

# SDN IXP

**Marc Bruyere**

The University of Tokyo



東京大学  
THE UNIVERSITY OF TOKYO

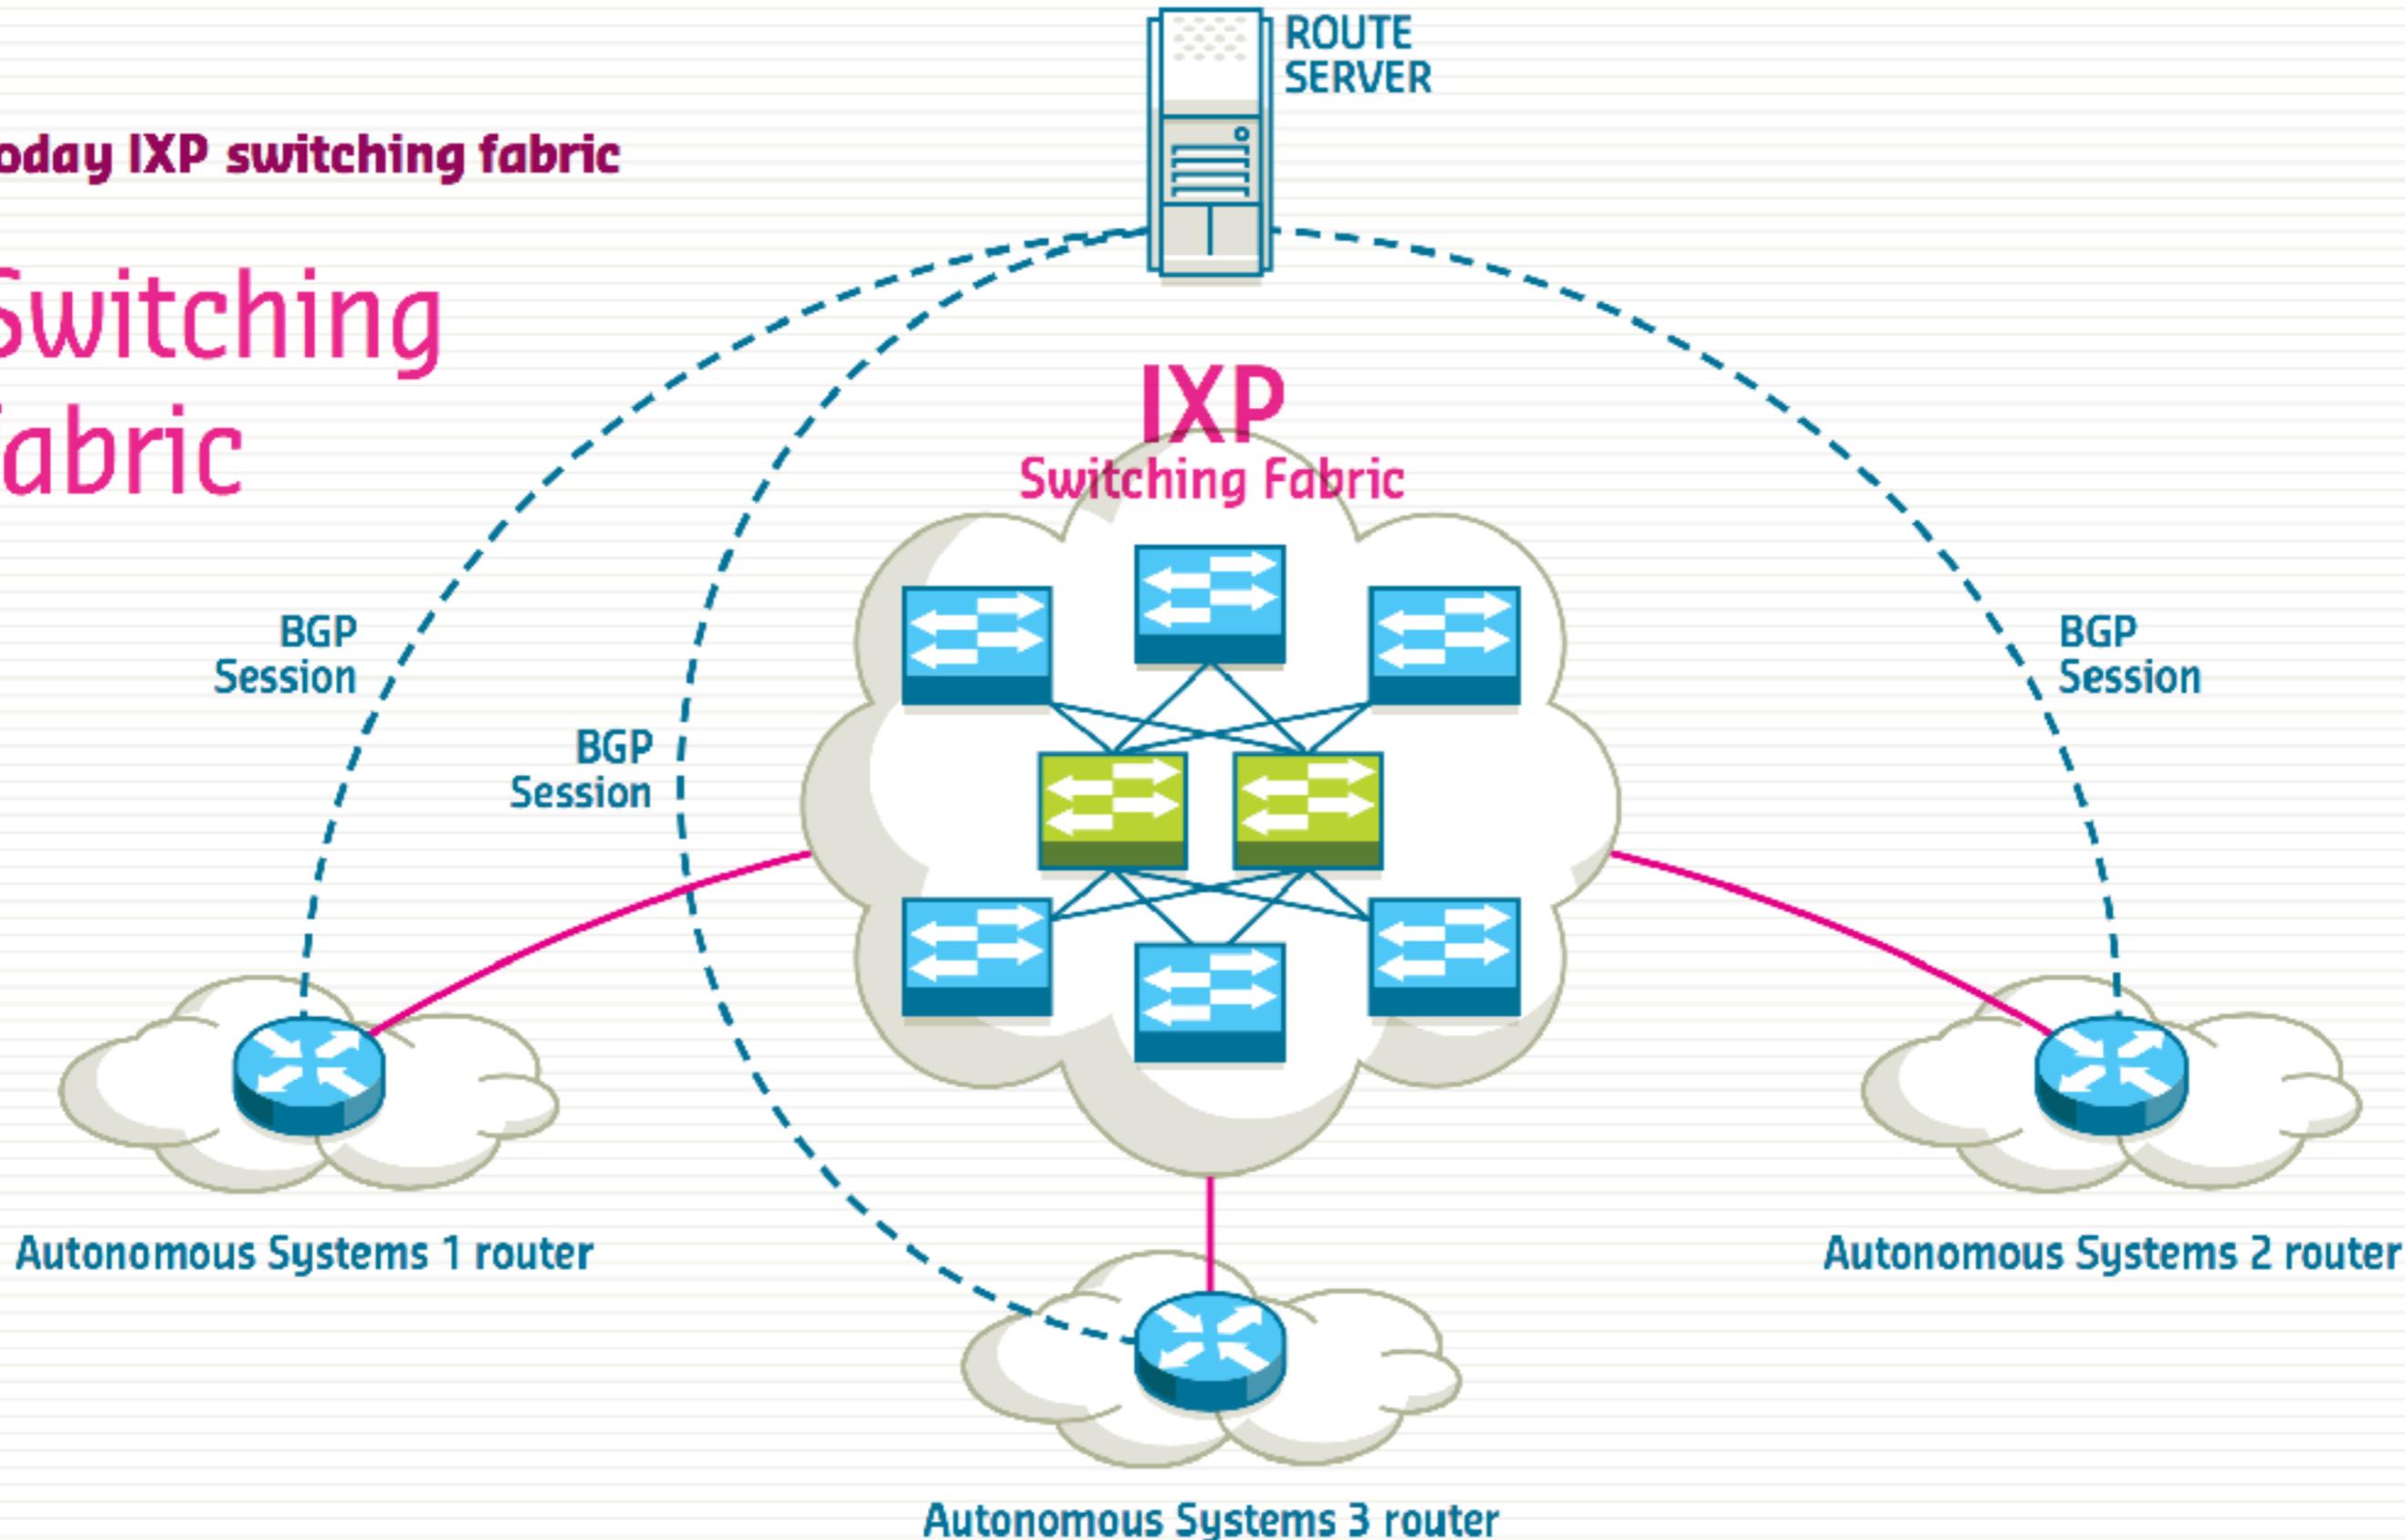


# Agenda

- SDN momentum for IXPs - Umbrella
- Toulouse IXP - TouIX to TouSIX
- Tokyo IXP - DIX-IE to PIX-IE
- Osaka - NSPIXP-3 to FAUCET Umbrella

