





# Peering and Transit on Whitebox Switches

FrNOG // October 6, 2023


# AGENDA

- 01** Introduction à l'Open Networking
- 02** Les ASICs génériques
- 03** Peering et Transit sur ASICs génériques


# Modèle Traditionnel




# Open Networking



Software




Hardware



Optics



# Qu'est ce qui fait le succès de ce modèle ?



Diversification des options




Sécurisation des délais d'approvisionnement



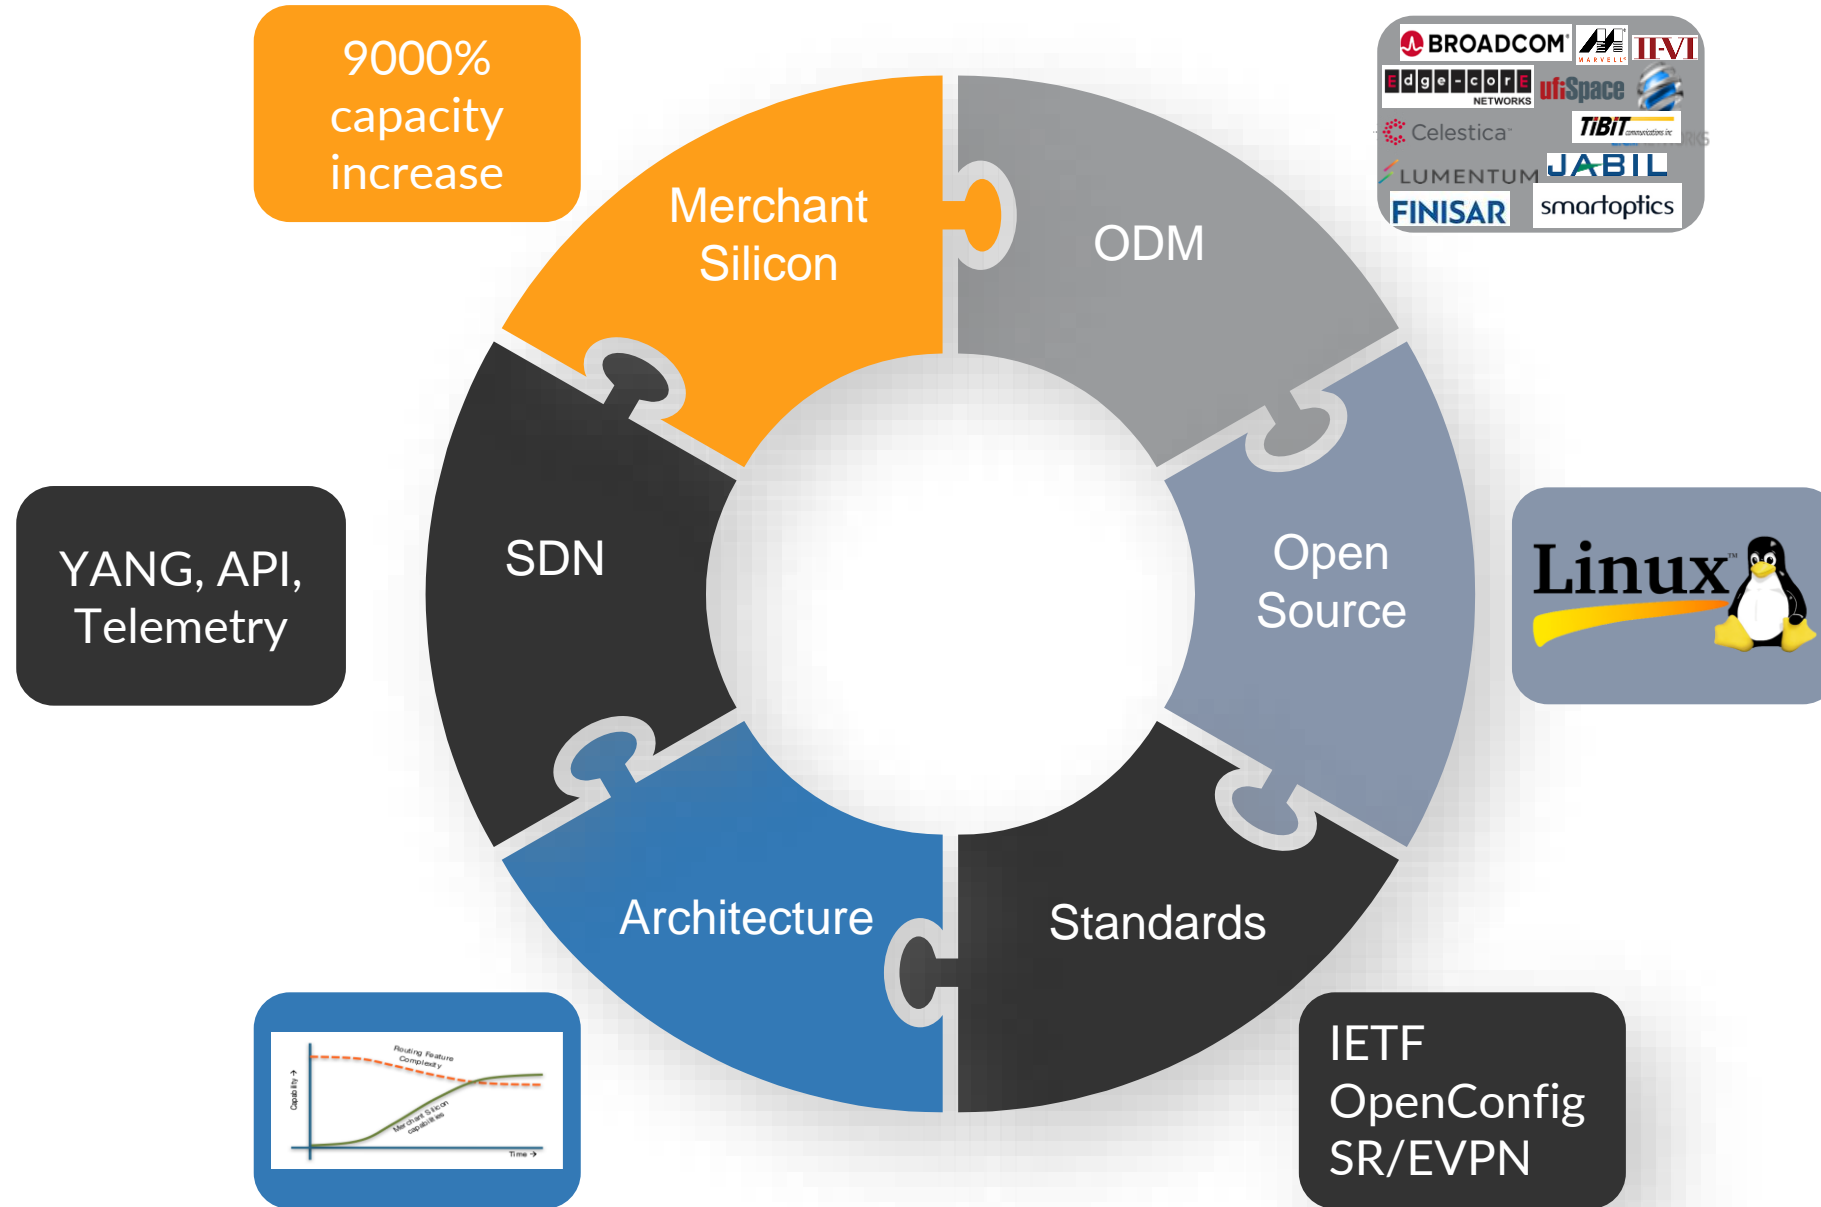
Réduction des couts > 50%



Déverrouillage de fournisseurs



# Pourquoi ce modèle est maintenant mature ?



# Un écosystème riche

## SOFTWARE

ipinfusion™

exaware

DRIVENETS

metaswitch

SONiC

## HARDWARE

Edge-core NETWORKS

ufiSpace

DELL Technologies

DELTA

Celestica™




\*Liste non exhaustive

# Peering et Transit

The background of the slide is a dark blue aerial night view of a city. The city lights are visible, including a large bridge over a river on the left. Overlaid on the city is a network of glowing blue lines and nodes, representing a network topology. The lines form a complex web of connections across the city, with several prominent nodes or hubs. The overall aesthetic is futuristic and technological.

