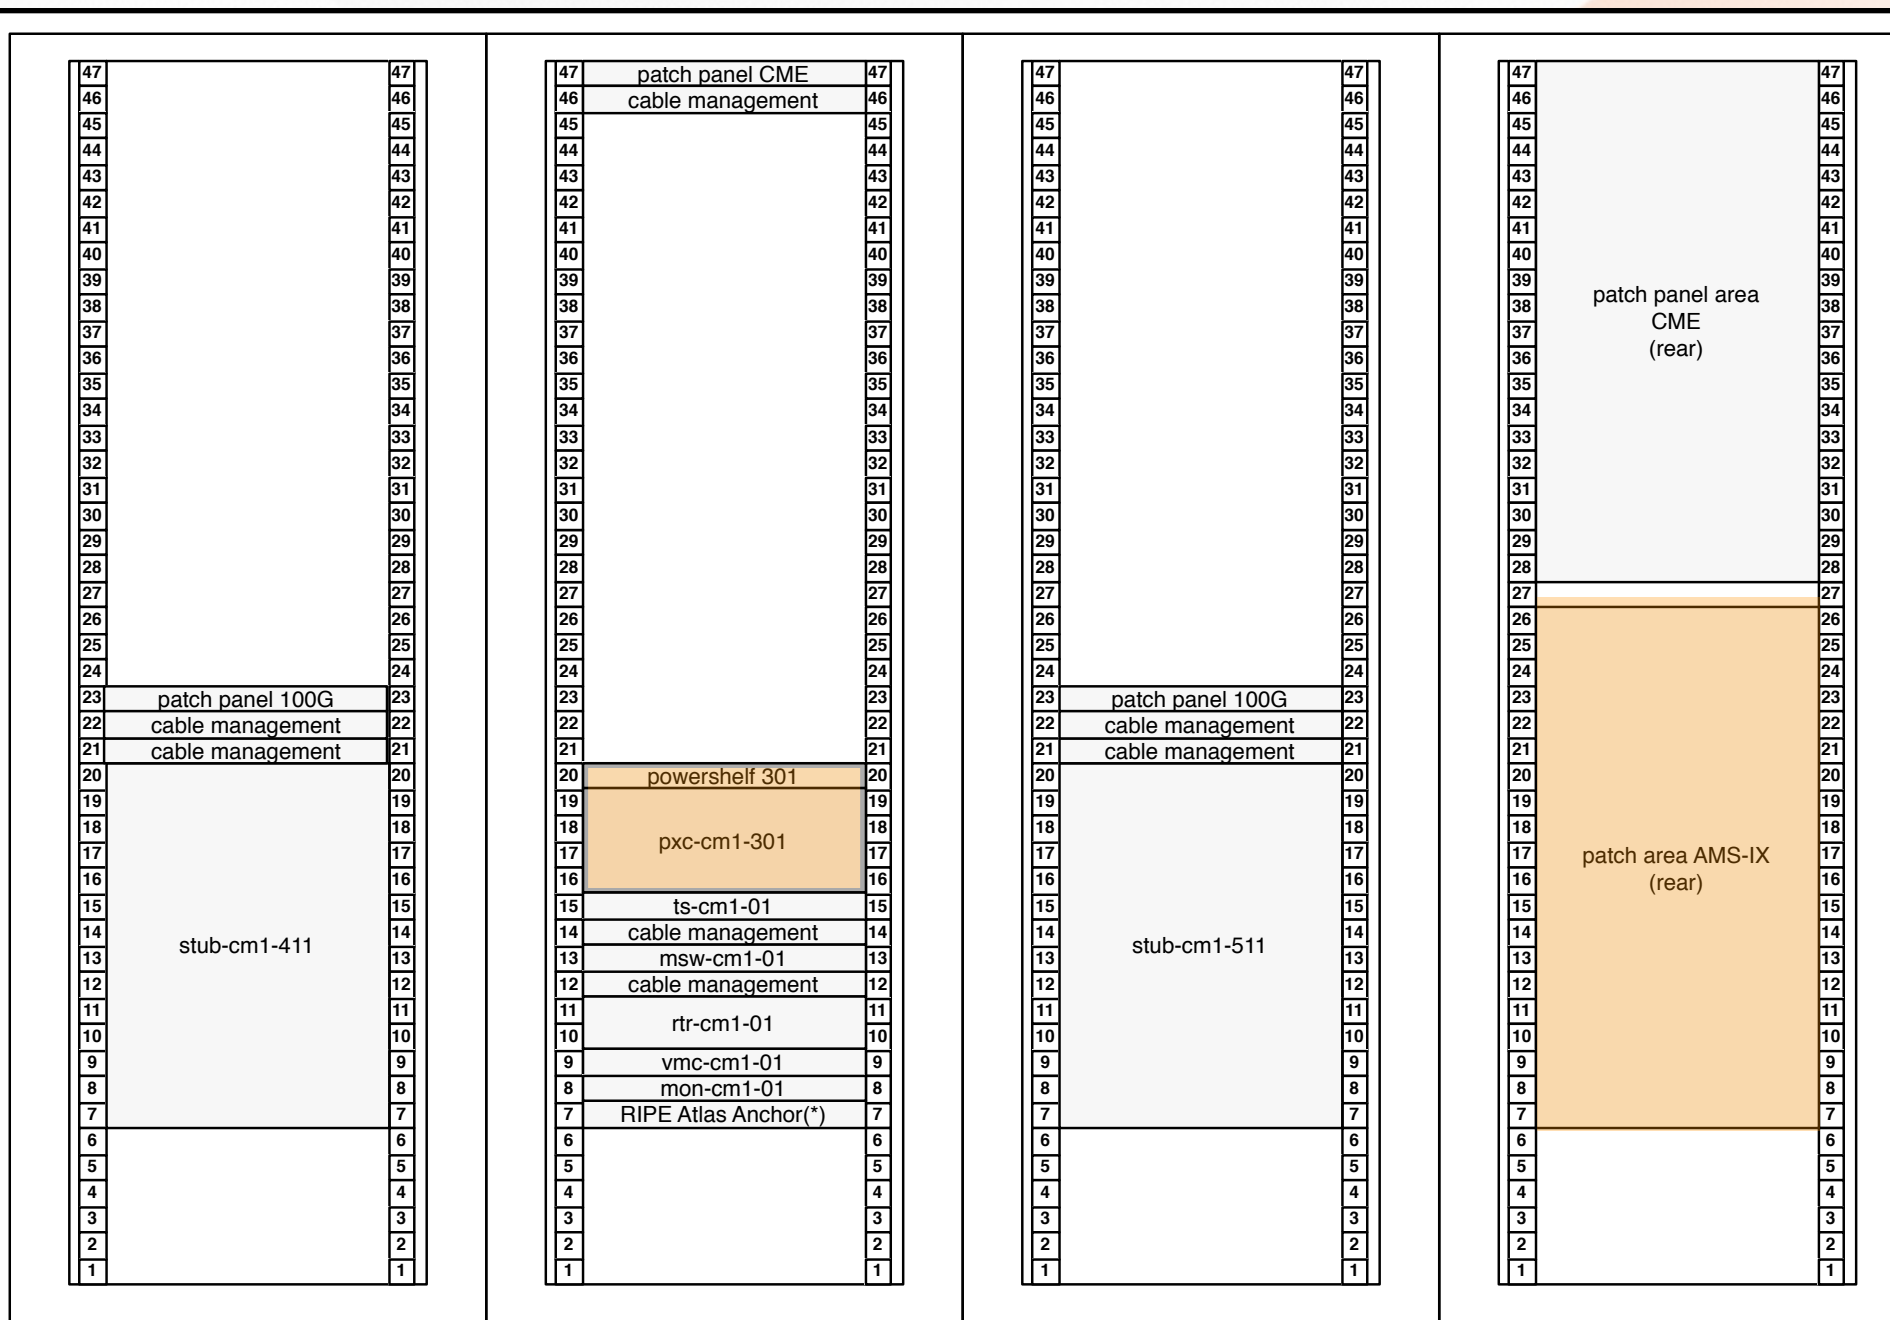
# Scalability



Cold Isle

