SP Edge Security Service
Scale-out Services Architectures

Melchior Aelmans
maelmans@juniper.net

JUNIPER NETWORKS

Forward-Looking Statements

This presentation contains forward-looking statements within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended, which statements involve substantial risks and uncertainties. Except for historical information contained herein, all statements could be deemed forward-looking statements, including, without limitation, Juniper Networks' views concerning our business, economic and market outlook; our expectations with respect to market trends; our product development; the strength of certain use-cases and customer segments; the introduction of future products; the strength of our solution portfolio; the timing of recovery from COVID-19 on customer demand and resolution of supply challenges; and overall future prospects.

Actual results or events could differ materially from those anticipated in those forward-looking statements as a result of several factors, including: general economic and political conditions globally or regionally; the duration of the effects of the COVID-19 pandemic; business and economic conditions in the networking industry; changes in the financial stability of and overall technology spending by our customers; the network capacity requirements of our customers and, in particular, cloud and communication service providers; the timing of orders and their fulfillment; manufacturing and supply chain constraints, changes or disruptions in our business operations caused by, among other things, armed conflicts, cyberwarfare, political tensions, natural disasters and climate change; availability of product components; delays in scheduled product availability; adoption of regulations or standards affecting Juniper Networks' products, services or the networking industry; the impact of inflationary pressures; executive orders, tariffs, governmental sanctions, changes in laws or regulations and accounting rules, or interpretations thereof; and other factors listed in Juniper Networks' most recent reports on Form 10-Q and 10-K filed with the Securities and Exchange Commission. These forward-looking statements are not guarantees of future performance and speak only as of the date of this presentation. Juniper Networks undertakes no obligation to update the information in this presentation in the event facts or circumstances subsequently change.

Statement of Product Direction. Juniper Networks may disclose information related to development and plans for future products, features or enhancements, known as a Plan of Record ("POR"). These details provided are based on Juniper's current development efforts and plans. These development efforts and plans are subject to change at Juniper's sole discretion, without notice. Except as may be set forth in definitive agreements, Juniper Networks provides no assurances and assumes no responsibility to introduce products, features or enhancements described in this presentation. Purchasing decisions by third-parties should not be based on this POR and no purchases are contingent upon Juniper Networks delivering any feature or functionality depicted in this presentation.

Company Logos. Juniper Networks, the Juniper Networks logo, Juniper, Junos, and Mist AI are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. Other names and/or logos may be trademarks of their respective owners, and Juniper Networks' use hereof does not imply an affiliation with, or endorsement by, the owners of these trademarks or logos.

JAN 2024

INTERNET USE TIMELINE

NUMBER OF INDIVIDUALS USING THE INTERNET OVER TIME (IN MILLIONS)

