



HolistIX: Software & Intend Based Networking within IXPs

Marc Bruyere & Christoff Visser @ IIJ Lab





HOLISTIX

How IXPs can minimize Effort, Cost and Risk ?

IXP's "Costs and Efforts"


is multi-dimensional



**Administrative
Commercial**



**Engineering
Technical**



**Hardware
Infrastructure**

IXP's "Costs and Efforts"


is multi-dimensional



**Administrative
Commercial**



**Engineering
Technical**



**Hardware
Infrastructure**

IXP's "Costs and Efforts"


is multi-dimensional



**Administrative
Commercial**



**Engineering
Technical**



**Hardware
Infrastructure**

IXP's "Costs and Efforts"


is multi-dimensional



**Administrative
Commercial**



**Engineering
Technical**



**Hardware
Infrastructure**



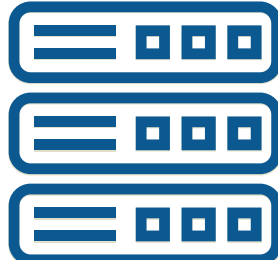


Question: The effort and cost for connecting a new IXP member?

IXP's "Costs"

Cost level to connect a new member





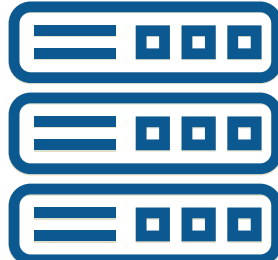
	Request To connect	Provisioning	Setup	Validation	Maintenance
 Administration Commercial	Medium	Low	None	None	None
 Engineering Technical	Low to none	Medium	High	High	Medium to high
 Hardware Infrastructure	None	Medium to high	Medium	Low to none	Low to none

IXP's "Costs"

Cost level to connect a new member









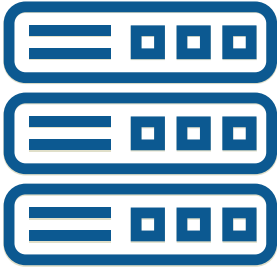
HOLISTIX

	Request To connect	Provisioning	Setup	Validation	Maintenance
 Administration Commercial	Medium	Low	None	None	None
 Engineering Technical	Low to none	Medium	High	High	Medium to high
 Hardware Infrastructure	None	Medium to high	Medium	Low to none	Low to none

IXP's "Costs"

Our Reduction Goals



	Request To connect	Provisioning	Setup	Validation	Maintenance
 Administration Commercial	Medium	Low	None	None	None
 Engineering Technical	Low to none	 Medium	 High	 High	 Medium to high
 Hardware Infrastructure	None	Medium to high	Medium	Low to none	Low to none



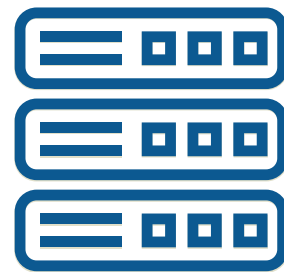


Question: Cost to change the infrastructure architecture?

IXP's "Costs"

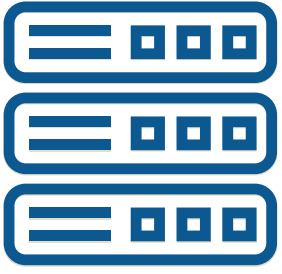


IXP infrastructure cost level



		Design	Staging	Validation Pre production	Maintenance
	Admin	Medium	None	None	None
	Technical Human cost	High	High	High	Medium to High
	Physical Equipment	High	Medium	High	None

IXP's "Costs"

IXP infrastructure cost level




	Design	Stagging	Validation Pre production	Maintenance
Admin	Medium	None	None	None
Technical Human cost	High	High	High	Medium to High
Physical Equipment	High	Medium	High	None

IXP's "Costs"

Reduction goals



	Design	Stagging	Validation Pre production	Maintenance
 Admin	✓ Medium	None	None	None
 Technical Human cost	✓ High	✓ High	✓ High	✓ Medium to High
 Physical Equipment	✓ High	✓ Medium	✓ High	None