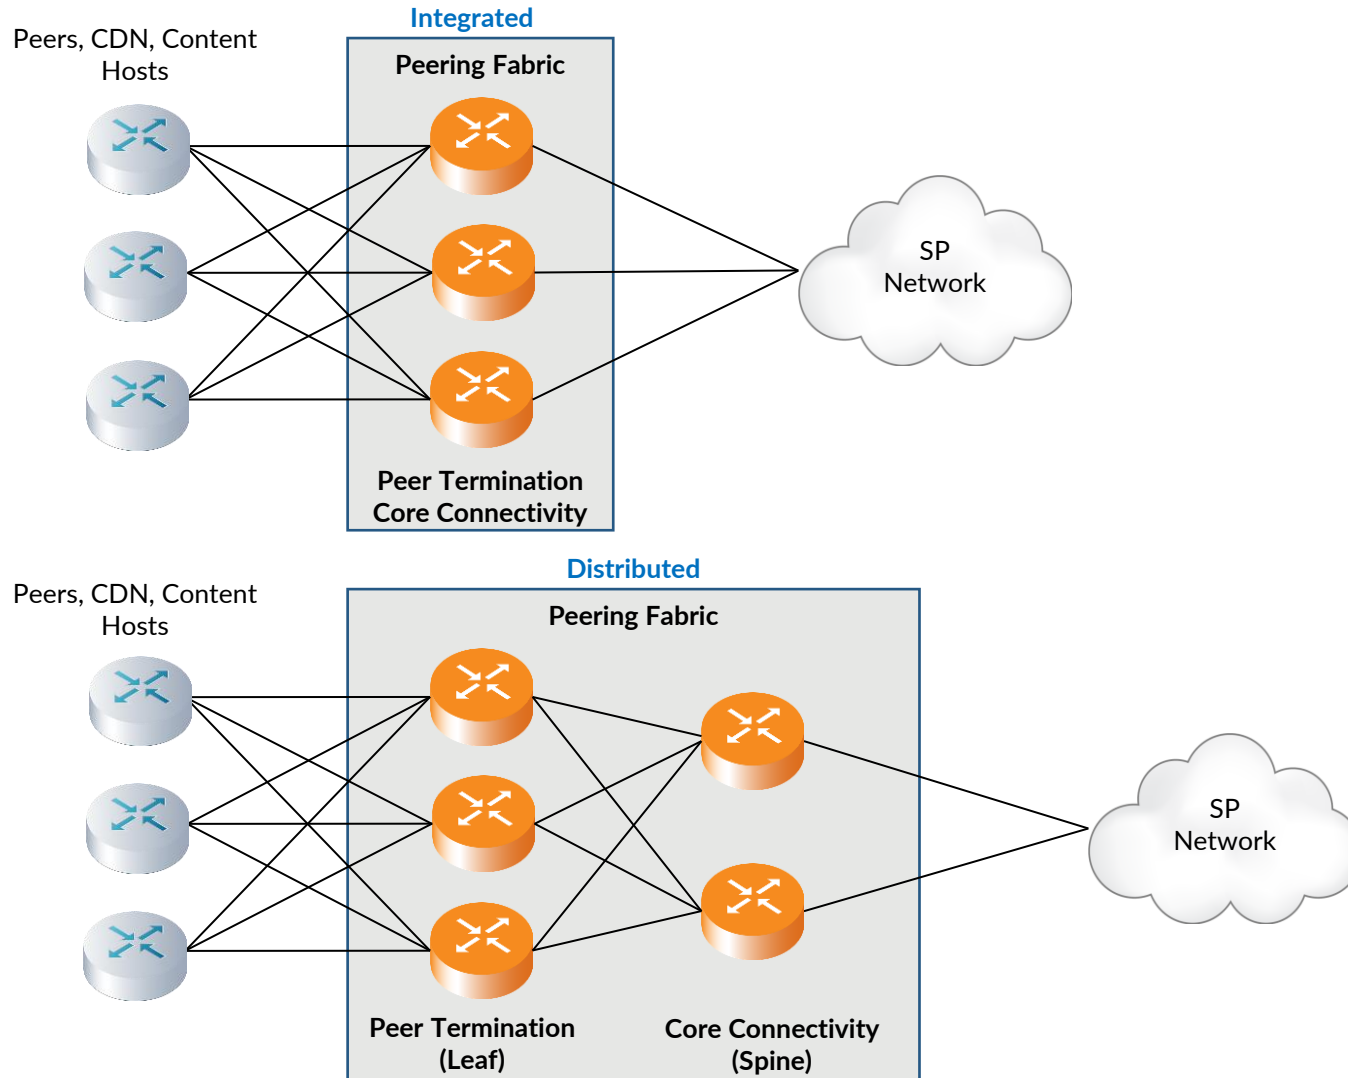
# Table Internet toujours en croissance

- IPv4 table growing at about ~9% and is projected to peak at 1.2 Million prefixes in 2028
- IPv6 is no longer growing at a linear rate
  - Projections of both linear and current 30% growth rate (EXP)
- On chip, FIB, capabilities growing at 50%
- With **external TCAM you can have 20M routes in hardware today!**





# Vers une architecture de Peering distribuée



## Internet Peering Support

- 100s of BGP Peers, 2M+ routes
- ECMP
- EVPN VXLAN and EVPN MPLS Service Overlays for exit point selection
- Support SDN-based traffic engineering and path selection to support CDN
  - SR-ISIS
  - BGP-LU
  - PCEP
  - BGP Ext GW
  - sFlow Monitoring

# Profil de forwarding du Broadcom Qumran2A

## Broadcom Qumran2A

Parameter	Forwarding Profile				
	Balanced	Balanced P-and-p	I2-xl	I3-xl	Balanced Extended
MAC	355K	372K	590K	77.5K	253K
SRV6-SID	30K	30K	30K	37K	30K
LPM - Public	1048K	917K	131K	1441K	917K
LPM- Private	1048K	917K	131K	1441K	917K



# Profil de forwarding du Broadcom Qumran2C

## Broadcom Qumran2C

Parameter	Forwarding Profile				
	Balanced	Balanced P-and-p	I2-xl	I3-xl	Balanced Extended
MAC	710K	744K	1180K	155K	506K
SRV6-SID	60K	60K	60K	75K	60K
LPM - Public	2097K	1832K	131K	2883K	1832K
LPM- Private	2097K	1832K	131K	2883K	1832K



# Broadcom Qumran2C avec TCAM externe

## Broadcom Qumran2C avec TCAM externe

Parameter	Forwarding Profile				
	Balanced	Balanced P-and-p	I2-xl	I3-xl	Balanced Extended