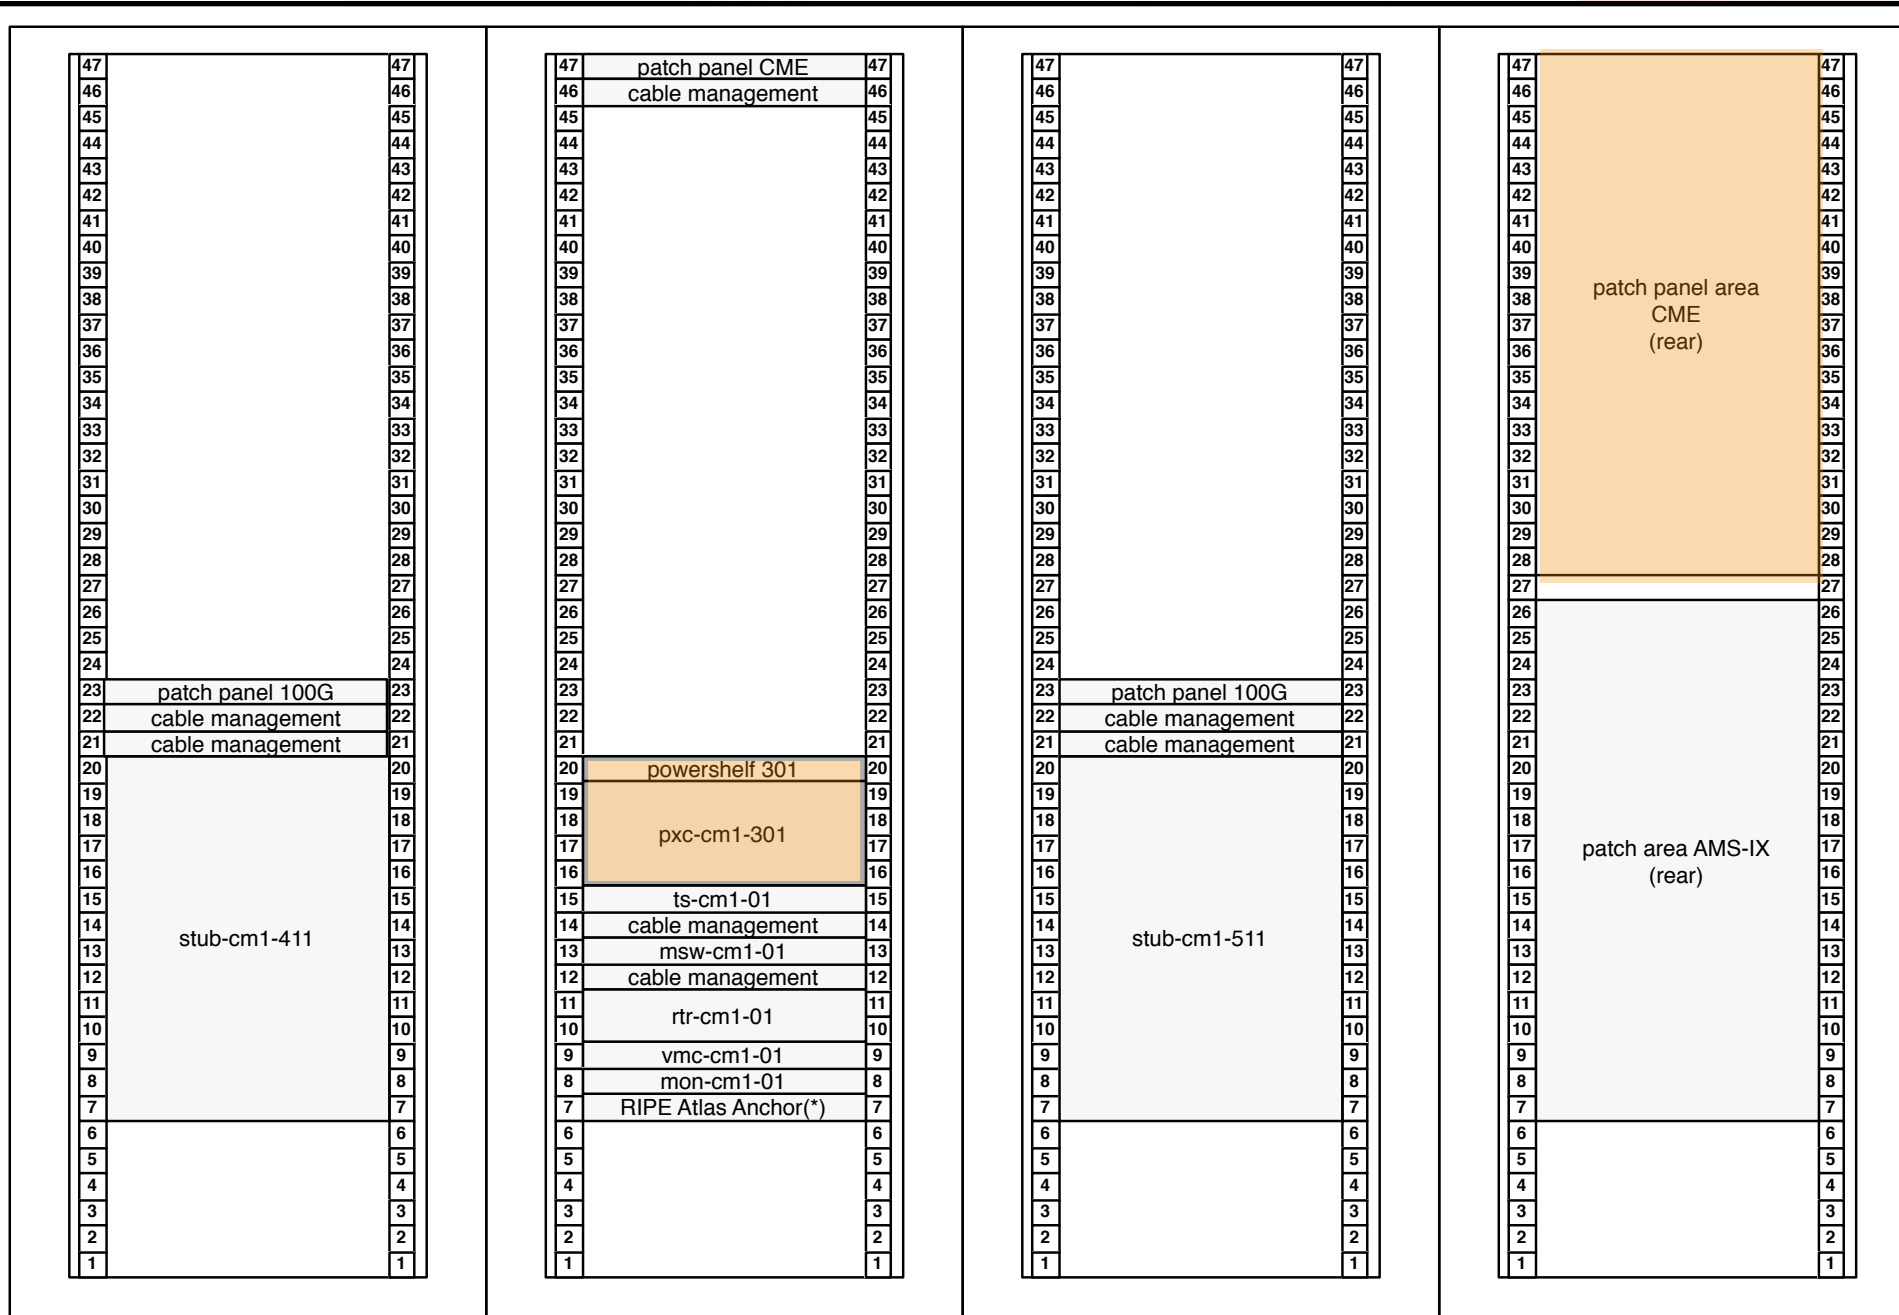


# Scalability



Cold Isle

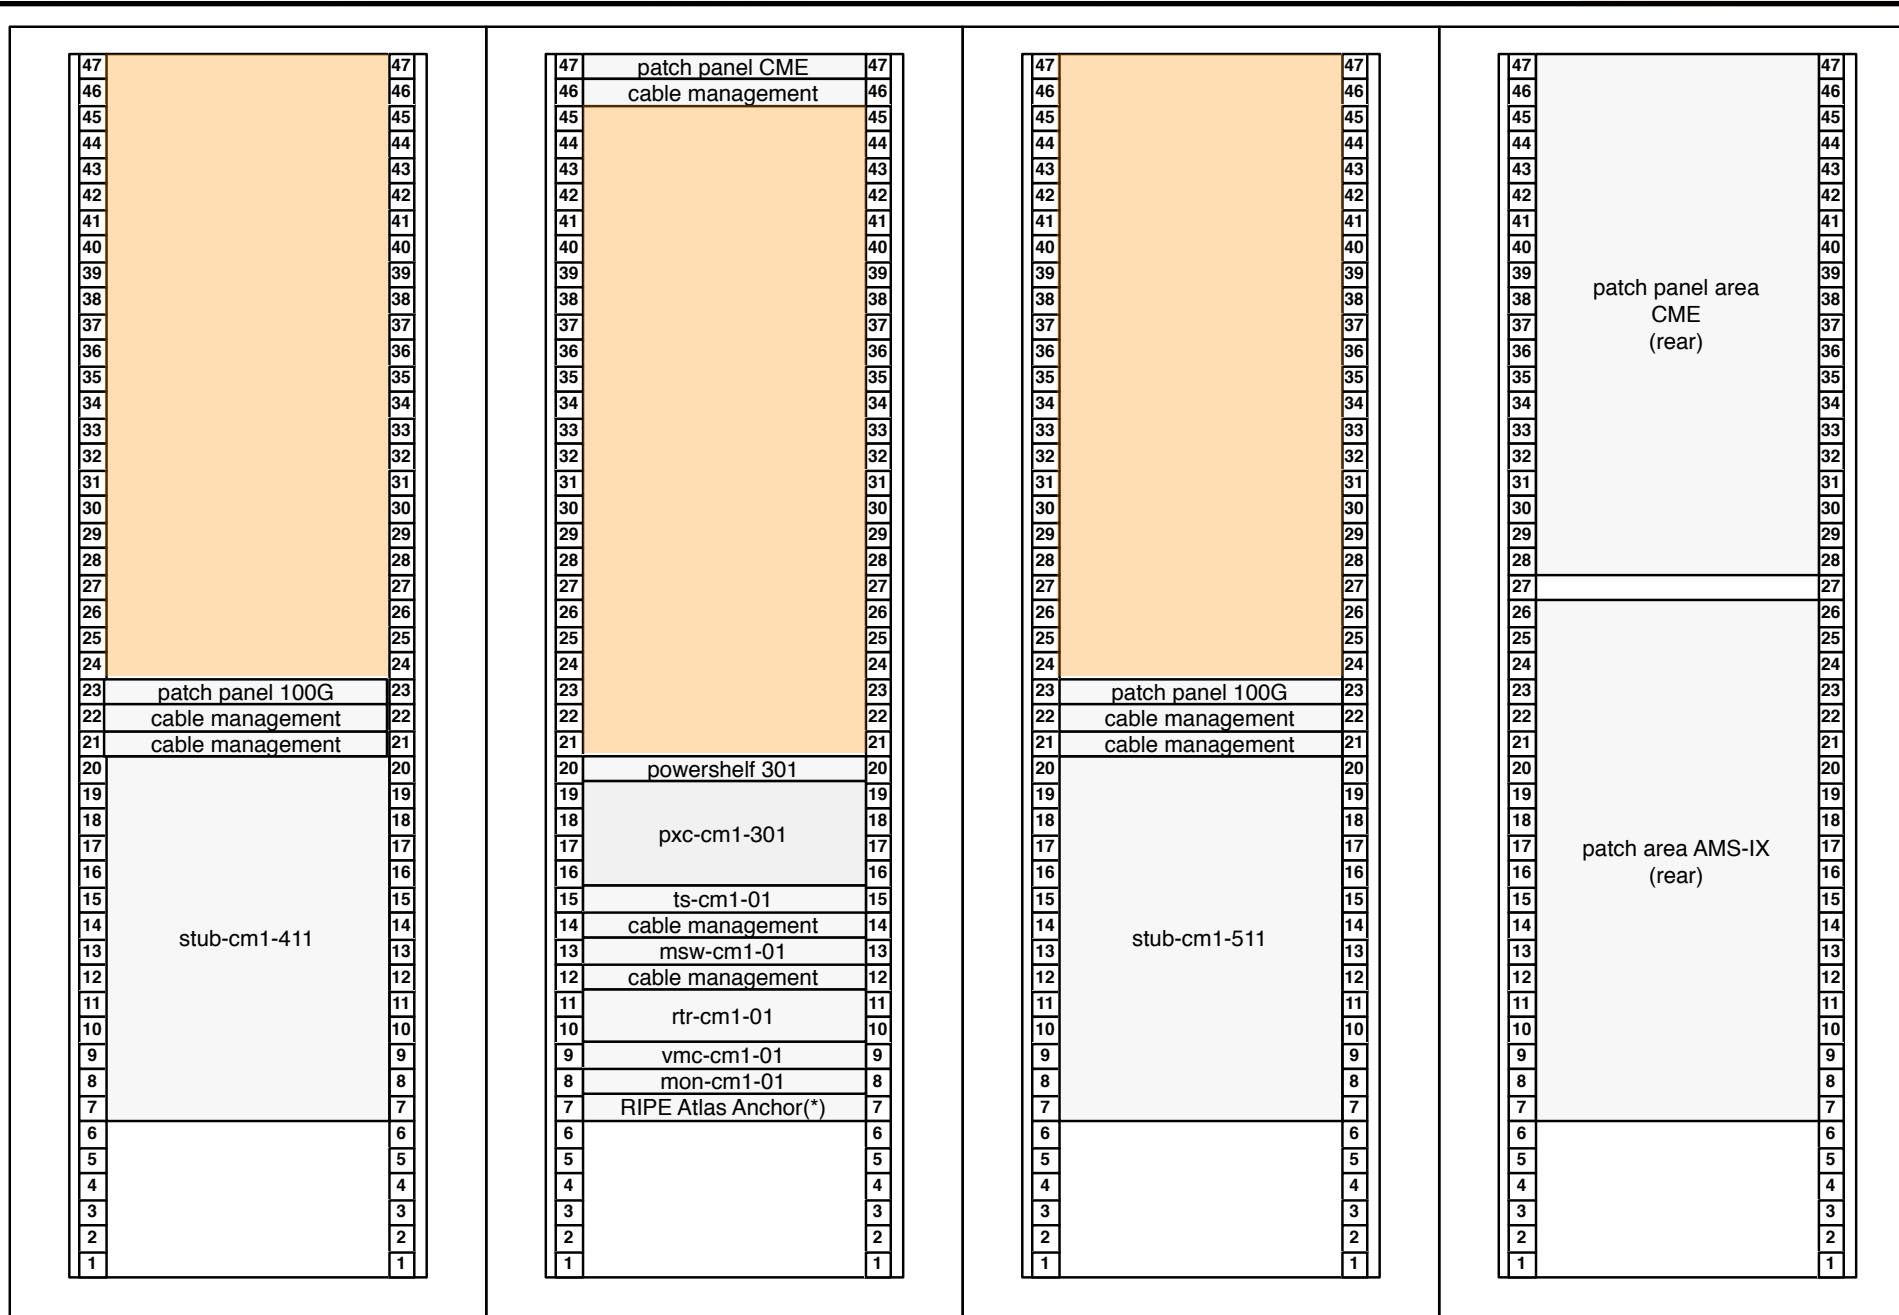