Today IXP switching fabric

Switching fabric



# Issues with today IXP switching fabric

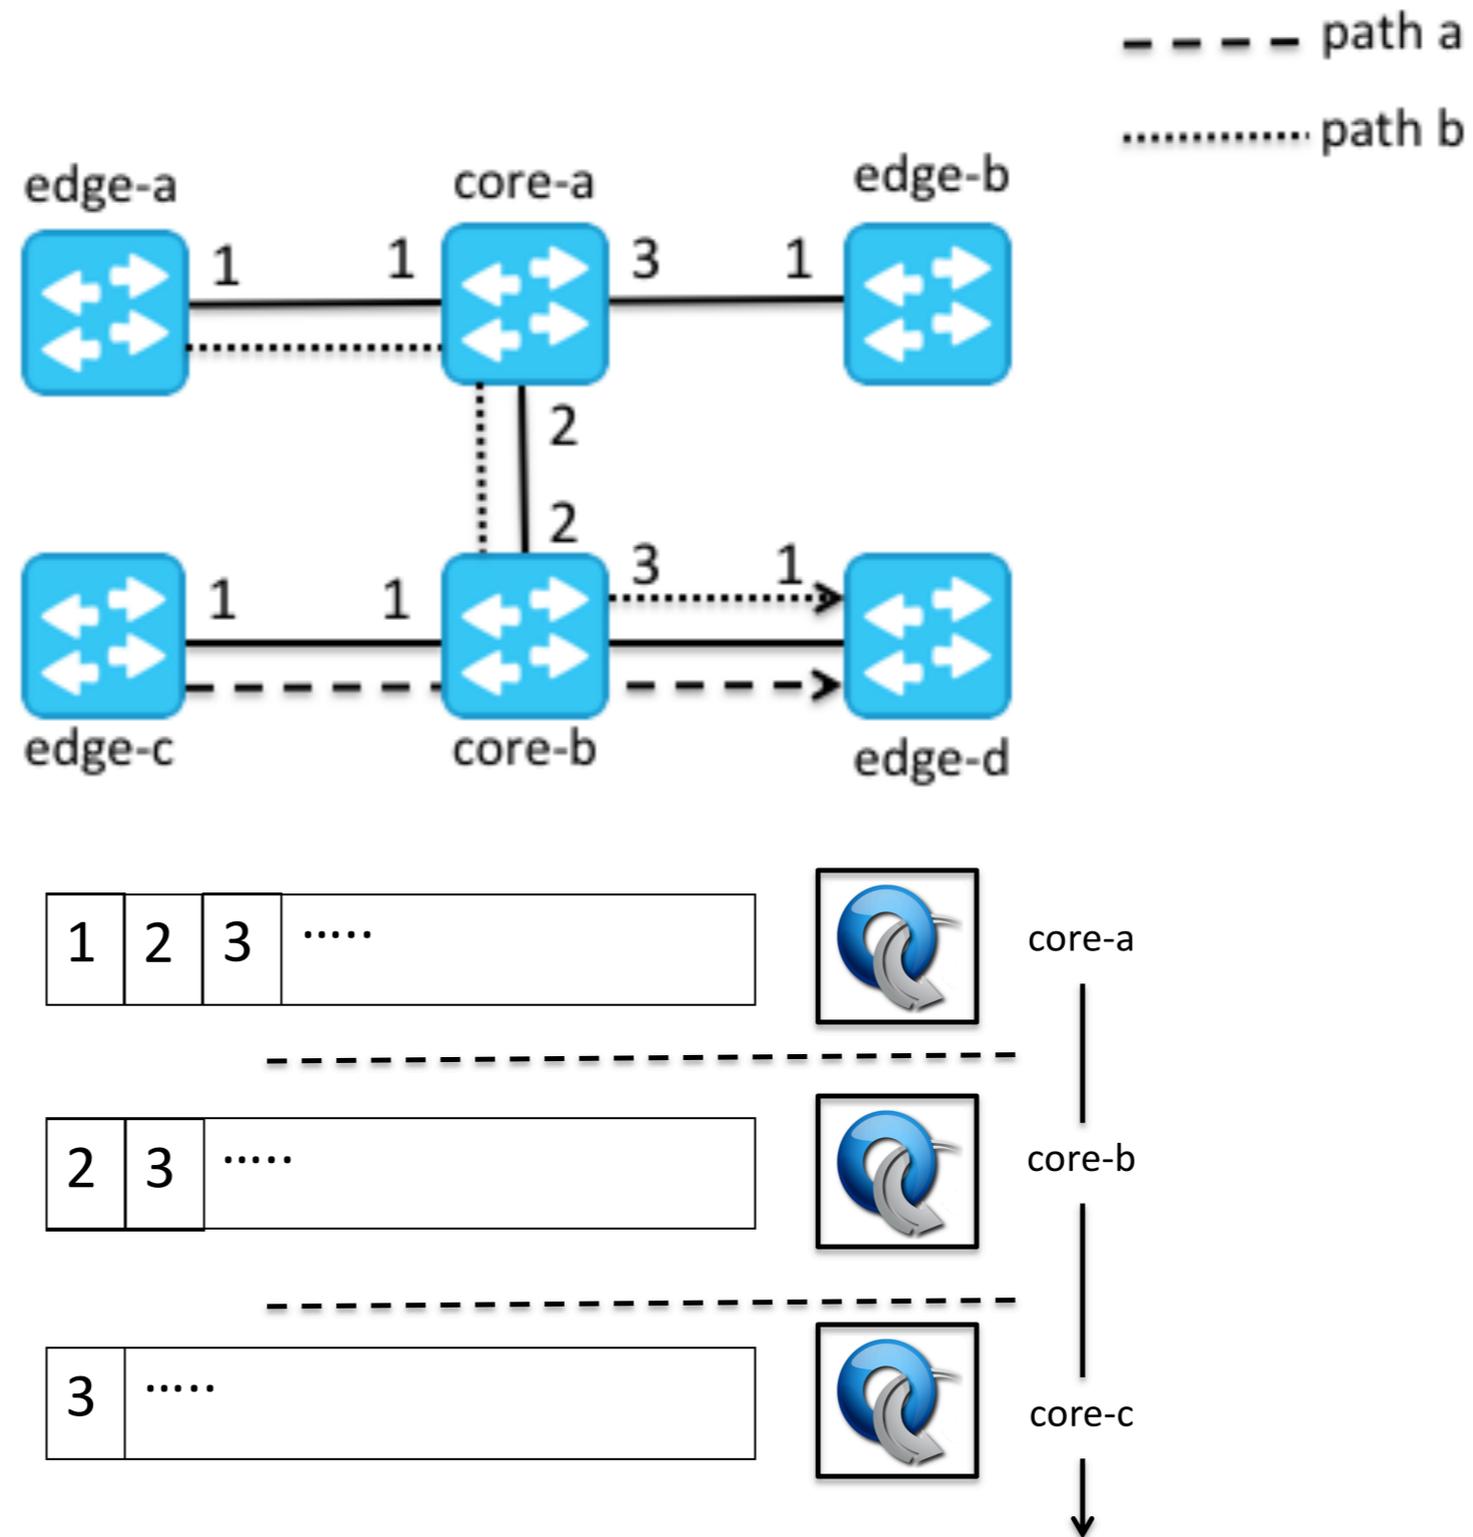
*IXP switching fabric are shared Layer 2 broadcast domain*

- Broadcast traffic can weaken router CPU or even neutralize the entire IXP
- Loop Free solutions are not perfect
- Hard to scale up
- Undesired traffic are hard to be kept out
- Monitoring is too limited or too complex

# Umbrella architecture

- No more Broadcast and perfect edge filtering
- Pseudo Wire
- Can run even if the control plane is down
- Works even without OpenFlow switch in the core
- Fined-grained monitoring with OpenFlow
- Scalable for more PoPs and IXPs Members
- Open to future applications Oriented IXP Customer

