

marcel.koch@dd-ix.net

MARCEL KOCH


October 2024

IXP

FROM

RIPE 89

SCRATCH



# What is IXP from Scratch?

## IXP-from-Scratch: Building a New IX



Thomas Liske — 12 Aug 2024

Contributors: [Marcel Koch](#), [Tassilo Tanneberger](#), [Matthias Wählich](#)  
5 min read

[tools](#) [guides](#) [routing](#) [community](#) [competition](#)

152 0



In this new series "IXP from Scratch", we report about our endeavor of building DD-IX, an Internet Exchange located in the city of Dresden, Germany. DD-IX is driven by a grassroots community that loves contributing to a resilient and efficient Internet. We start with some basic background and will share technical design decisions in upcoming articles.

## Believe it or not, the Internet is still a community project

The Internet is a network of networks and Internet Exchange Points (IXP) are a crucial part of the underlying infrastructure. They facilitate public and local interconnections between networks, which increases resiliency because of higher interconnectivity and reduces delays because of vicinity. They are a place where networks meet technically and non-technically.

### More from this author

[View more](#)



#### IXP-from-Scratch: The Peering LAN

Thomas Liske  
30 Sep 2024



#### IXP-from-Scratch: Network and Security Design

Thomas Liske  
26 Aug 2024

# Goals of this RIPE Labs Series




- document how we've build DD-IX
- self-reflection on design choices
- gather feedback from the community
- encourage others to do the same

# Dresden Internet Exchange

2005 past efforts: IX in Dresden

2023 DD-IX e.V. gets formed



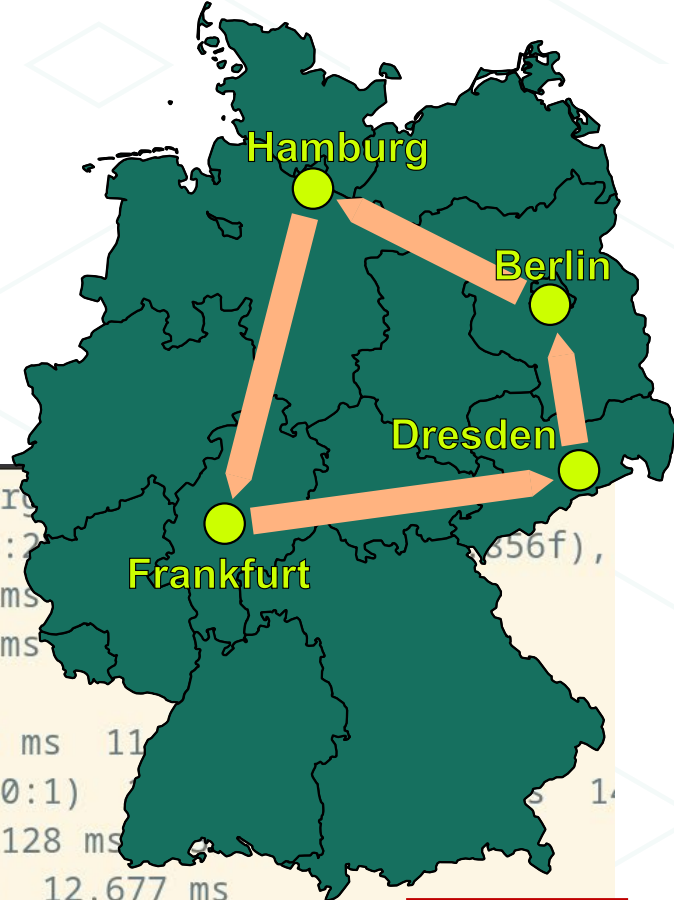


# Dresden Internet Exchange

- 2024 operational
- 2024 multiple  
"Open Tech Meetups"
- 2024 3 peers
- 2024 first Gbits/s peek
- 2024 additional peers awaiting  
for X-Connect