# Scalability



Cold Isle




# Scalability



Cold Isle

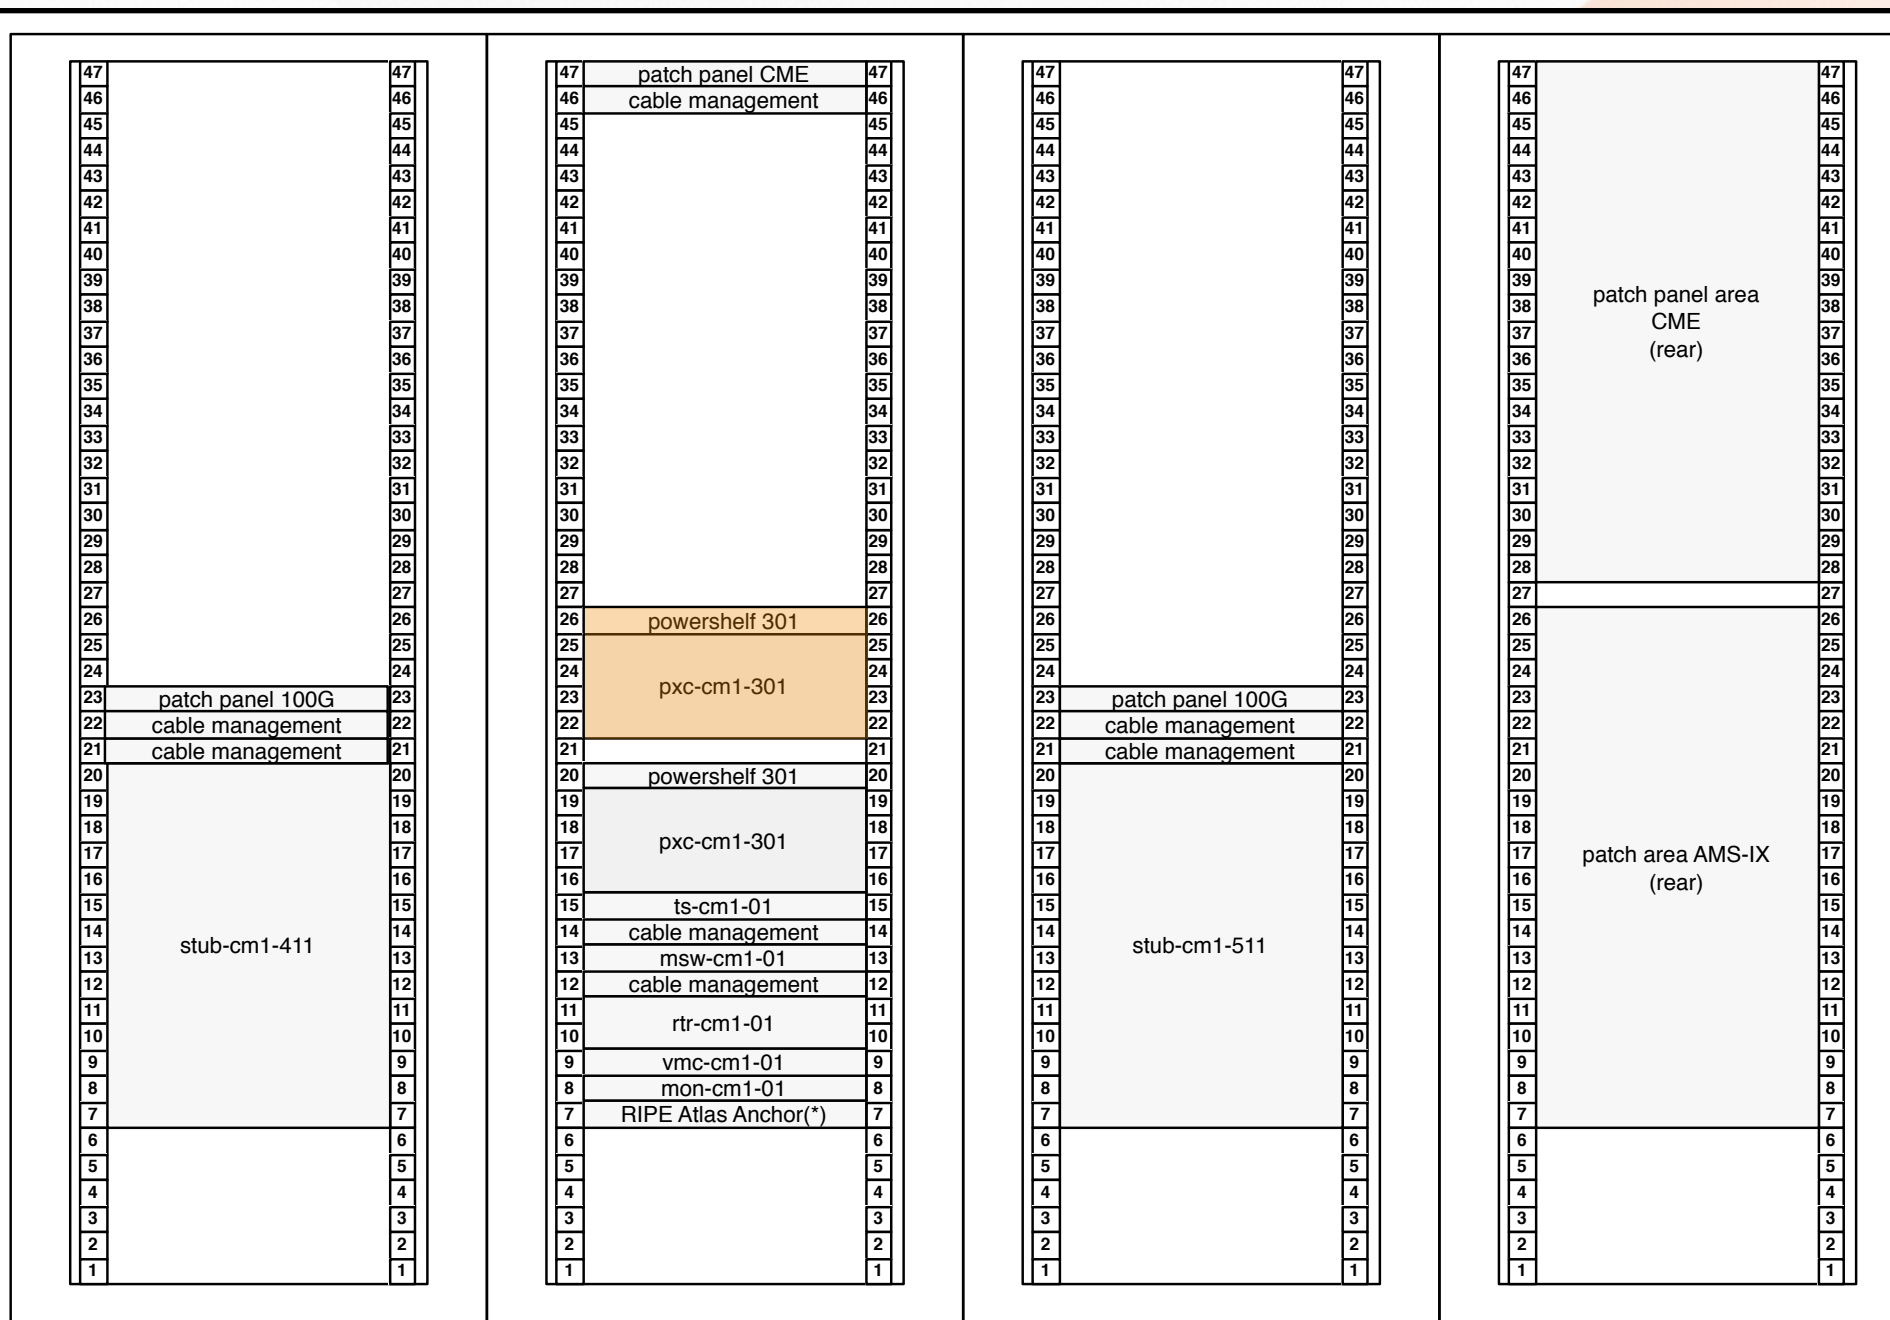
# Scalability



Cold Isle

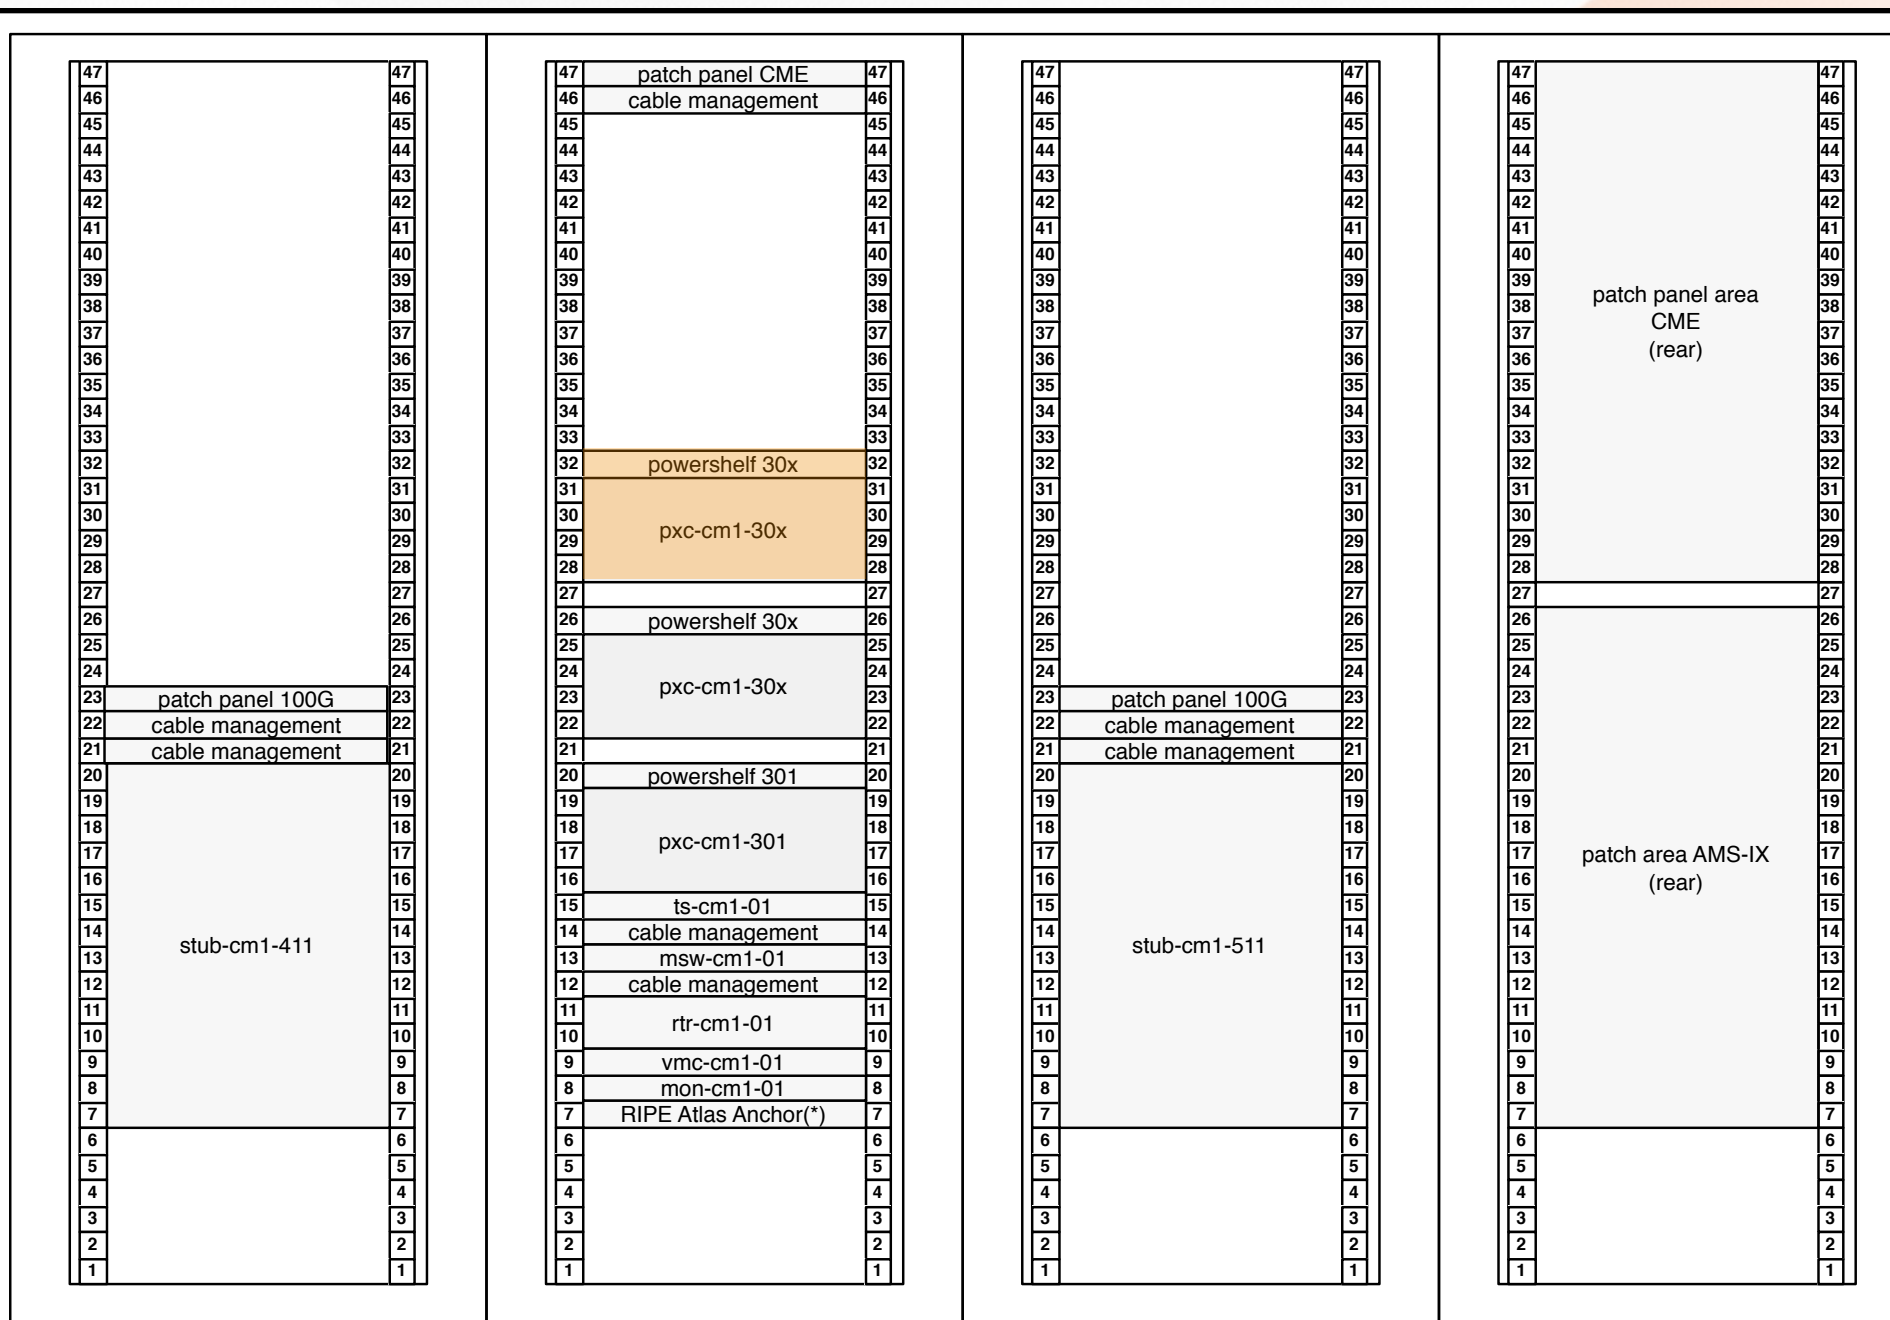