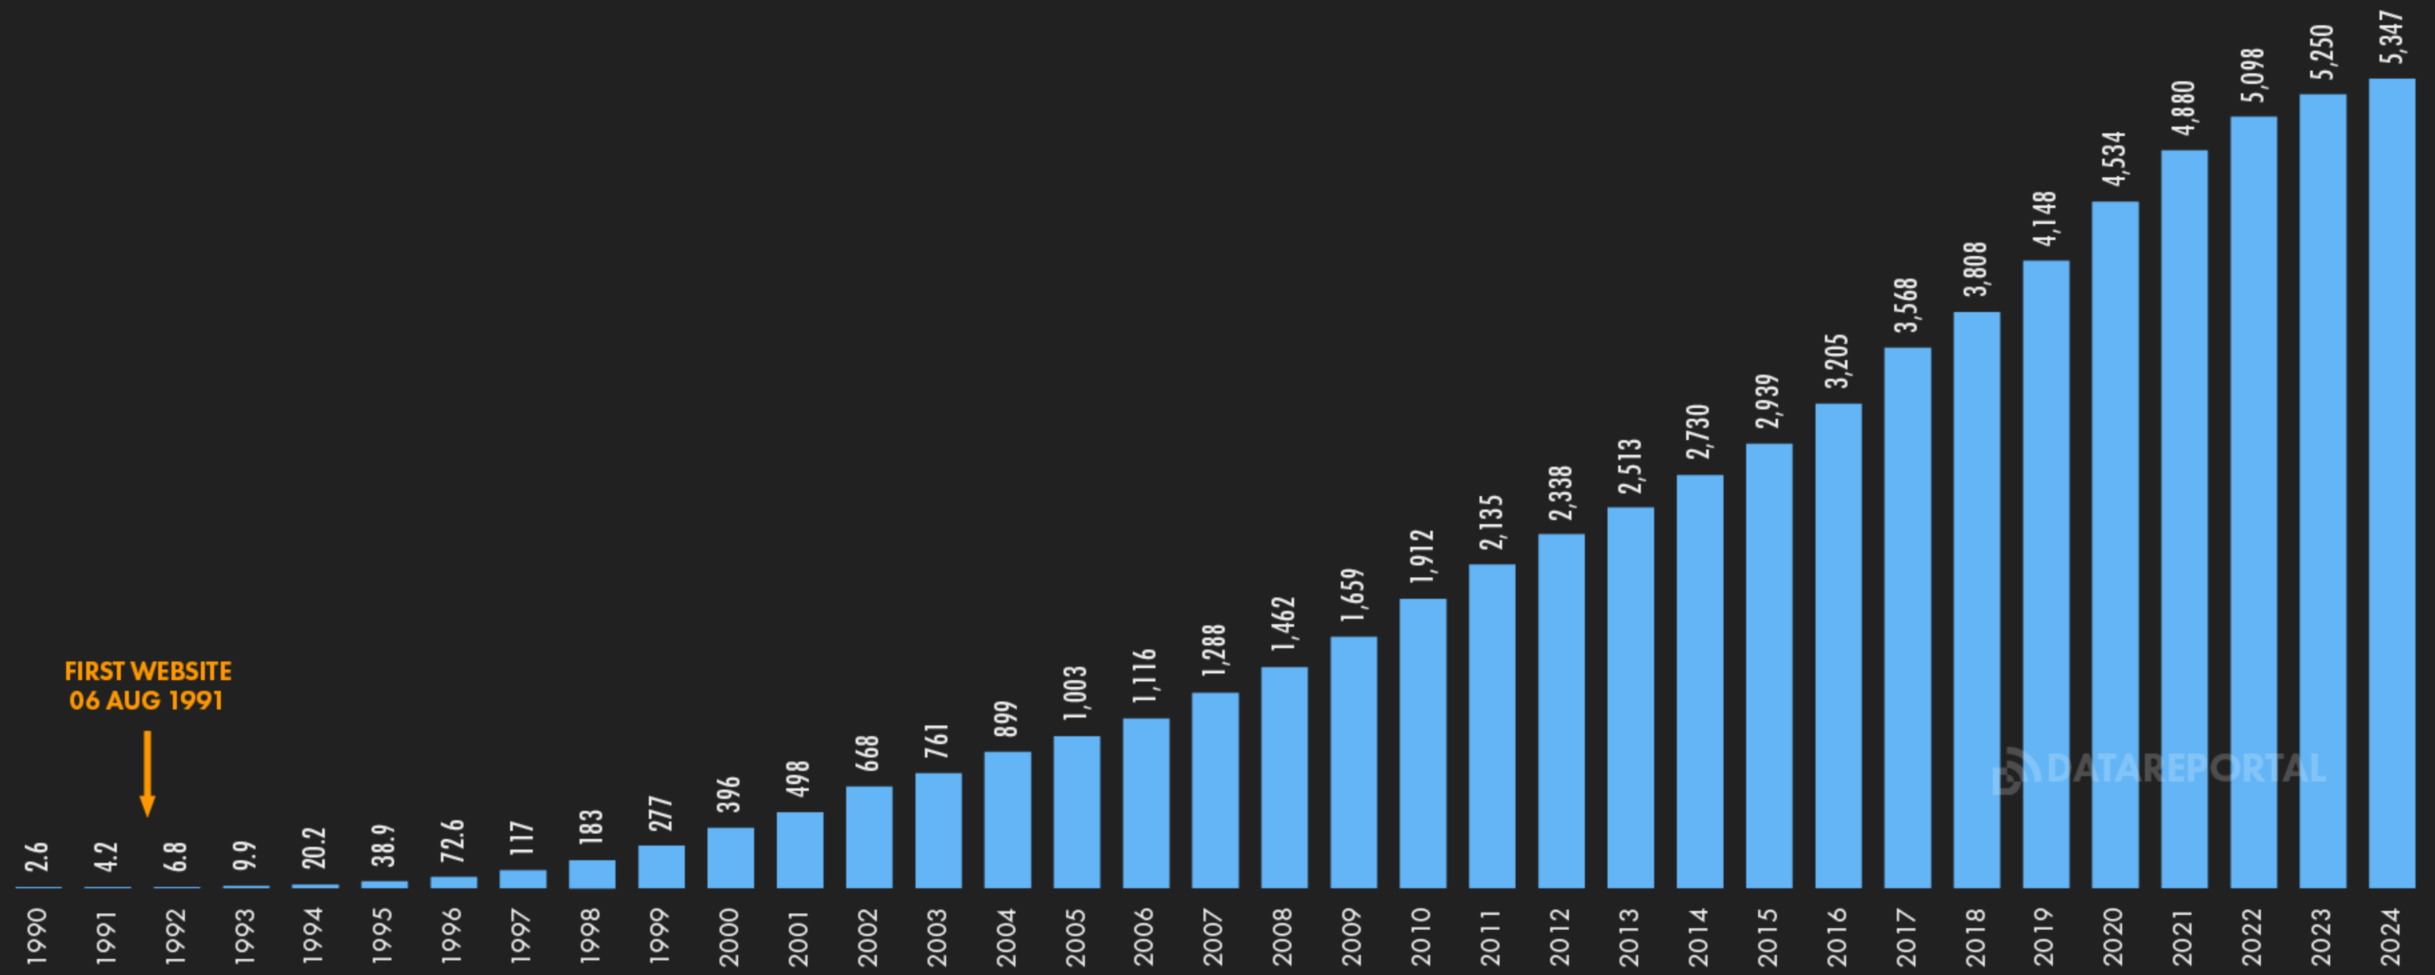

GLOBAL OVERVIEW

SOURCES: KEPIOS ANALYSIS; ITU; GSMA INTELLIGENCE; EUROSTAT; GOOGLE'S ADVERTISING RESOURCES; CNNIC; KANTAR & IAMAI; GOVERNMENT RESOURCES; UNITED NATIONS. **COMPARABILITY:** SOURCE AND BASE CHANGES. ALL FIGURES USE THE LATEST AVAILABLE DATA, BUT SOME SOURCES DO NOT PUBLISH REGULAR UPDATES, SO FIGURES FOR RECENT PERIODS MAY UNDER-REPRESENT ACTUAL USE. SEE NOTES ON DATA.

we are social   Meltwater

OVERVIEW OF INTERNET USE

ESSENTIAL INDICATORS OF INTERNET ADOPTION AND USE

INDIVIDUALS USING THE INTERNET	INDIVIDUALS USING THE INTERNET AS A PERCENTAGE OF TOTAL POPULATION	YEAR-ON-YEAR CHANGE IN THE NUMBER OF INDIVIDUALS USING THE INTERNET	PERCENTAGE OF THE TOTAL FEMALE POPULATION THAT USES THE INTERNET	PERCENTAGE OF THE TOTAL MALE POPULATION THAT USES THE INTERNET
5.35 BILLION	66.2% YOY: +0.9% (+60 BPS)	+1.8% +97 MILLION	63.5% YOY: +4.9% (+304 BPS)	68.8% YOY: +4.2% (+285 BPS)