# Before DD-IX



```
[root@micky:~]# traceroute upstream4.core.zentralwerk.org
traceroute to upstream4.core.zentralwerk.org (2a00:8180:2000:37:9e8:1602:c2f3:856f),
 1  dc-c1-cs1-v10.ibh.net (2a01:7700:0:1001::2)  1.679 ms
 2  rtr-c1-bg1.ibh.net (2a01:7700::3)  0.249 ms  0.251 ms
 3  * * *
 4  e0-35.core2.ham1.he.net (2001:470:0:134::2)  11.025 ms  11.025 ms
 5  ipv6.de-cix.fra.de.as16205.dsi.net (2001:7f8::3f4d:0:1)  13.141 ms  13.128 ms
 6  2a00:8180:0:1::4 (2a00:8180:0:1::4)  12.785 ms  12.807 ms  12.677 ms
 7  2a00:8180::18 (2a00:8180::18)  12.785 ms  12.807 ms  12.677 ms
 8  upstream4.core.zentralwerk.org (2a00:8180:2000:37:9e8:1602:c2f3:856f)  12.943 ms
```

```
[root@micky:~]#
```

# Using DD-IX



```
[root@micky:~]# traceroute upstream4.core.zentralwerk.org
traceroute to upstream4.core.zentralwerk.org (2a00:8180:2000:37:9e8:1602:c2f3:856f),
 1  dc-c1-cs1-v10.ibh.net (2a01:7700:0:1001::2)  0.744 ms  0
 2  rtr-c1-bg1.ibh.net (2a01:7700::3)  0.235 ms  0.231 ms  0
 3  dsinet.cc.peer dd-ix.net (2001:7f8:79::3f4d:1)  1.003 ms  0
 4  2a00:8180::18 (2a00:8180::18)  0.623 ms  0.604 ms  0.577 ms
 5  upstream4.core.zentralwerk.org (2a00:8180:2000:37:9e8:1602:c2f3:856f)  0.757 ms
```

```
[root@micky:~]#
```