# Scalability



Cold Isle

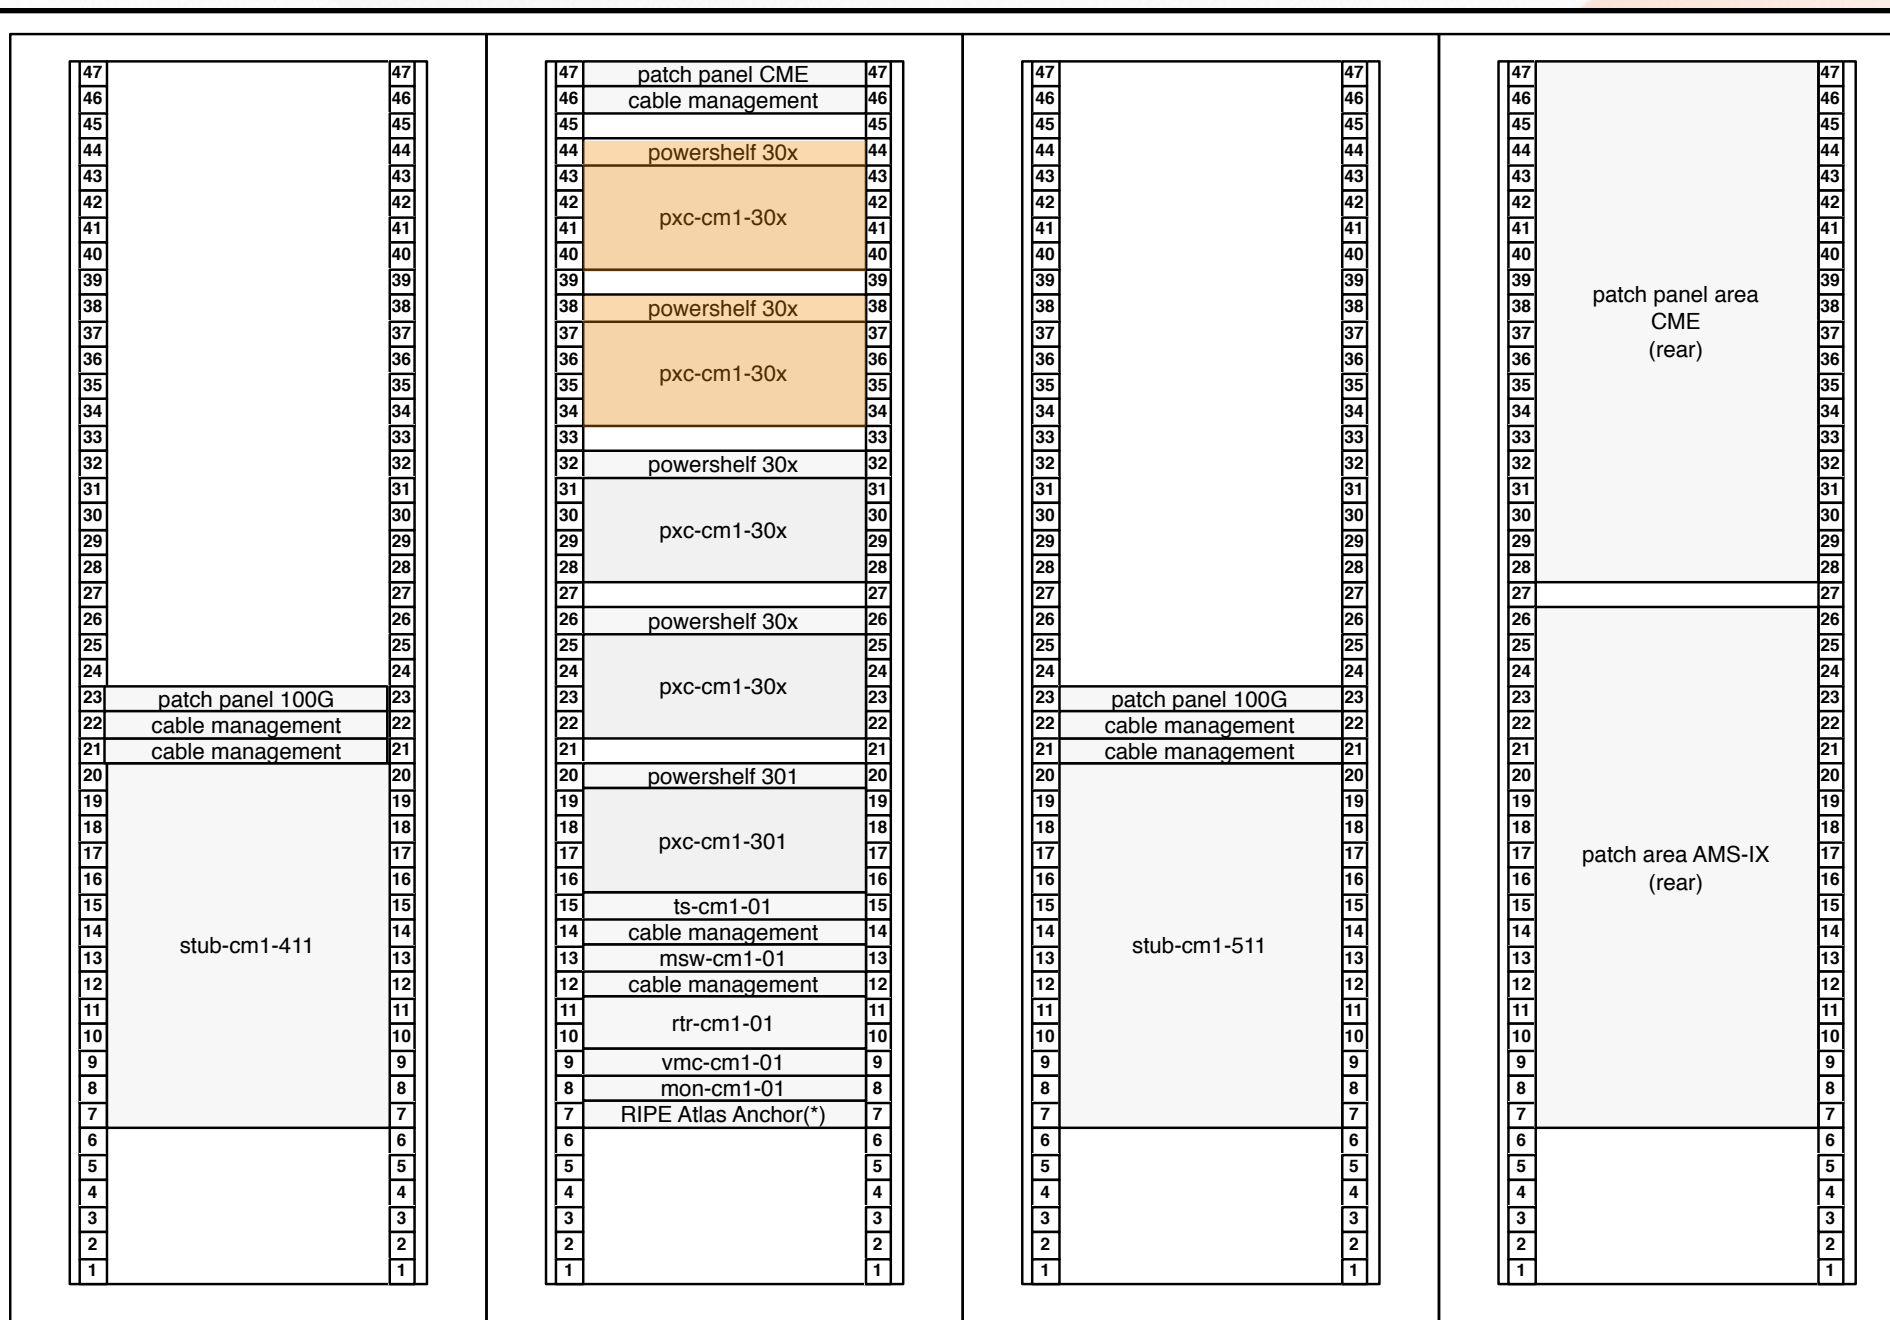