AVERAGE DAILY TIME SPENT USING THE INTERNET BY EACH INTERNET USER	PERCENTAGE OF USERS ACCESSING THE INTERNET VIA MOBILE PHONES	PERCENTAGE OF USERS ACCESSING THE INTERNET VIA LAPTOPS AND DESKTOPS	PERCENTAGE OF THE TOTAL URBAN POPULATION THAT USES THE INTERNET	PERCENTAGE OF THE TOTAL RURAL POPULATION THAT USES THE INTERNET
6H 40M YOY: +0.8% (+3 MINS)	96.5% YOY: +4.6% (+420 BPS)	61.8% YOY: -5.8% (-380 BPS)	78.8% YOY: +3.2% (+252 BPS)	48.9% YOY: +7.2% (+340 BPS)

SOURCES: KEPIOS ANALYSIS; ITU; GSMA INTELLIGENCE; EUROSTAT; GOOGLE'S ADVERTISING RESOURCES; CNNIC; KANTAR & IAMAI; GOVERNMENT RESOURCES; UNITED NATIONS. TIME SPENT AND MOBILE SHARE DATA FROM GWI (Q3 2023). SEE GWI.COM. **NOTES:** GENDER DATA ARE ONLY AVAILABLE FOR "FEMALE" AND "MALE". PERCENTAGE CHANGE FIGURES SHOW RELATIVE YEAR-ON-YEAR CHANGE. "BPS" FIGURES REPRESENT BASIS POINTS, AND SHOW ABSOLUTE YEAR-ON-YEAR CHANGE. **COMPARABILITY:** SOURCE AND BASE CHANGES. ALL FIGURES USE THE LATEST AVAILABLE DATA, BUT SOME SOURCES DO NOT PUBLISH REGULAR UPDATES, SO FIGURES MAY UNDER-REPRESENT ACTUAL USE. SEE NOTES ON DATA.

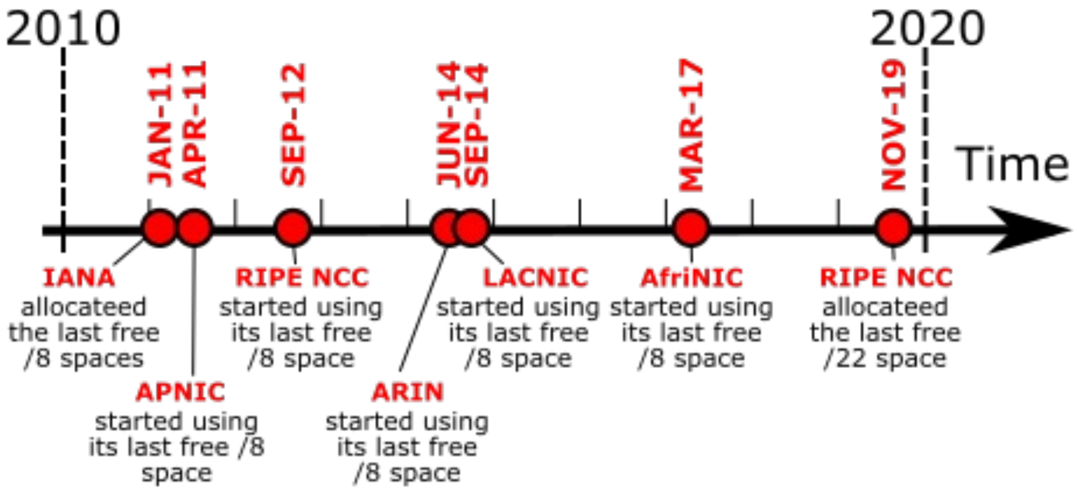
2010
2020
JAN-11
APR-11
SEP-12
JUN-14
SEP-14
MAR-17
NOV-19
Time
IANA allocateed the last free /8 spaces
APNIC started using its last free /8 space
RIPE NCC started using its last free /8 space
ARIN started using its last free /8 space
LACNIC started using its last free /8 space
AfriNIC started using its last free /8 space
RIPE NCC allocateed the last free /22 space