HOLISTIX

HolistIX: Full Automation Stack



HOLISTIX

HolistIX: IXP-Manager Add-on



HOLISTIX




HOLISTIX

HolistIX: IXP-Manager Add-on



HOLISTIX





HolistIX: IXP-Manager Add-on



+





HolistIX: IXP-Manager Add-on



+



=



HolistIX: IXP-Manager Add-on



+



=





- **Free & Open Source Software Platform for IXPs**
- **Teaches and implements best practice**
- **MANRS Compliant**
- **Full stack management platform**
- **Do more with less**
- **Route server configuration**



Question :
What bring HolistIX ?



HOLISTIX

HolistIX



- IXP Manager
- Miru



HOLISTIX

HolistIX



- IXP Manager
- Miru



- Athos



HOLISTIX


HolistIX



- IXP Manager
- Miru



- Athos



- Cerberus



HolistIX


- Introduce automation from the top down for IXPs
 - Plan > Verify > Deploy
- Based on the Software Designed Network **Umbrella** switching fabric.
 - Change broadcast packets to unicast ones
 - No more quarantine time
 - Not all vendor switch can support Umbrella

Miru: Provisioning & Planning





HOLISTIX



Dashboard

Overall Member Numbers

Member Type	Count
Internal	1
Full	20

Members by VLAN

We count full and pro-bono members with at least one connected physical interface.

VLAN	Members	Percentage
Peering_VLAN	25	100%
NSPIPX3_VLAN	3	12%

Members by Location

Location	Members
KDDI Otemachi	20
NTTCom Otemachi	8
NTTData Otemachi	1

Member Ports by Location

Location	100 Mbps	1 Gbits	10 Gbits	Total
NTTCom Otemachi	0	6	2	8
NTTData Otemachi	0	0	1	1
KDDI Otemachi	5	11	4	20
Totals	5	17	7	29

Virtual Interfaces / List


Show 10 entries

Search:

Member	Facility	Switch	Port(s)	Speed	Raw Speed	Action
Acme Internet Access	Facility 1	s1	port1.0.1	1 Gbits	1000	
Alpha Corp	Facility 1	s1	port1.0.2	1 Gbits	1000	
Charlie Internet Access	Facility 1	s1	port1.0.3	1 Gbits	1000	
Delta Internet Access	Facility 1	s2	port1.0.1	1 Gbits	1000	
Echo Internet Access	Facility 1	s2	port1.0.2	1 Gbits	1000	
Foxtrot internet	Facility 1	s2	port1.0.3	1 Gbits	1000	
Golf Electric	Facility 1	s3	port1.0.1	1 Gbits	1000	



HOLISTIX



Customer	Interface(s)	VLAN	IPv4	IPv6	MAC Address	Manufacturer	Actions
Acme Internet Access	s1::port1.0.1	peering	10.0.0.1	fd00::1	000000000001	Unknown	
Acme Internet Access	s1::port1.0.1	Vlan2	10.0.1.1	2001:db8:1::1	000000000021	Unknown	
Acme Internet Access	s1::port1.0.1	vlan3	10.0.2.1	2001:db8:2::	000000000022	Unknown	
Acme Internet Access	s1::port1.0.1	vlan4	10.0.3.1	2001:db8:3::	000000000023	Unknown	
Alpha Corp	s1::port1.0.2	peering	10.0.0.2	fd00::2	000000000002	Unknown	
Charlie Internet Access	s1::port1.0.3	peering	10.0.0.3	fd00::3	000000000003	Unknown	
Delta Internet Access	s2::port1.0.1	peering	10.0.0.4	fd00::4	000000000004	Unknown	
Echo Internet Access	s2::port1.0.2	peering	10.0.0.5	fd00::5	000000000005	Unknown	
Foxtrot internet	s2::port1.0.3	peering	10.0.0.6	fd00::6	000000000006	Unknown	
golf	s3::port1.0.1	peering	10.0.0.7	fd00::7	000000000007	Unknown	