# Keep Local Traffic Local



# IXP from Scratch

## Published


1. Building a New IX
2. Network and Security Design
3. The Peering LAN

## In progress

- Automation
- Monitoring & Traffic Analysis
- <your input>

## First RIPE Labs article:

<https://labs.ripe.net/author/liske/ixp-from-scratch-building-a-new-ix/>



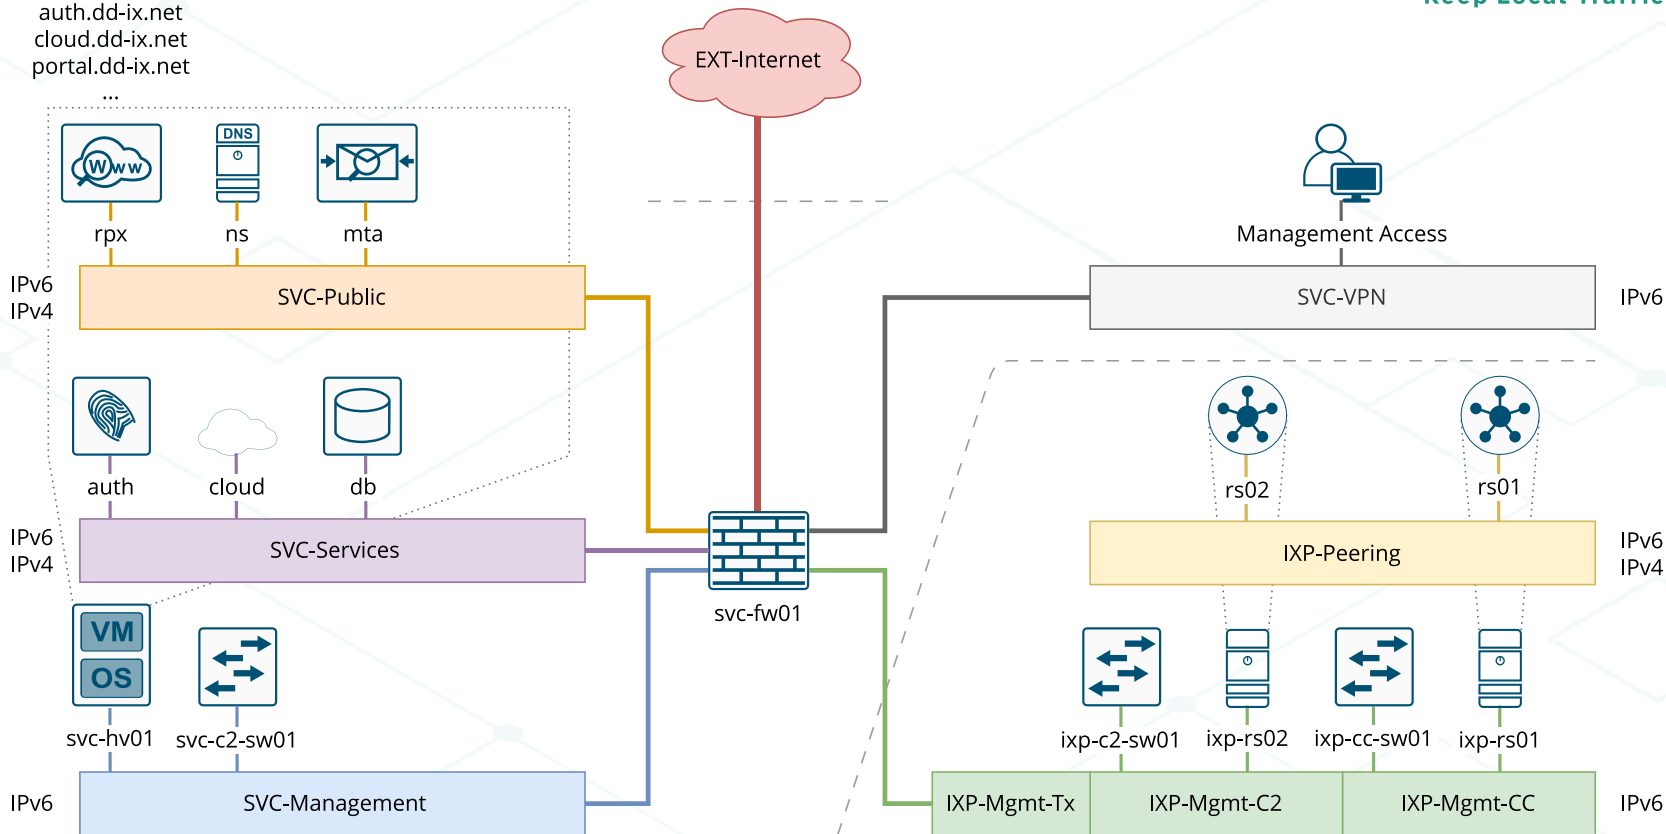
# Network and Security Design

# Services of IXP Operations

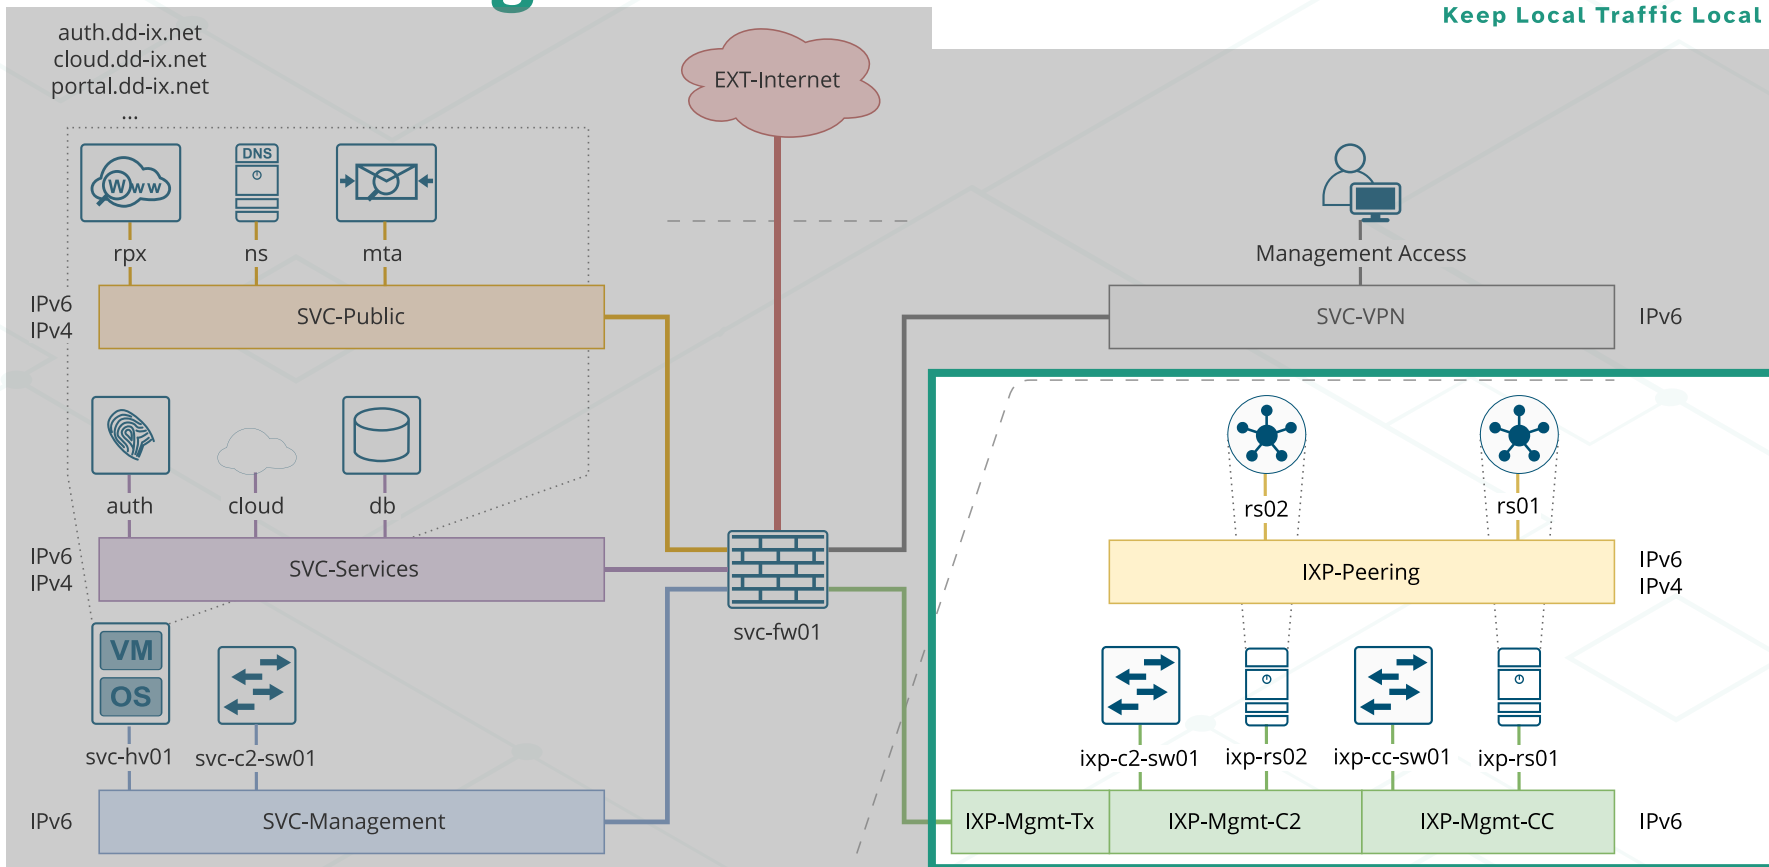
- IXP services
  - peering LAN
  - IXP management
  - monitoring
- office and association
  - paperwork
  - accounting