MAC	710K	744K	1180K	155K	506K
SRV6-SID	60K	60K	60K	75K	60K
LPM - Public	20480K	20480K	20480K	20480K	20480K
LPM- Private	20480K	20480K	20480K	20480K	20480K



# On a testé deux variantes en labo

- Deux méthodes:
  - Spirent
  - goBGP + fichiers MRT
- Spirent utilisée avec le Qumran2c (sans TCAM externe)
- goBGP utilisée avec le Qumran2a
  - Méthode Open Source pour générer la table de routage internet
  - Inclut les attributs de communauté et de route



Spirent

Qumran2c – UfiSpace 9600-72XC



goBGP +  
MRT files

Qumran2a – UfiSpace 9510-28DC



# Résultats avec le Qumran2c

## IP routing table name is Default-IP-Routing-Table(0)

IP routing table maximum-paths : 8

Total number of IPv4 routes : 1600005

Total number of IPv4 paths : 4800005

Pending routes (due to route max reached): 0

Route Source Networks

connected 5

bgp 1600000

Total 1600005

FIB 1600005

ECMP statistics (active in ASIC):

Total number of IPv4 ECMP routes : 1600000

Total number of IPv4 ECMP paths : 4800000

Number of routes with 3 ECMP paths: 1600000

OcNOS#show bgp summary

BGP router identifier 40.40.40.1, local AS number 100

BGP table version is 19

1 BGP AS-PATH entries

0 BGP community entries

3 Configured ebgp ECMP multipath: Currently set at 3

Neighbor	V	AS	MsgRcv	MsgSen	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.10.10.2	4	200	1645	1607	19	0	0	00:07:55	1600000
20.20.20.2	4	200	1645	1608	19	0	0	00:07:55	1600000
30.30.30.2	4	200	1647	21	19	0	0	00:07:55	1600000