# Scalability



Cold Isle

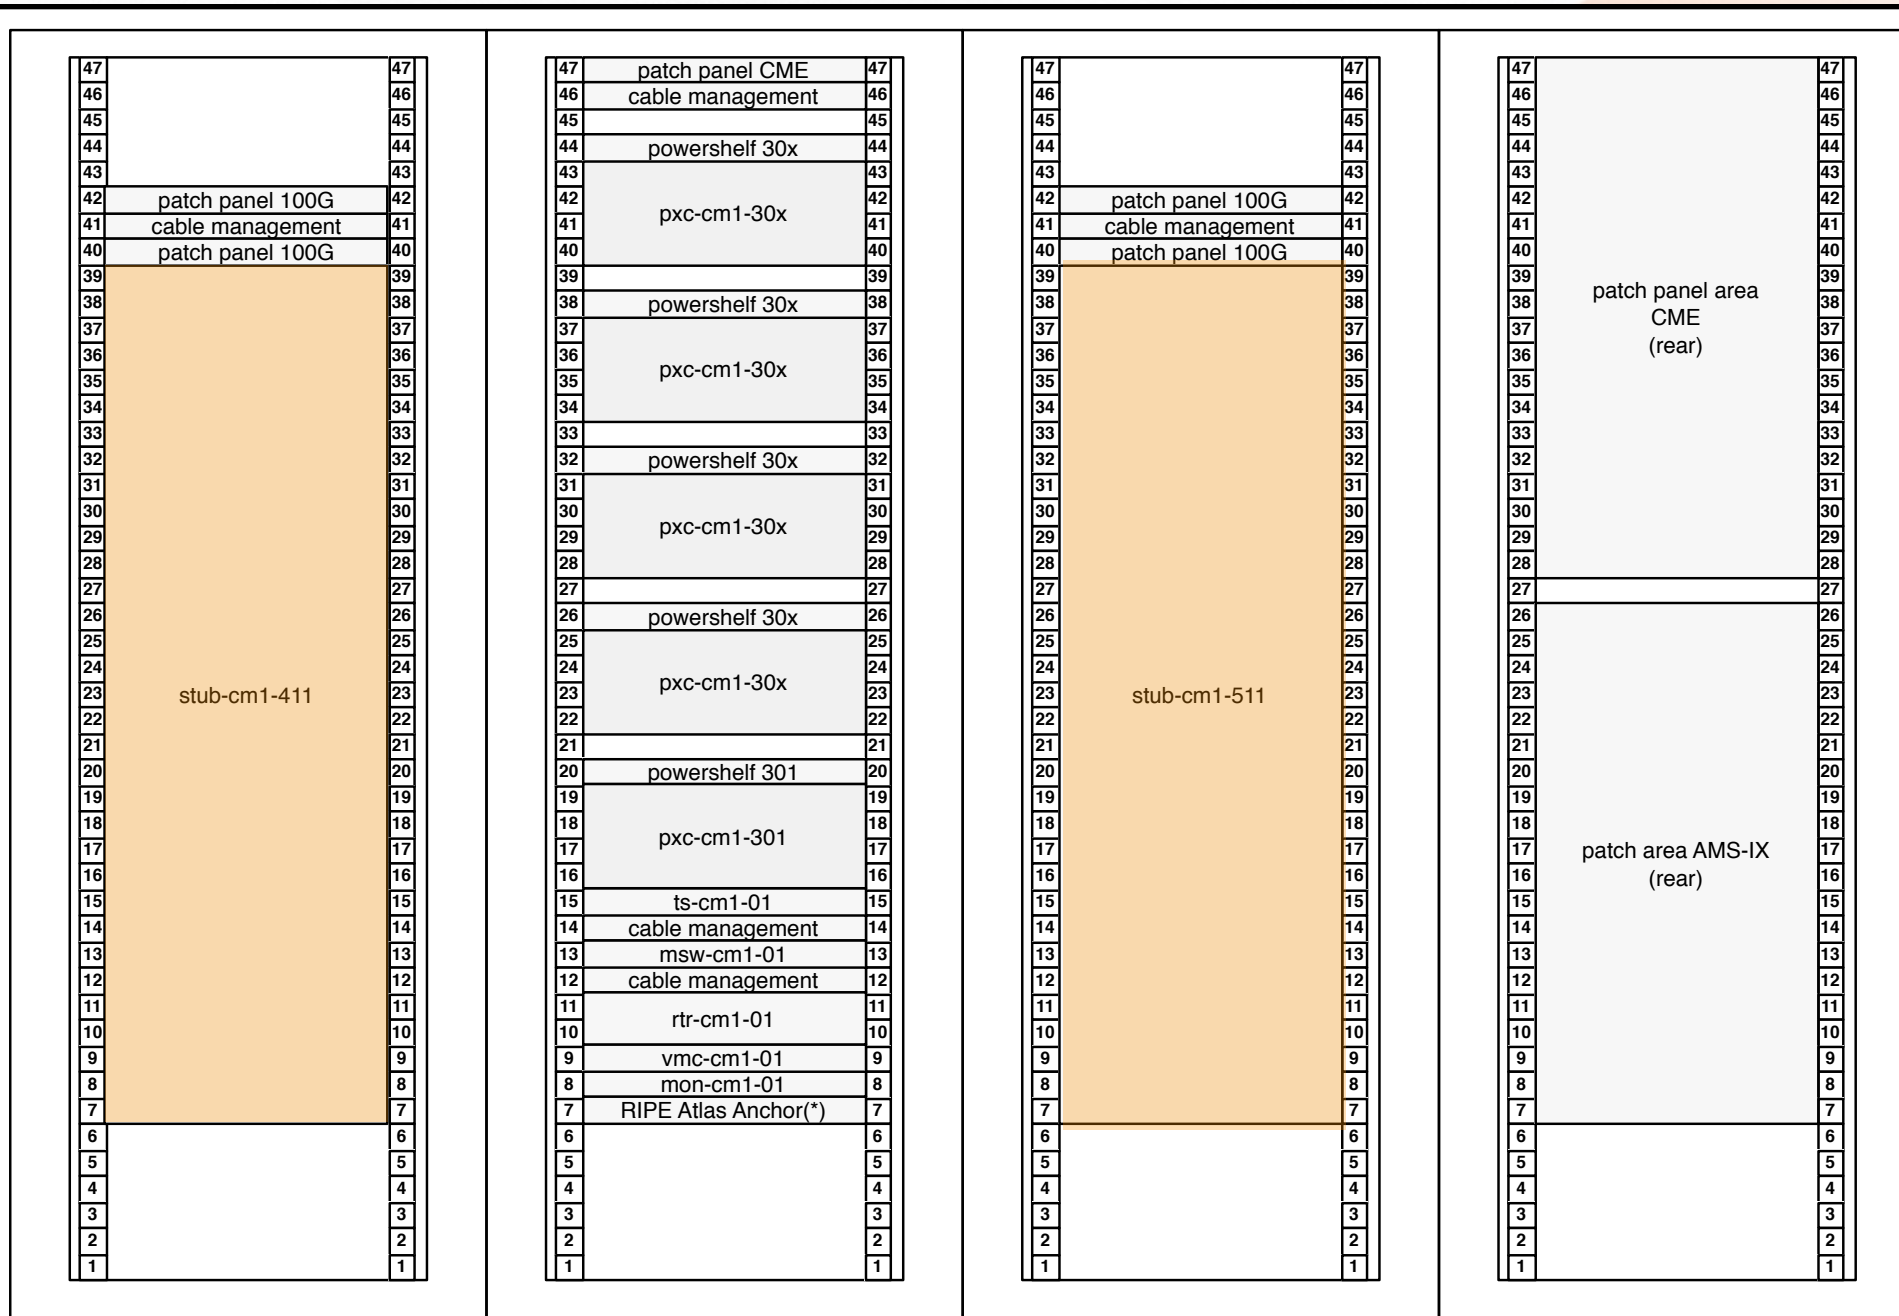