HOLISTIX

Plan

Verify

Deploy

Miru

Miru

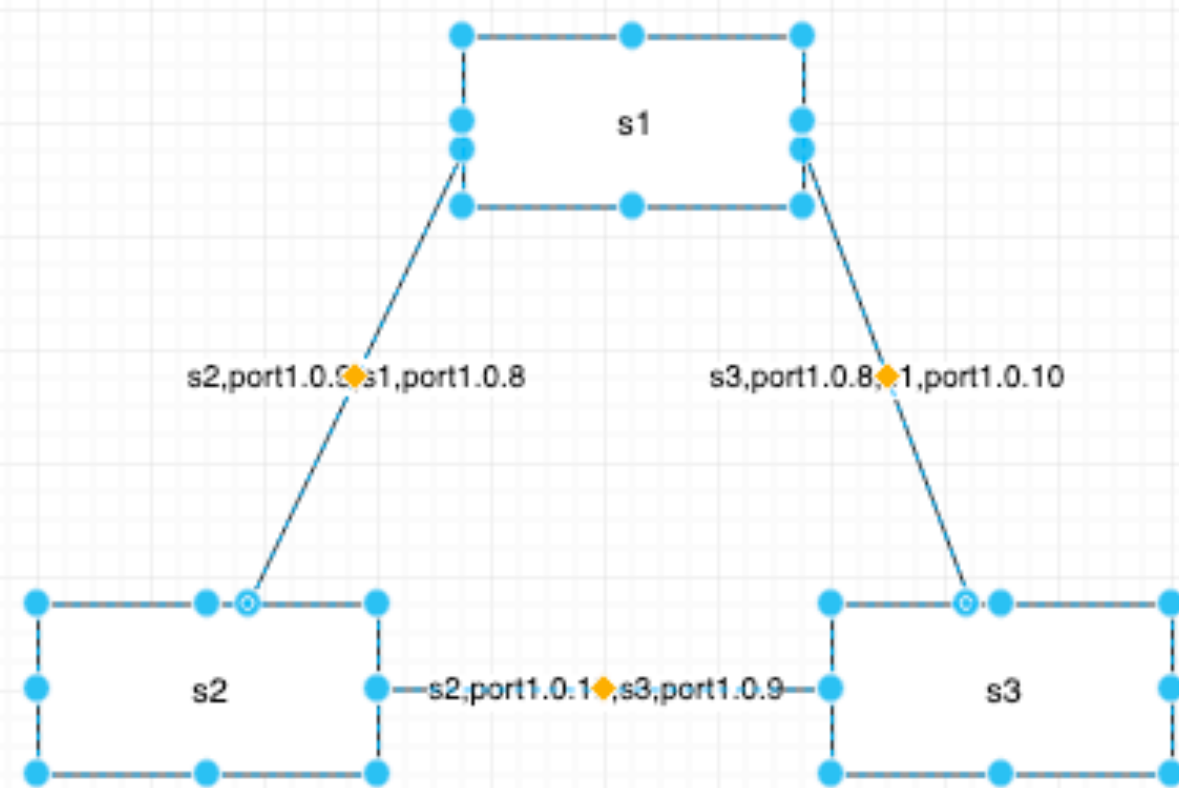
File Edit View Arrange Extras Help

100%

General



Switches



Miru

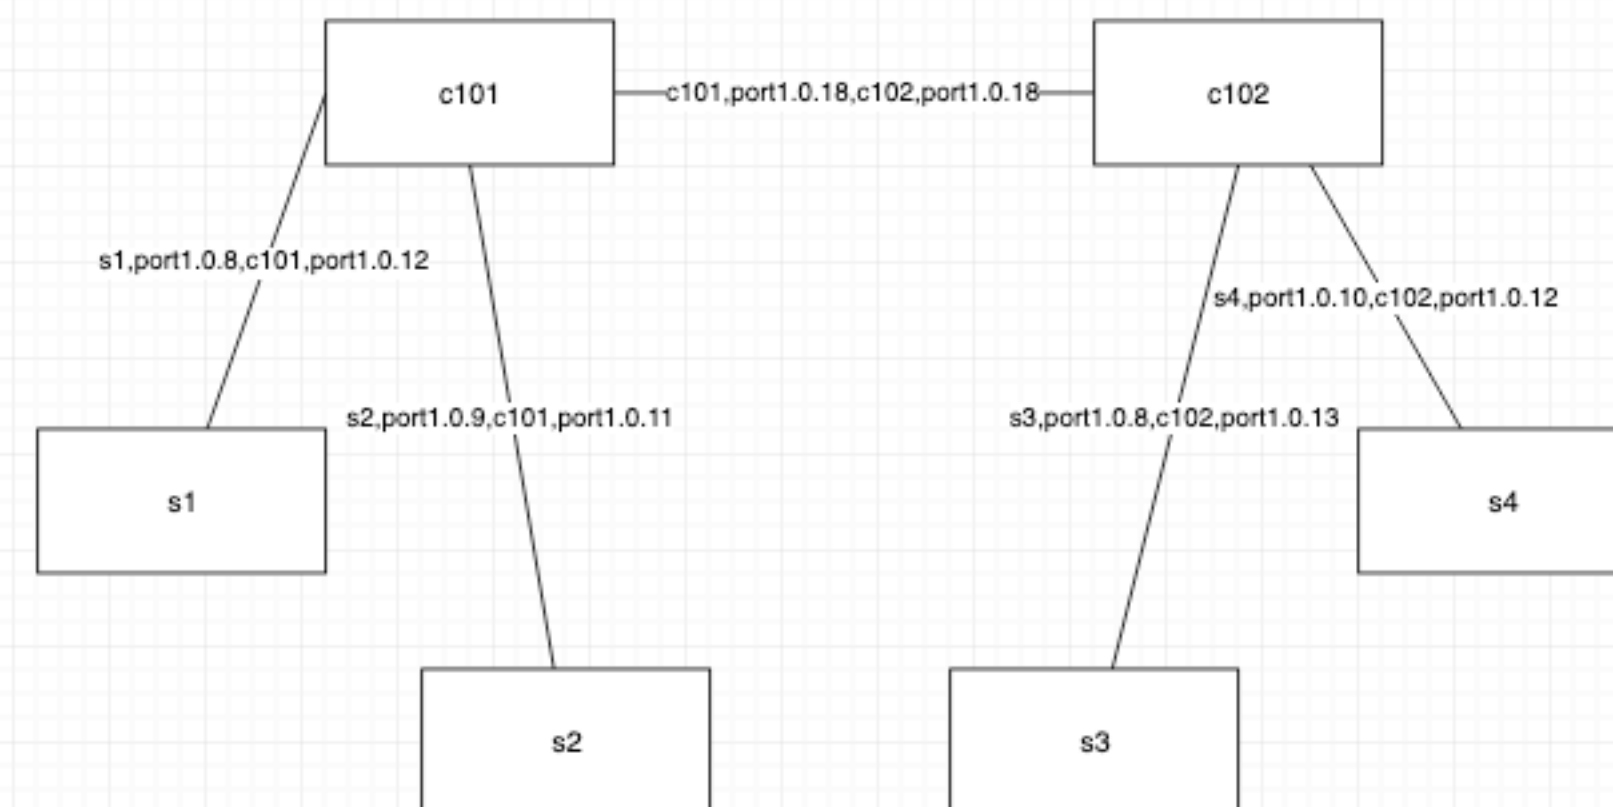
File Edit View Arrange Extras Help


100%

General



Switches





Miru

- Visual network planning
- Drag and drop diagramming
- Acts as the networks source of truth
- Generates network configurations
- Emulate and test your network before deploying it
- Deploy with a single click



HOLISTIX

Athos: No more
risk with verification





Athos

- Emulates configured network
- Test reachability between members
- Validates network redundancy
- OpenFlow support on edge Switches
- P4 support for core switches
- Docker support