# Umbrella architecture

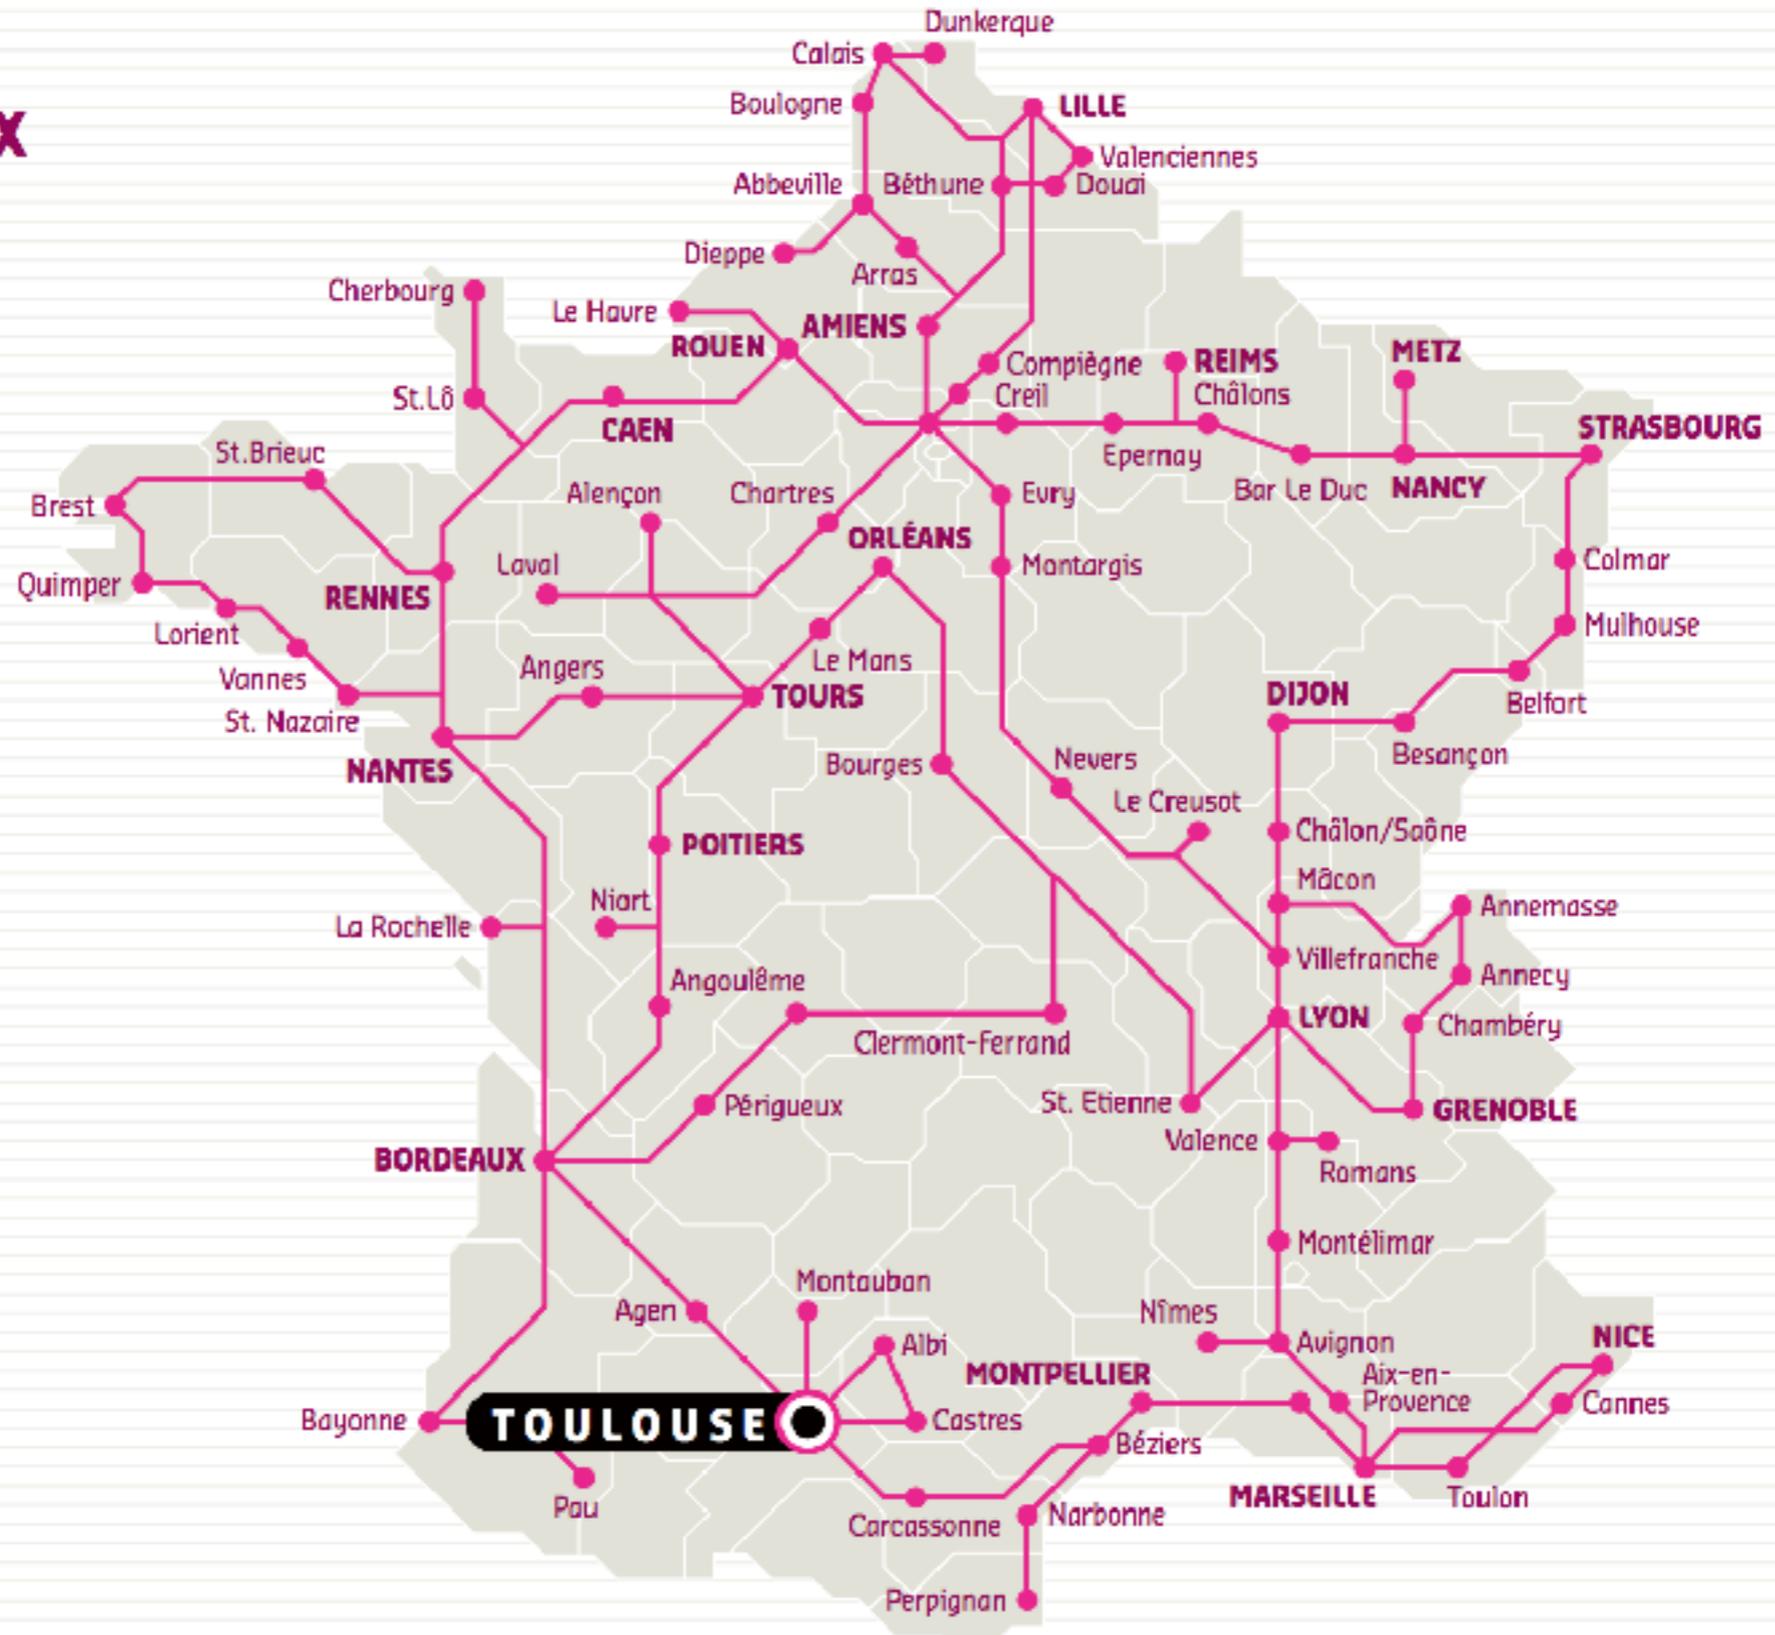




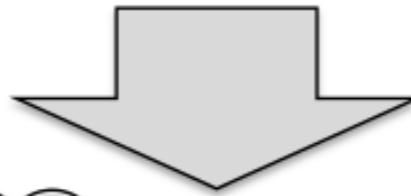
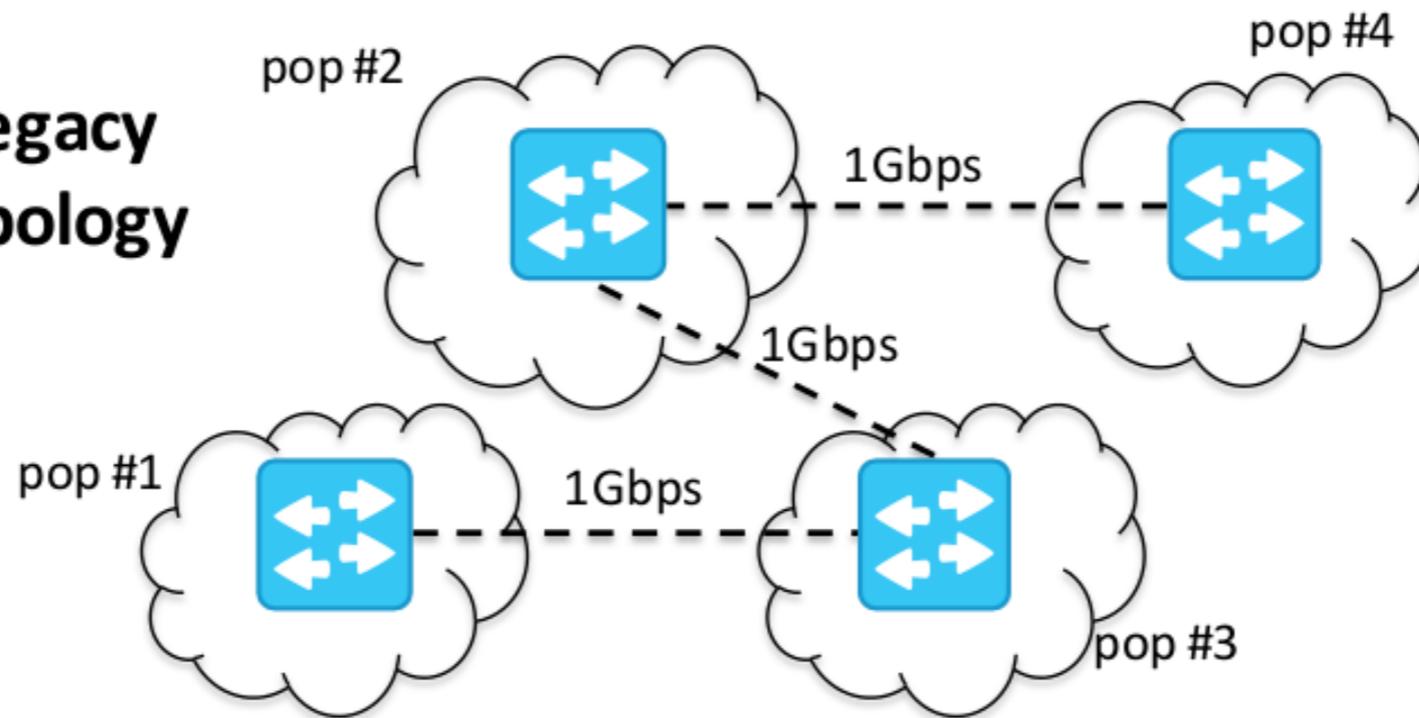
- Founded in 2006
- TouIX is a EURO-IX member
- 4 PoPs
- 10 active members
- Interconnected to France-IX and Lyon-IX