Total number of neighbors 3

Total number of Established sessions 3

OcNOS#show ip route summary

- ufiSpace 9600-72xc
- 3 eBGP neighbors sending 1.6m routes each
- 1.6m routes in FIB
- 4.8m routes in RIB



# Résultats avec le Qumran2a

CE2#sh ip bgp summary

BGP router identifier 10.255.255.2, local AS number 65004

BGP table version is 5

121887 BGP AS-PATH entries

204 BGP community entries

Neighbor	V	AS	MsgRcv	MsgSen	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.255.255.1		4	65555	140443	7	5	0	0 00:02:03	926264

Total number of neighbors 1

Total number of Established sessions 1

CE2#sh ip route summary

-----  
IP routing table name is Default-IP-Routing-Table(0)  
-----

IP routing table maximum-paths : 8

Total number of IPv4 routes : 926266

Total number of IPv4 paths : 926266

Pending routes (due to route max reached): 0

Route Source Networks

connected 2

bgp 926264

Total 926266

FIB 926266

ECMP statistics (active in ASIC):

Total number of IPv4 ECMP routes : 0

Total number of IPv4 ECMP paths : 0

- ufiSpace 9510-28DC
- 1 eBGP neighbors sending actual internet table based on latest MRT file
- **926k routes programmed in FIB**
- **926k routes in RIB**
- L3-XL profile required
- Balanced profile maxes at 820k



Le modèle Open Networking est mature

Les ASICs génériques apportent les performances

Les bénéfices **diversification - flexibilité - coûts** sont importants

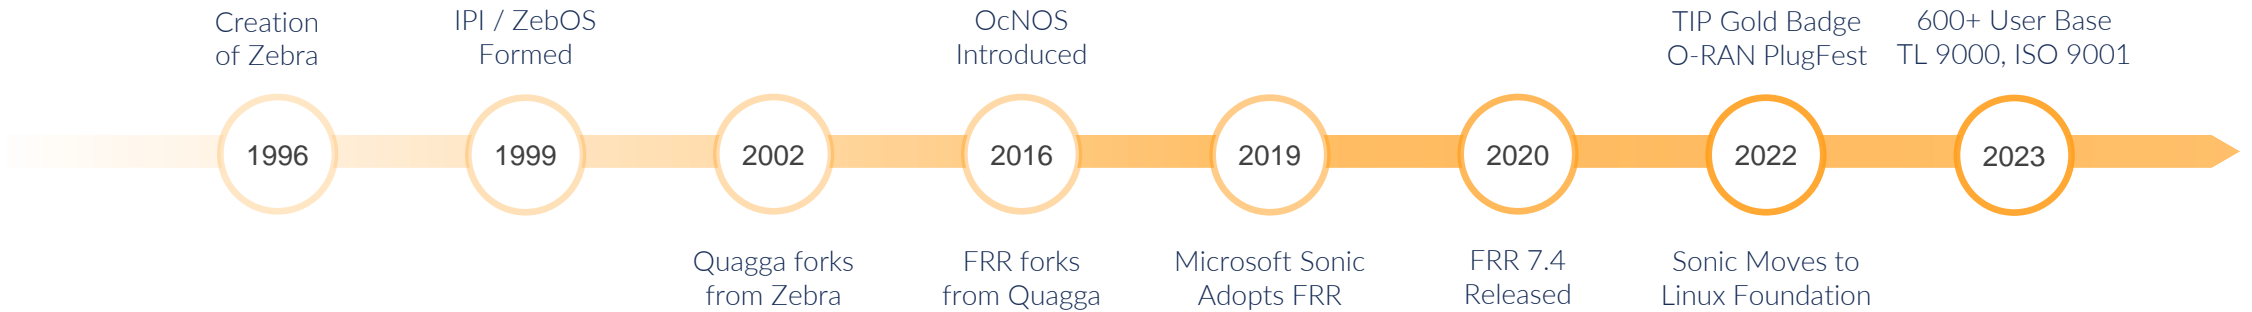


# IP Infusion (IPI) OcNOS

The background of the slide is a dark blue aerial night view of a city. The city lights are visible, including a large bridge over a river on the left. Overlaid on the city is a network diagram consisting of several large, overlapping, semi-transparent blue circles. These circles are interconnected by thin, glowing blue lines, creating a mesh-like structure that represents a network topology. The overall aesthetic is futuristic and technological.