Athos output

```
Echo Internet Access -> Acme Internet Access Alpha Corp Charlie Internet Access Delta Internet
Access Foxtrot internet golf IJ SCIX test VEON Group
Foxtrot internet -> Acme Internet Access Alpha Corp Charlie Internet Access Delta Internet
Access Echo Internet Access golf IJ SCIX test VEON Group
golf -> Acme Internet Access Alpha Corp Charlie Internet Access Delta Internet Access Echo
Internet Access Foxtrot internet IJ SCIX test VEON Group
IJ -> Acme Internet Access Alpha Corp Charlie Internet Access Delta Internet Access Echo
Internet Access Foxtrot internet golf SCIX test VEON Group
SCIX -> Acme Internet Access Alpha Corp Charlie Internet Access Delta Internet Access Echo
Internet Access Foxtrot internet golf IJ test VEON Group
test -> Acme Internet Access Alpha Corp Charlie Internet Access Delta Internet Access Echo
Internet Access Foxtrot internet golf IJ SCIX VEON Group
VEON Group -> Acme Internet Access Alpha Corp Charlie Internet Access Delta Internet Access
Echo Internet Access Foxtrot internet golf IJ SCIX test
*** Results: 0% dropped (110/110 received)
*** Stopping 1 controllers
faucet
*** Stopping 16 links
.....
*** Stopping 6 switches
c101 c102 s1 s2 s3 s4
*** Stopping 11 hosts
h1 h2 h3 h4 h5 h6 h7 h8 h9 h10 h11
*** Done
Success with no packet loss
```