## The Toulouse IXP : ToulIX

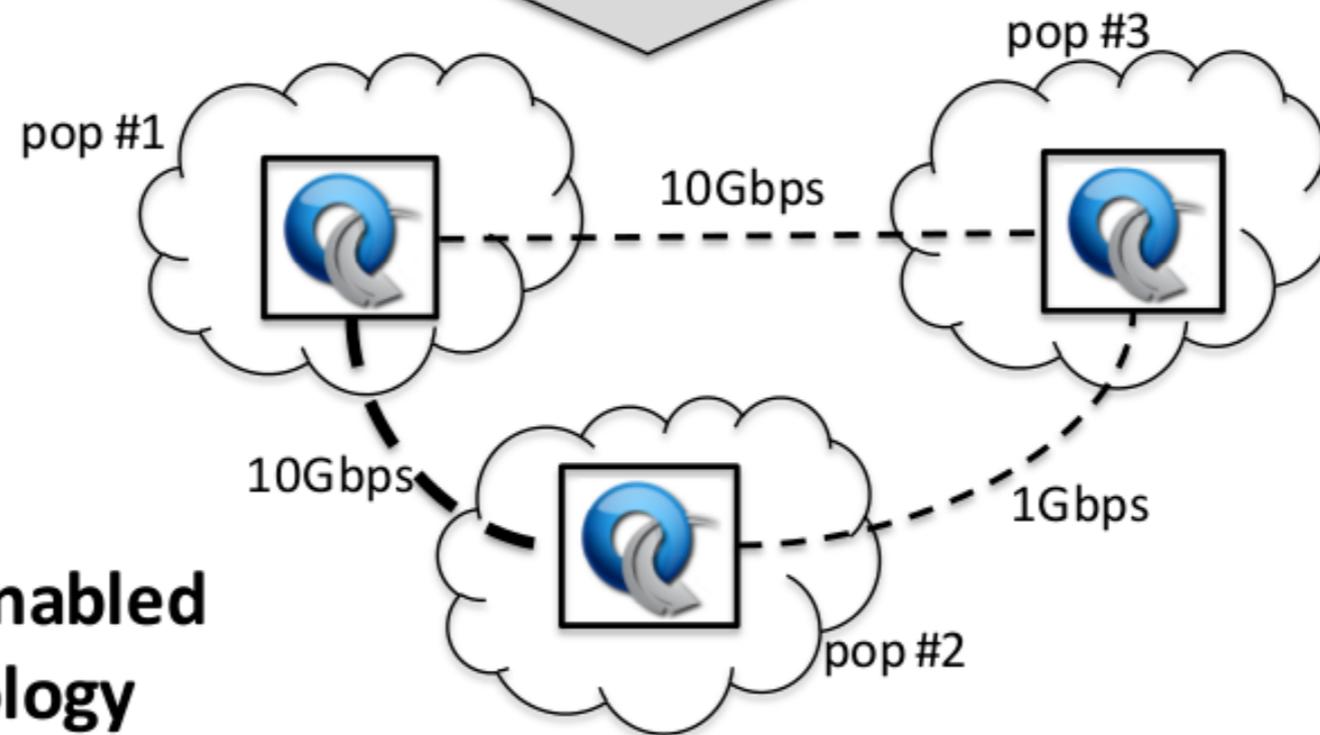
# Toulouse context



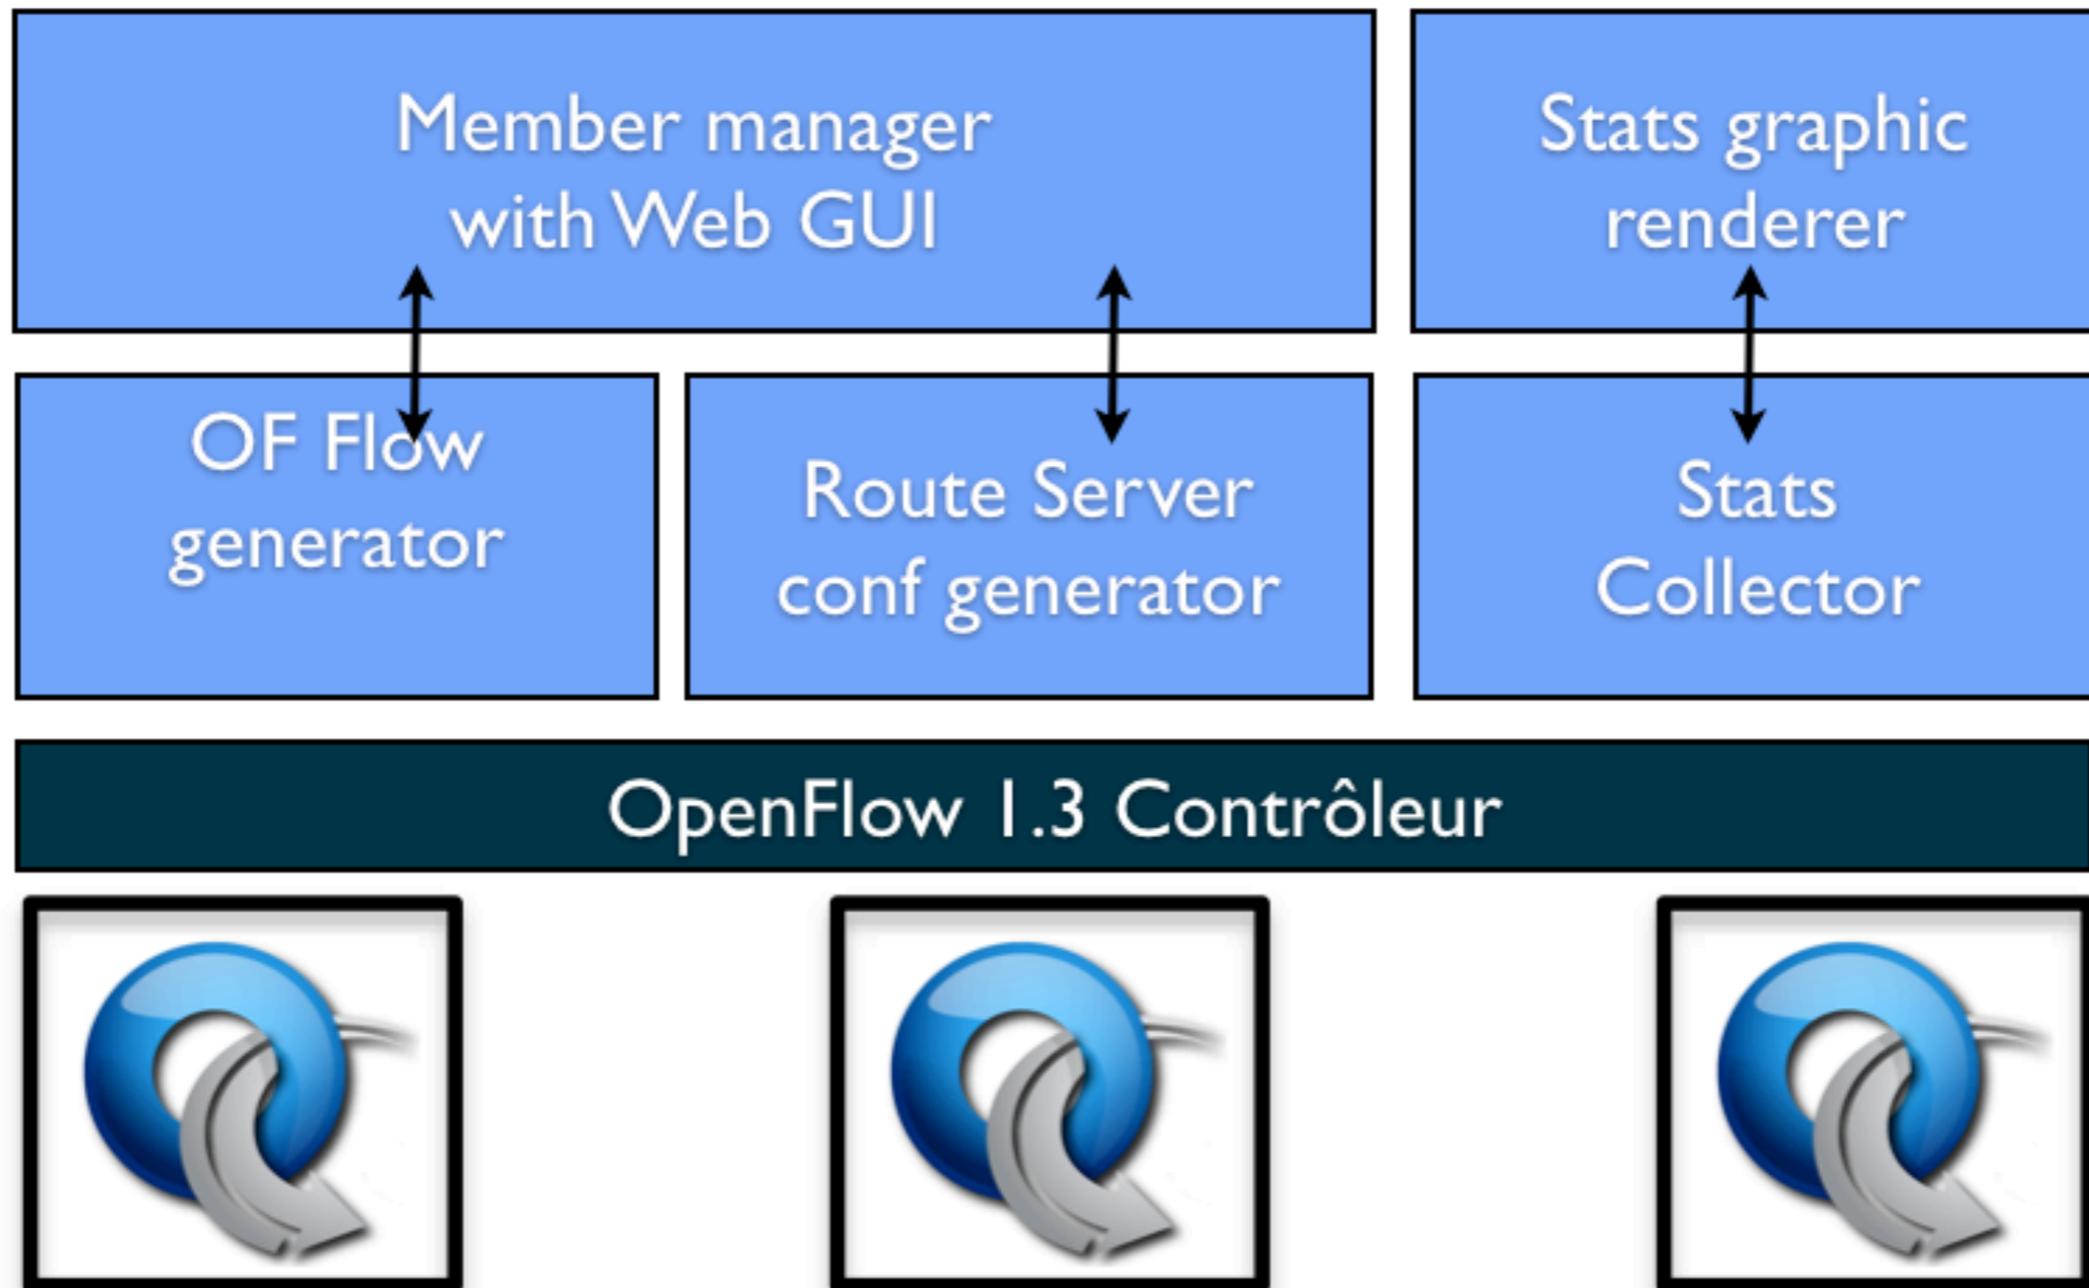
## Legacy topology



## SDN-enabled topology



# TouIX- TouSIX-Manager



**In live**

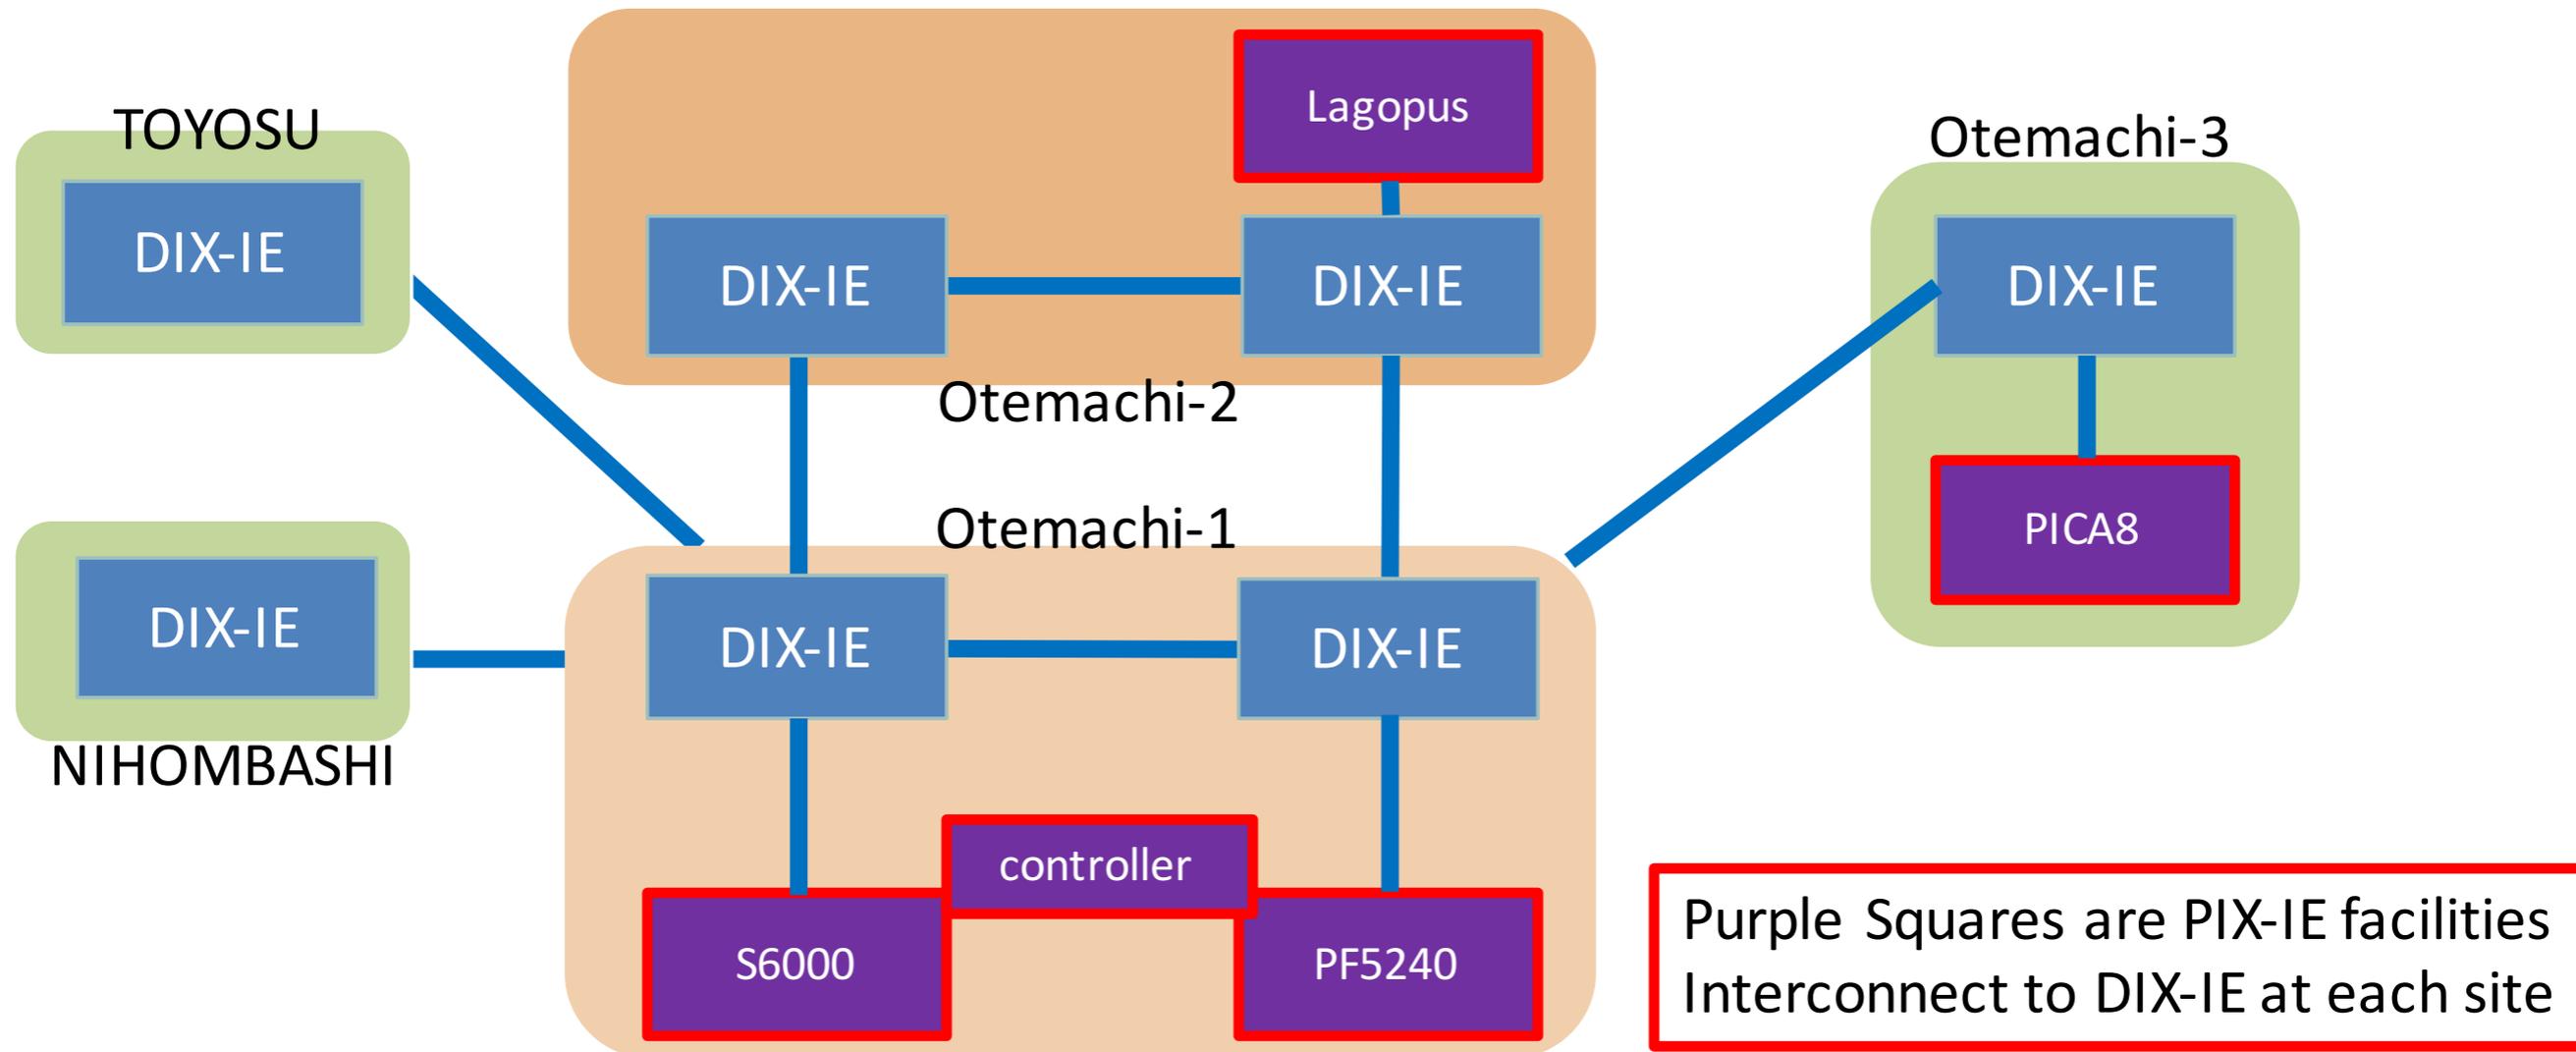
# TouSIX - Conclusion

- Umbrella proved to be a good architecture
- More than 2 years without any issue
- Creating an Open Source community is hard



- PIX-IE: A Programable Internet eXchange In Edo
- DDoS Mitigation functions
- On demand Path selection
- Proxy-ARP - external database

# PIX-IE



# PIX-IE Conclusion

- InterOp ShowNets Tokyo
- PIX-IE running in production for 2 years
- Positive combination OpenFlow and DDoS mitigation
- Proxy-ARP for IXP has important flexibility limitation

# NSPIXP-3 OSAKA

- National carrier connected to the DNS M roots server