# Pourquoi IP Infusion?

Insight on rapid growth and how to grow with our team



## Industry Name

24 Years of Innovative Network Software Solutions vs 7 Years average for other NOS Vendors

## Quality

Extensive user base for Control Plane NOS, including tens of Tier-1 Vendors

Rigorous software and platform testing and validation methodologies

## Innovation

Extensive routing, switching, and timing feature sets combined with high optimization.

Some of the widest platform support and solution ecosystems

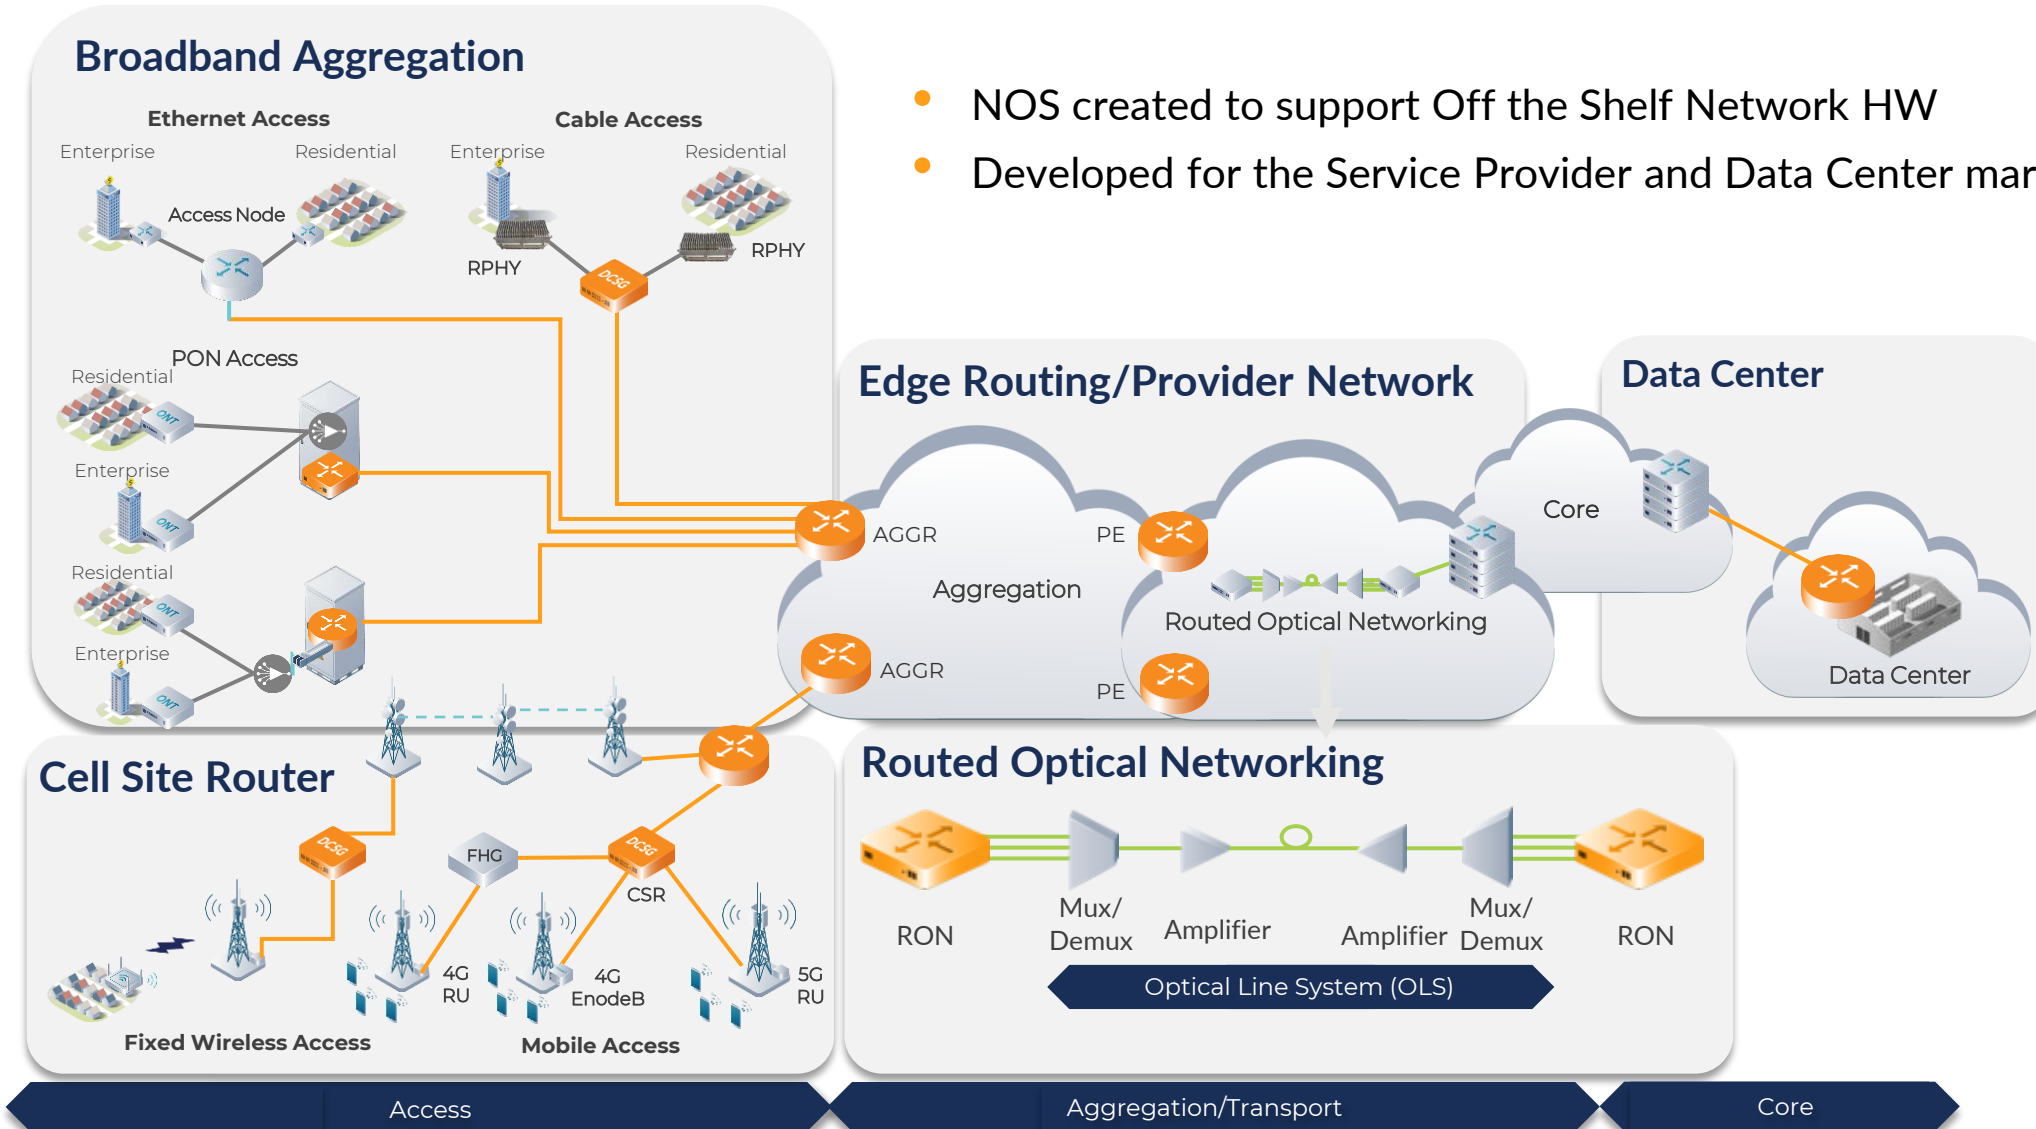
## Support

Mature global and regional channel map with local team support, resellers, and service providers

Well established one-stop shopping and service experience for NOS and hardware platforms

# OcNOS: Open compute Network Operating System

- NOS created to support Off the Shelf Network HW
- Developed for the Service Provider and Data Center markets





[contact@pine-networks.com](mailto:contact@pine-networks.com)

THANK YOU

[WWW.IPINFUSION.COM](http://WWW.IPINFUSION.COM)