HOLISTIX

Cerberus: Make it
work & maintained





Cerberus

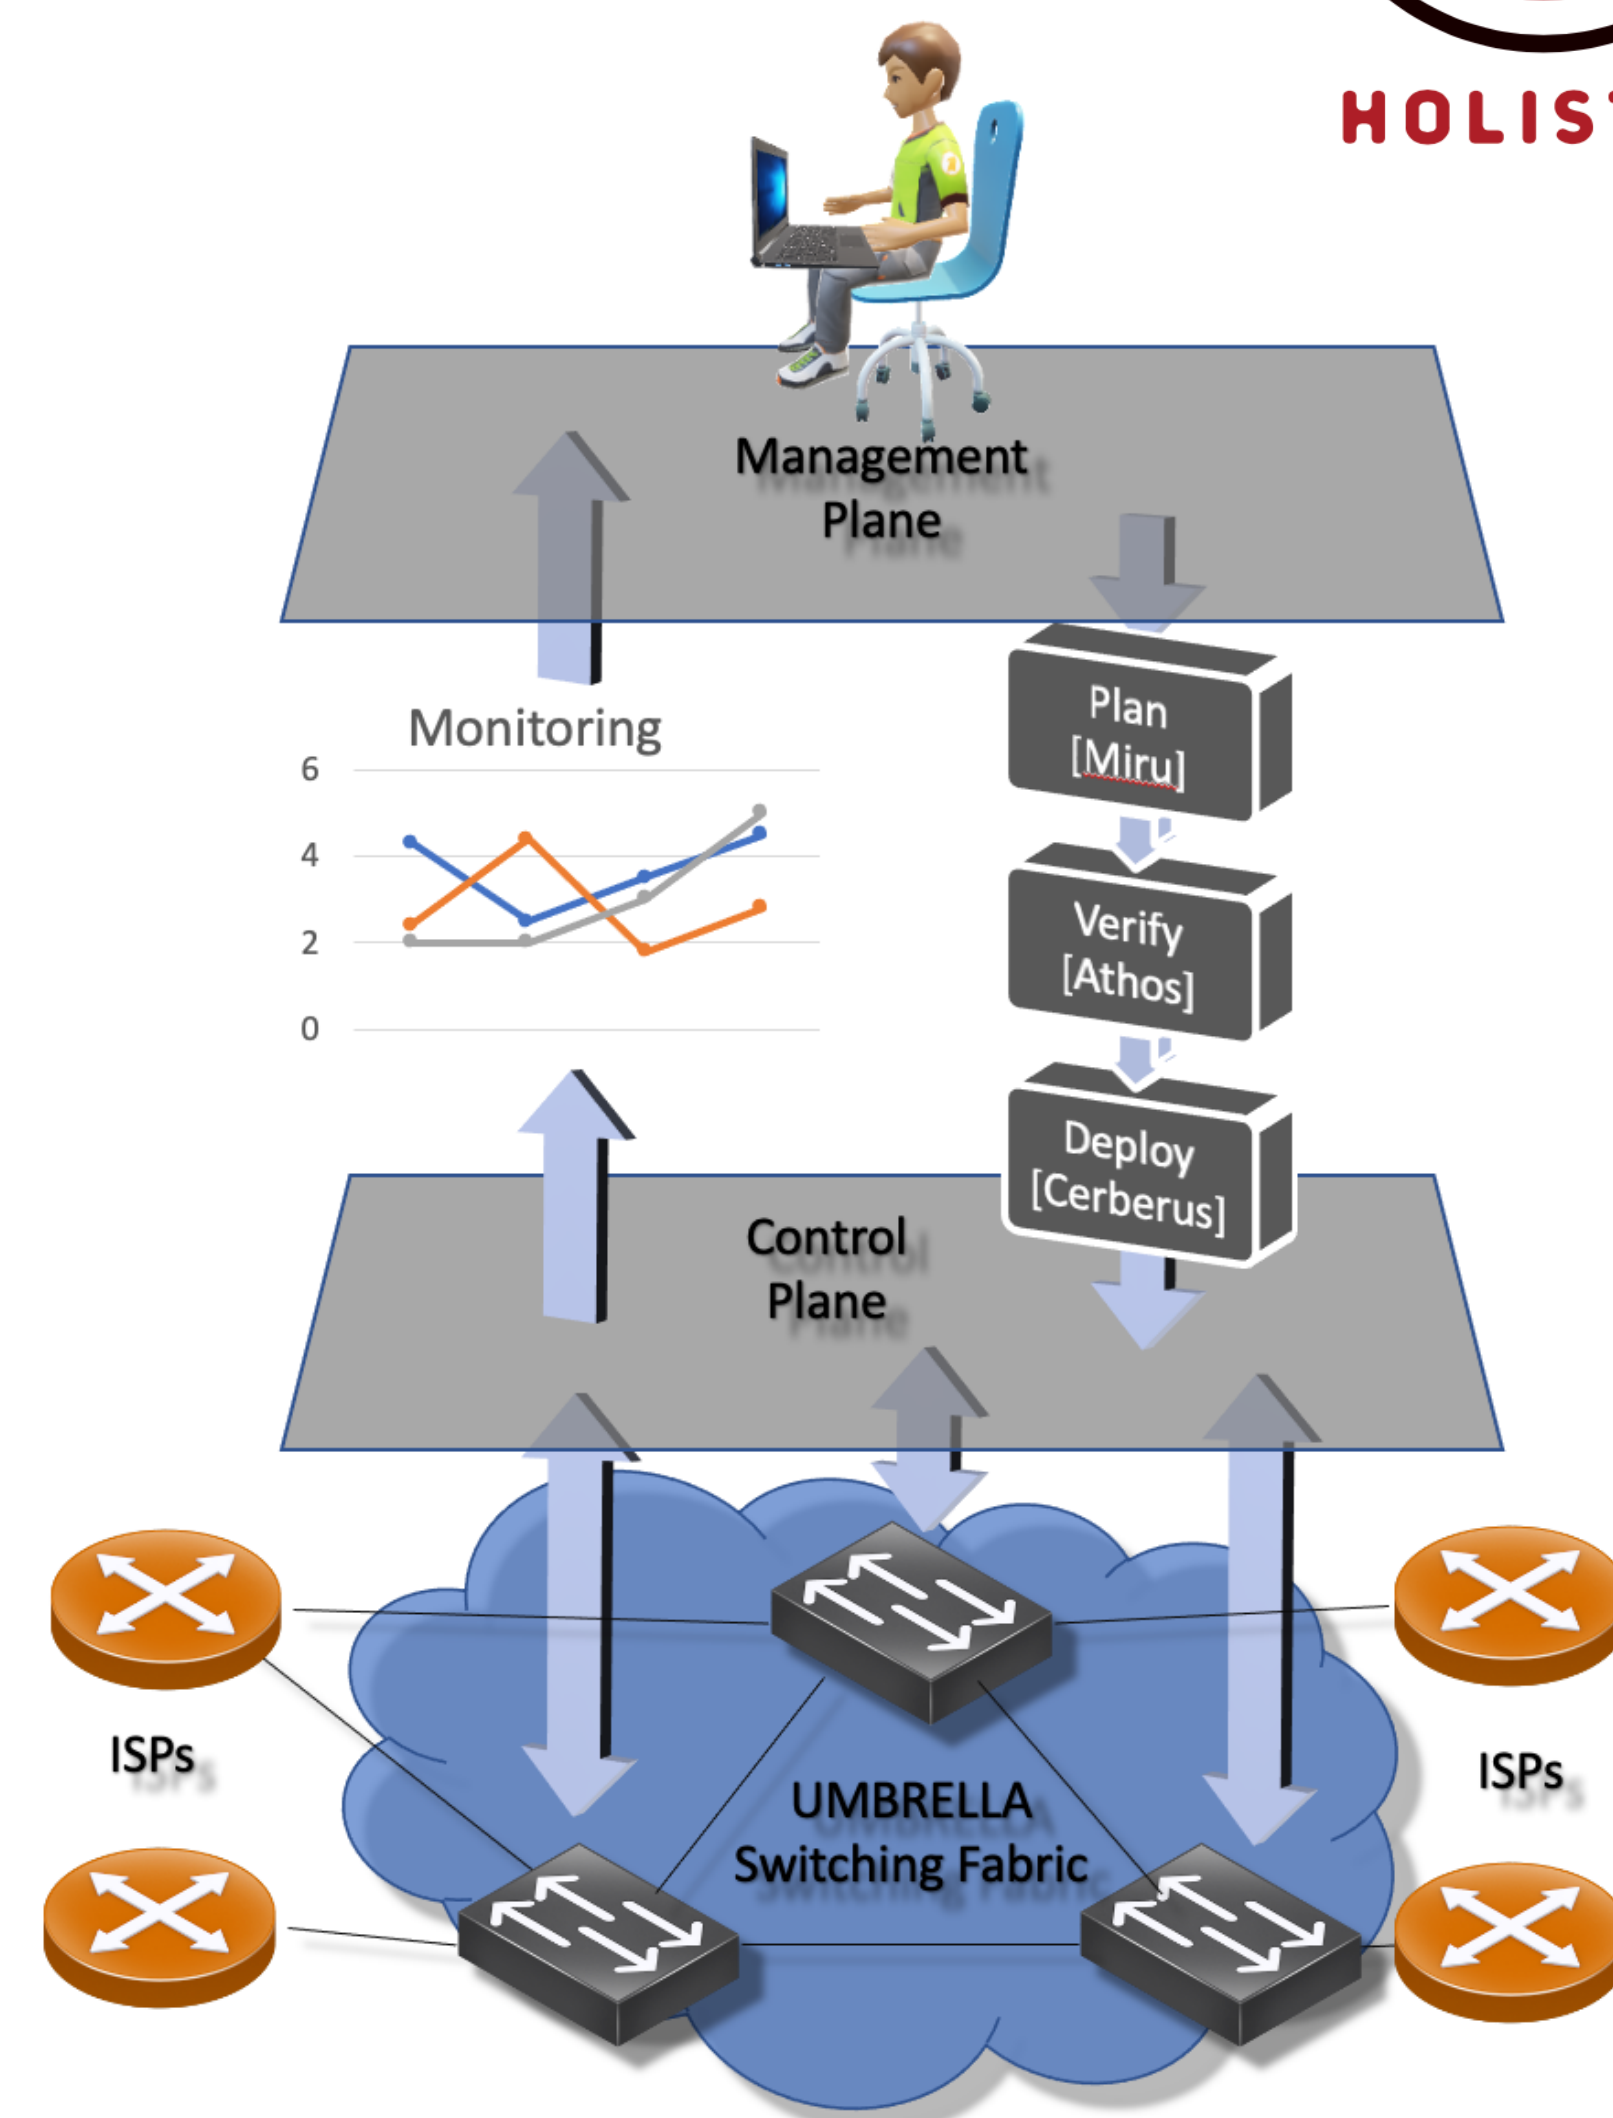
- API support to deploy from IXP Manager
- Rollback and fail state integration
- Transfer network config generation to the controller