# Network Segmentation

auth.dd-ix.net  
cloud.dd-ix.net  
portal.dd-ix.net

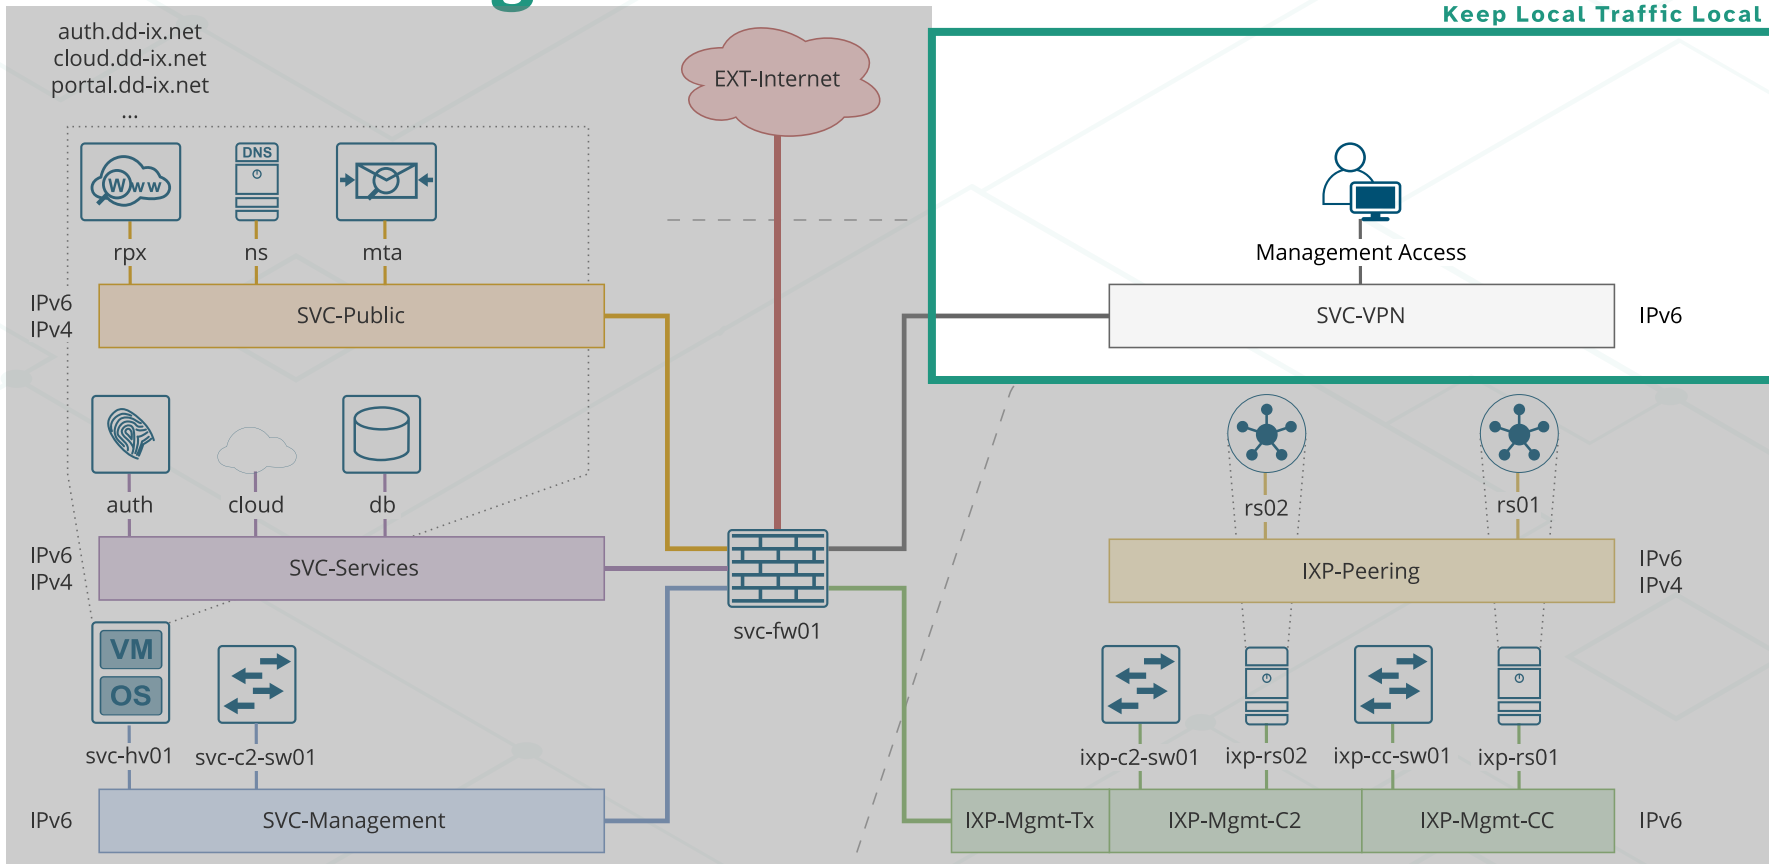


# Network Segmentation

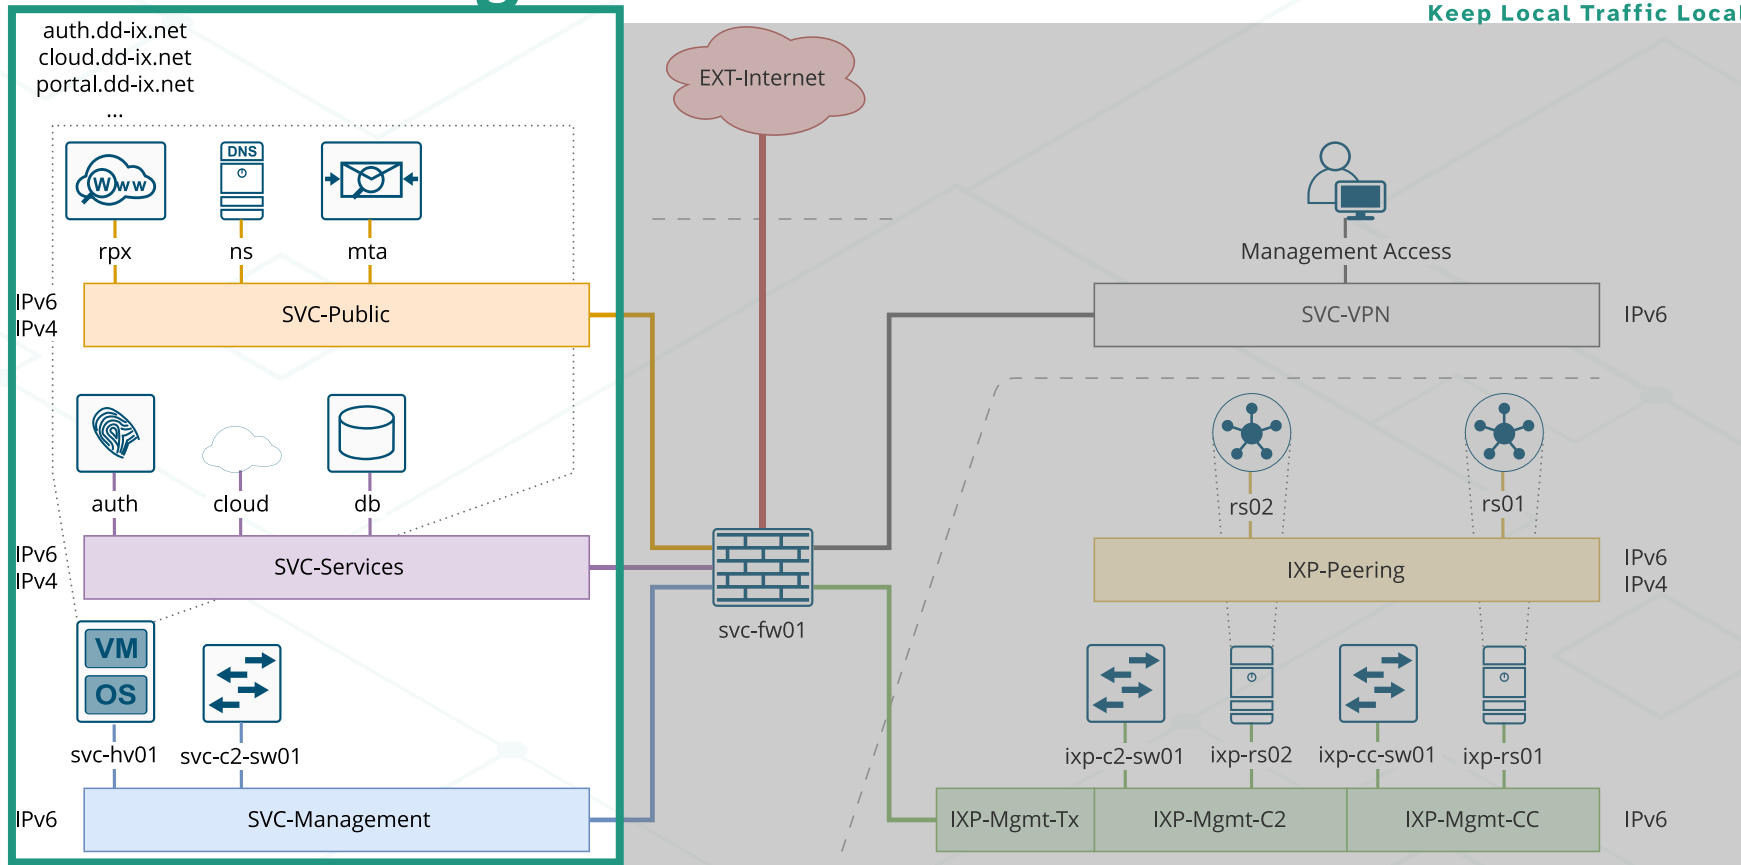




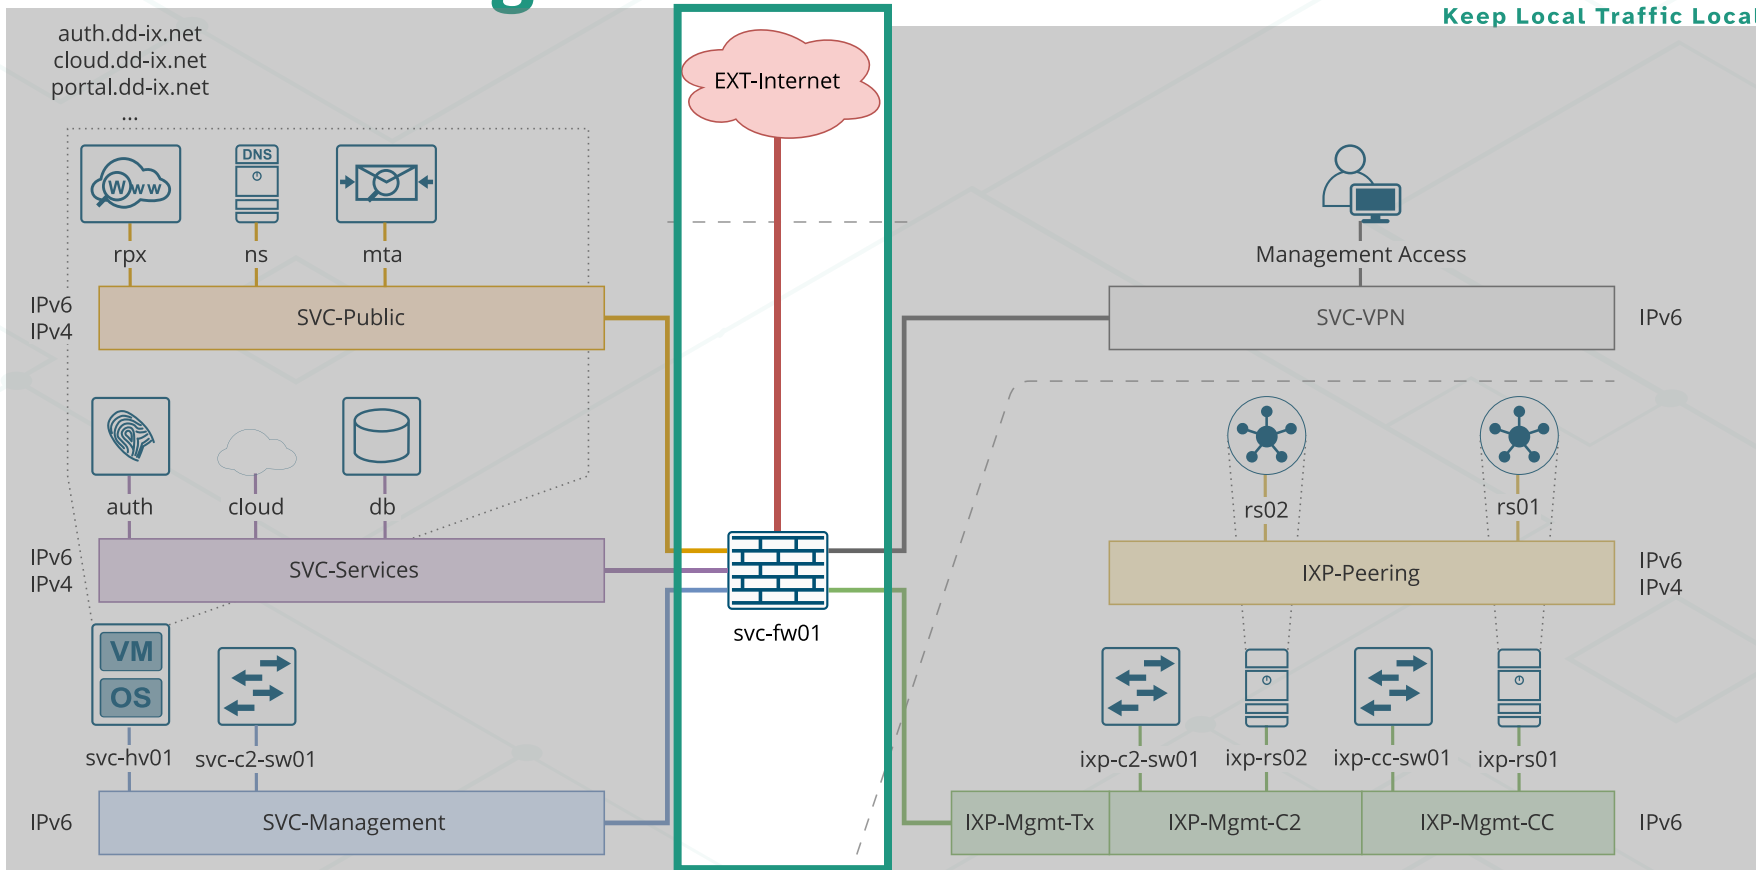
# Network Segmentation



# Network Segmentation



# Network Segmentation



# IPv6-only, we've tried very hard

- ☹️ some external APIs from IXP Manager and arouteserver are IPv4-only