- A single switch - 10 Operators
- Multi VLANs
- Can not be replace in one time

# FAUCET Umbrella

- Unicasting all ARP/ ICMPv6 broadcast - Done
- Group Fast Failover - Testing in progress
- Umbrella label switch encoding - WiP

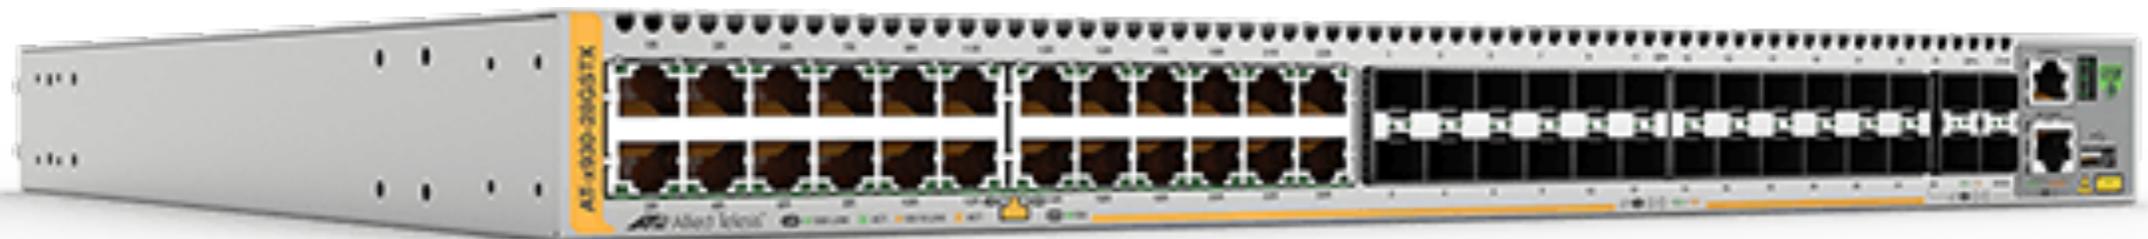
```
vlan:
  100:
    name: "dataplane_umbrella"
dps:
  dataplane:
    dp_id: 0x1
    interfaces:
      1:
        native_vlan: 100
        name: "port1"
        acl_in: 1
      2:
        native_vlan: 100
        name: "port2"
        acl_in: 1
acls:
  1:
    - rule:
      dl_type: 0x0806
      dl_dst: "ff:ff:ff:ff:ff:ff"
      arp_tpa: "10.0.0.1"
      actions:
        output:
          port: 1
    - rule:
      dl_type: 0x86dd
      ip_proto: 58
      icmpv6_type: 135
      ipv6_nd_target: "2001::1/128"
      actions:
        output:
          port: 1
```

# Allied Telesis x930-28GTX

All test OK !!

AlliedWare Plus file name : x930-5.4.6-2.6.rel

A an additional OpenFlow Licence was required.



Configuration file FAUCET.yaml  
Unicasting all broadcast  
Multi VLANs  
IPv4 and IPv6

Grafana  
Graphing Web front end  
InfluxDB  
Time Series DB

**VLAN to Tokyo**

FAUCET

GAUGE

Legacy  
Switch

OpenFlow  
Switch

Wide project  
Peering test  
Router

10Gbps

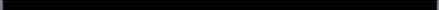
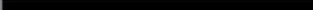
1Gbps

**VLAN peering**

Customer#1  
Router

...

Customer#2  
Router



**In live**

# Where are we

- NSPIXP-3 Full migration expected Early 2018
- TouSIX-Manager with FAUCET Umbrella End and Pica8 with AT switch expected at the end 2017
- Integrating FAUCET Umbrella in [IXP-Manager.org](http://IXP-Manager.org)



ENDEAVOUR



<http://www.h2020-endeavour.eu>