HOLISTIX

Resume

- Automated deployment
- No more manual configuration when making changes
- Push on Green
- Made for SDN Switching Fabric





Show time:
Real hardware demo

```
MemberA:# ip a
1: lo: <LOOPBACK> mtu 65536 qdisc noop state DOWN group default qlen 1000
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
   link/ether 3c:fd:fe:02:05:06 brd ff:ff:ff:ff:ff:ff
   inet 10.200.0.1/24 scope global eth1
       valid_lft forever preferred_lft forever
   inet6 fd80:9cb1:aeff:8181::10/64 scope global
       valid_lft forever preferred_lft forever
   inet6 fe80::3efd:feff:fe02:506/64 scope link
       valid_lft forever preferred_lft forever
MemberA:# ping 10.200.0.2
PING 10.200.0.2 (10.200.0.2) 56(84) bytes of data.
[]
```

```
MemberB:# ip a
1: lo: <LOOPBACK> mtu 65536 qdisc noop state DOWN group default qlen 1000
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth1.1234@eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether 00:0c:29:af:0a:bb brd ff:ff:ff:ff:ff:ff
   inet 10.200.0.2/24 scope global eth1.1234
       valid_lft forever preferred_lft forever
   inet6 fd80:9cb1:aeff:8181::11/128 scope global
       valid_lft forever preferred_lft forever
   inet6 fe80::8678:acff:fe3c:8b03/64 scope link
       valid_lft forever preferred_lft forever
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
   link/ether 84:78:ac:3c:8b:03 brd ff:ff:ff:ff:ff:ff
MemberB:# ping 10.200.0.3
PING 10.200.0.3 (10.200.0.3) 56(84) bytes of data.
[]
```

```
MemberC:# ip a
1: lo: <LOOPBACK> mtu 65536 qdisc noop state DOWN group default qlen 1000
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
9: ens9f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
   link/ether 00:0f:1f:67:32:ea brd ff:ff:ff:ff:ff:ff
   inet 10.200.0.3/24 scope global ens9f0
       valid_lft forever preferred_lft forever
   inet6 fd80:9cb1:aeff:8181::12/64 scope global
       valid_lft forever preferred_lft forever
MemberC:# ping 10.200.0.4
PING 10.200.0.4 (10.200.0.4) 56(84) bytes of data.
[]
```

```
MemberD:# ip a
1: lo: <LOOPBACK> mtu 65536 qdisc noop state DOWN group default qlen 1000
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: ens9.1234@ens9f1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
   link/ether b8:c2:53:30:ac:5b brd ff:ff:ff:ff:ff:ff
   inet 10.200.0.4/24 scope global ens9.1234
       valid_lft forever preferred_lft forever
   inet6 fd80:9cb1:aeff:8181::13/64 scope global
       valid_lft forever preferred_lft forever
   inet6 fe80::bac2:53ff:fe30:ac5b/64 scope link
       valid_lft forever preferred_lft forever
10: ens9f1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
   link/ether b8:c2:53:30:ac:5b brd ff:ff:ff:ff:ff:ff
MemberD:# ping 10.200.0.1
PING 10.200.0.1 (10.200.0.1) 56(84) bytes of data.
[]
```




Deployments and collaboration

- **Deployed at the Toulouse IX**
- **DIX-IE -> PIX-IE WIDE Project IXP**
- **France-IX HolistIX testbed**
- **Discussion with CIVIX and KINIX**



Sustaining the peering community

- Looking for testing and demo with IXPs and ISPs
- Aim to build a long-term initiative
- Aiming to publish Academic papers as:
<https://ieeexplore.ieee.org/document/9615540>

Questions?



Links and Contact

<https://holistix.iijlab.net>

contact@holistix.email



@IxHolist