- ☹️ NixOS heavily relies on GitHub, GitHub is missing AAAA records...
- ☹️ some private addresses (RFC1918) remain...

```
> host -t aaaa github.com  
github.com has no AAAA record
```

# Operating Systems on Servers

- **Alpine Linux for routers**
  - runs from tmpfs
  - small overhead
  - Linux routing: Kernel + BIRD
- **NixOS for everything else**
  - declarative and reproducible
  - rollbacks are easy
  - fast if you know how



alpine  
Linux

<https://github.com/dd-ix/ddix-ansible-ixp>  
<https://github.com/dd-ix/ddix-ansible-rs>




NixOS

<https://github.com/dd-ix/nix-config>



# Service Isolation

- currently, only single management host
- *MicroVM*'s for every service
  - KVM-based virtualization
  - diskless/tmpfs
  - mount points for persistent data
  - well integrated into NixOS thanks to @astro



# Any Questions



**DD-IX**  
Keep Local Traffic Local



[dd-ix.net](https://dd-ix.net)  
[contact@dd-ix.net](mailto:contact@dd-ix.net)

## Our Sustaining Members

**IBH**  
IT-SERVICE

 Sachsen  
GigaBit

**DSI**

## Our Patrons

**BCIX**

**I/P/B/**<sup>®</sup>  
Internet Provider Berlin

 TECHNISCHE  
UNIVERSITÄT  
DRESDEN

# Nix Store with MicroVMs