# Scalability



Cold Isle

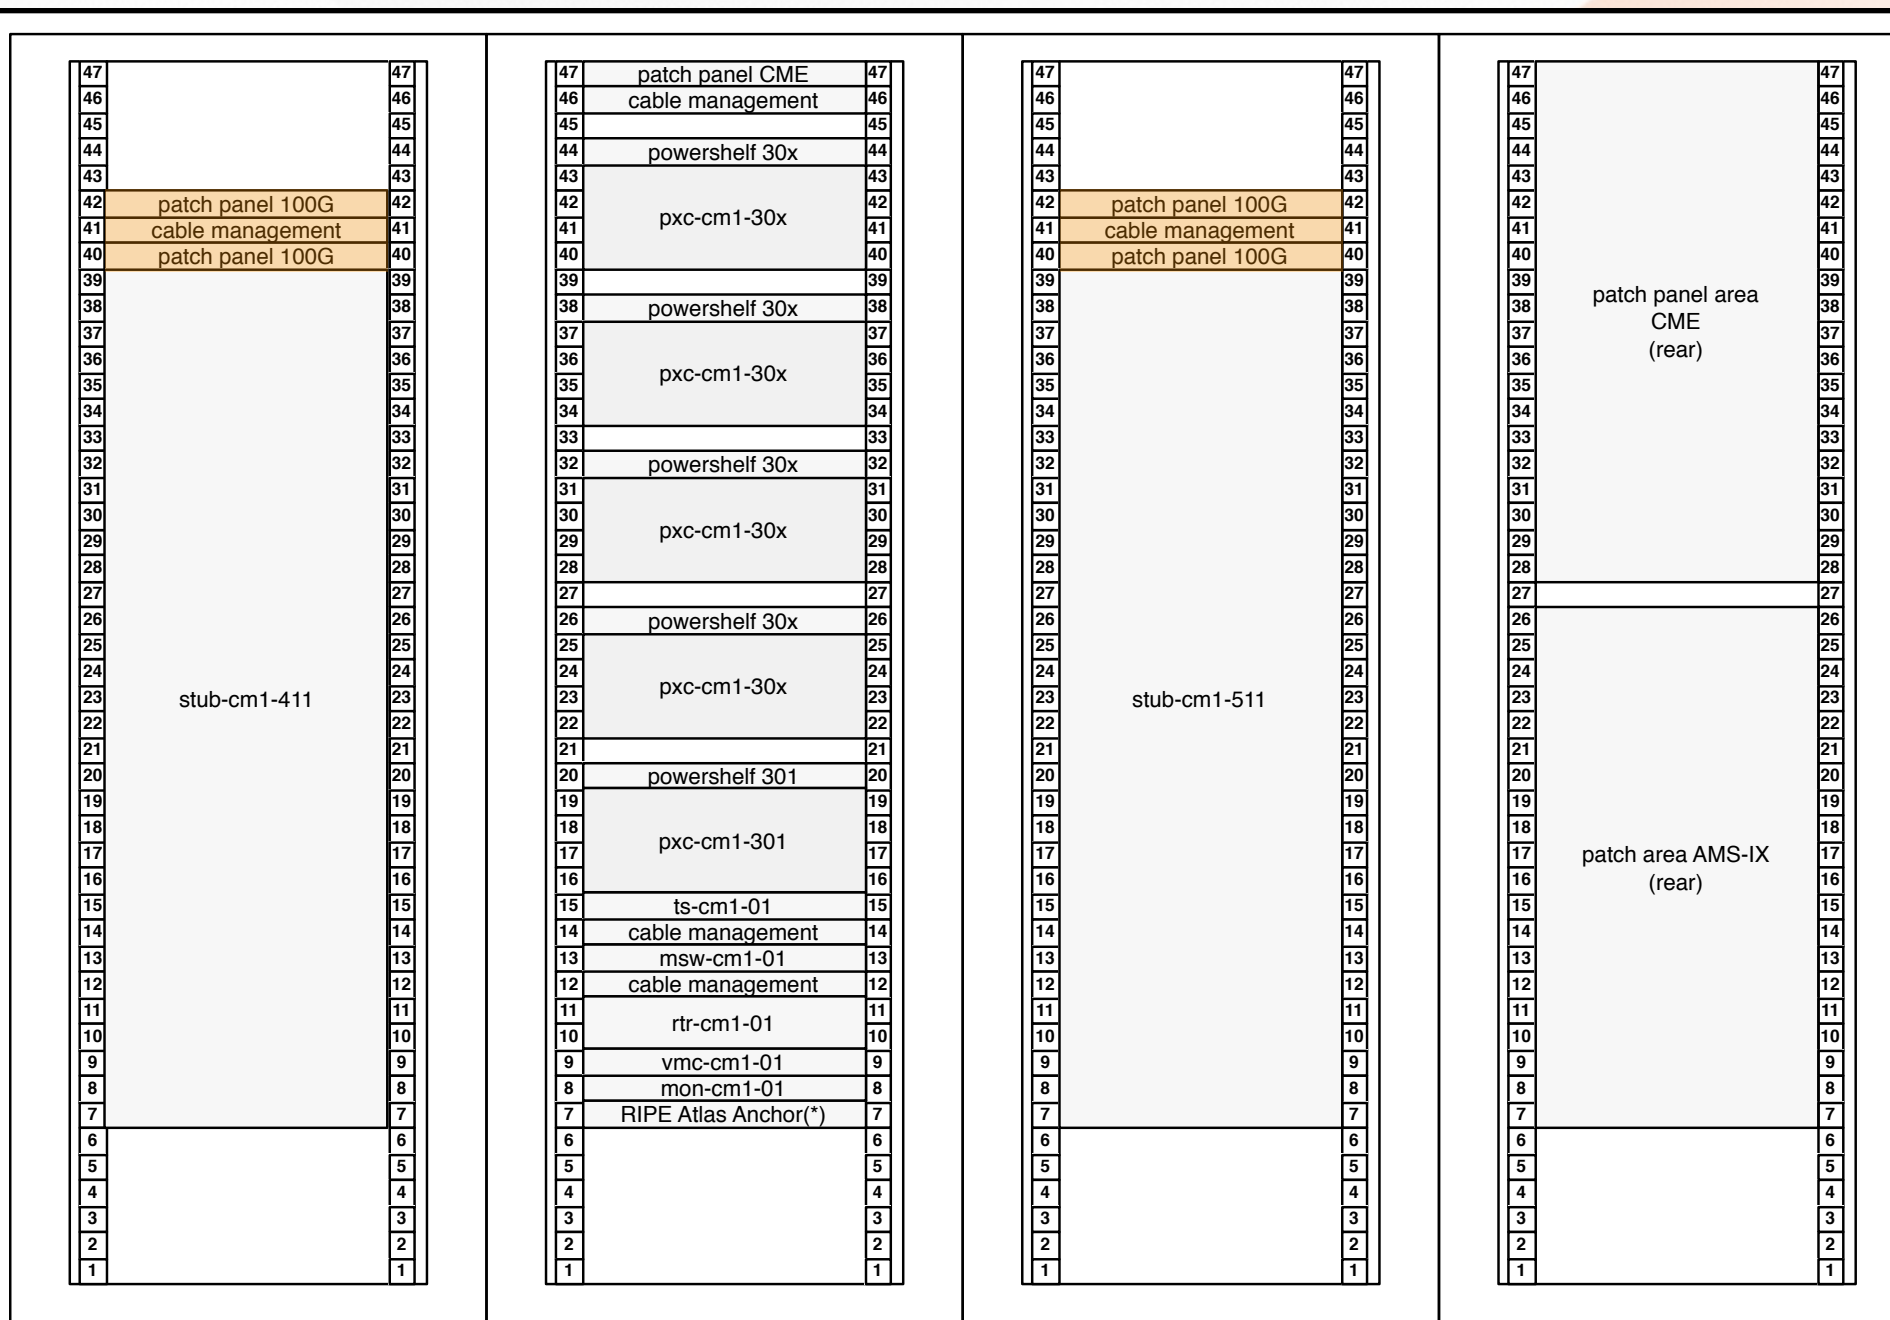
# Scalability



Cold Isle



# Scalability



Cold Isle



**amsix**  
amsterdam  
internet exchange

**End**