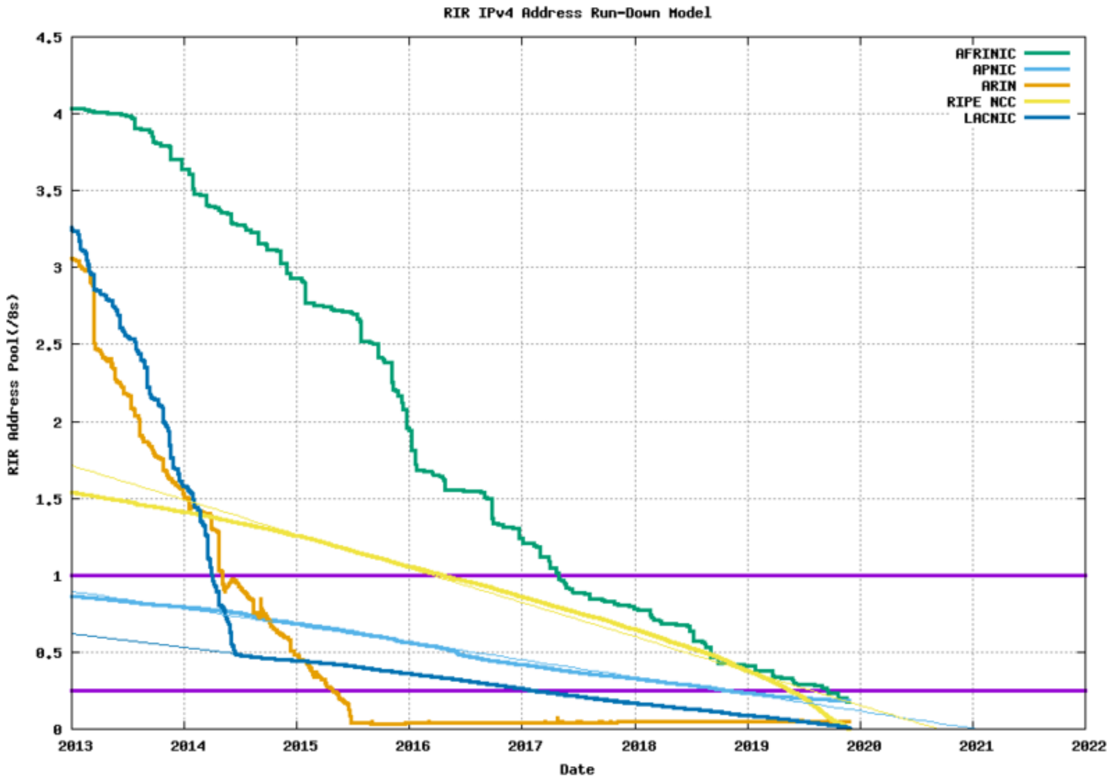
RIR IPv4 Address Run-Down Model
AFRINIC
APNIC
ARIN
RIPE NCC
LACNIC
RIR Address Pool(/8s)
Date
4.5
4
3.5
3
2.5
2
1.5
1
0.5
0
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022

IPv6 Adoption | Per-Country IPv6 adoption

IPv6 Adoption

We are continuously measuring the availability of IPv6 connectivity among Google users. The graph shows the percentage of users that access Google over IPv6.

Native: 43.36% 6to4/Teredo: 0.00% Total IPv6: 43.36% | Feb 25, 2024

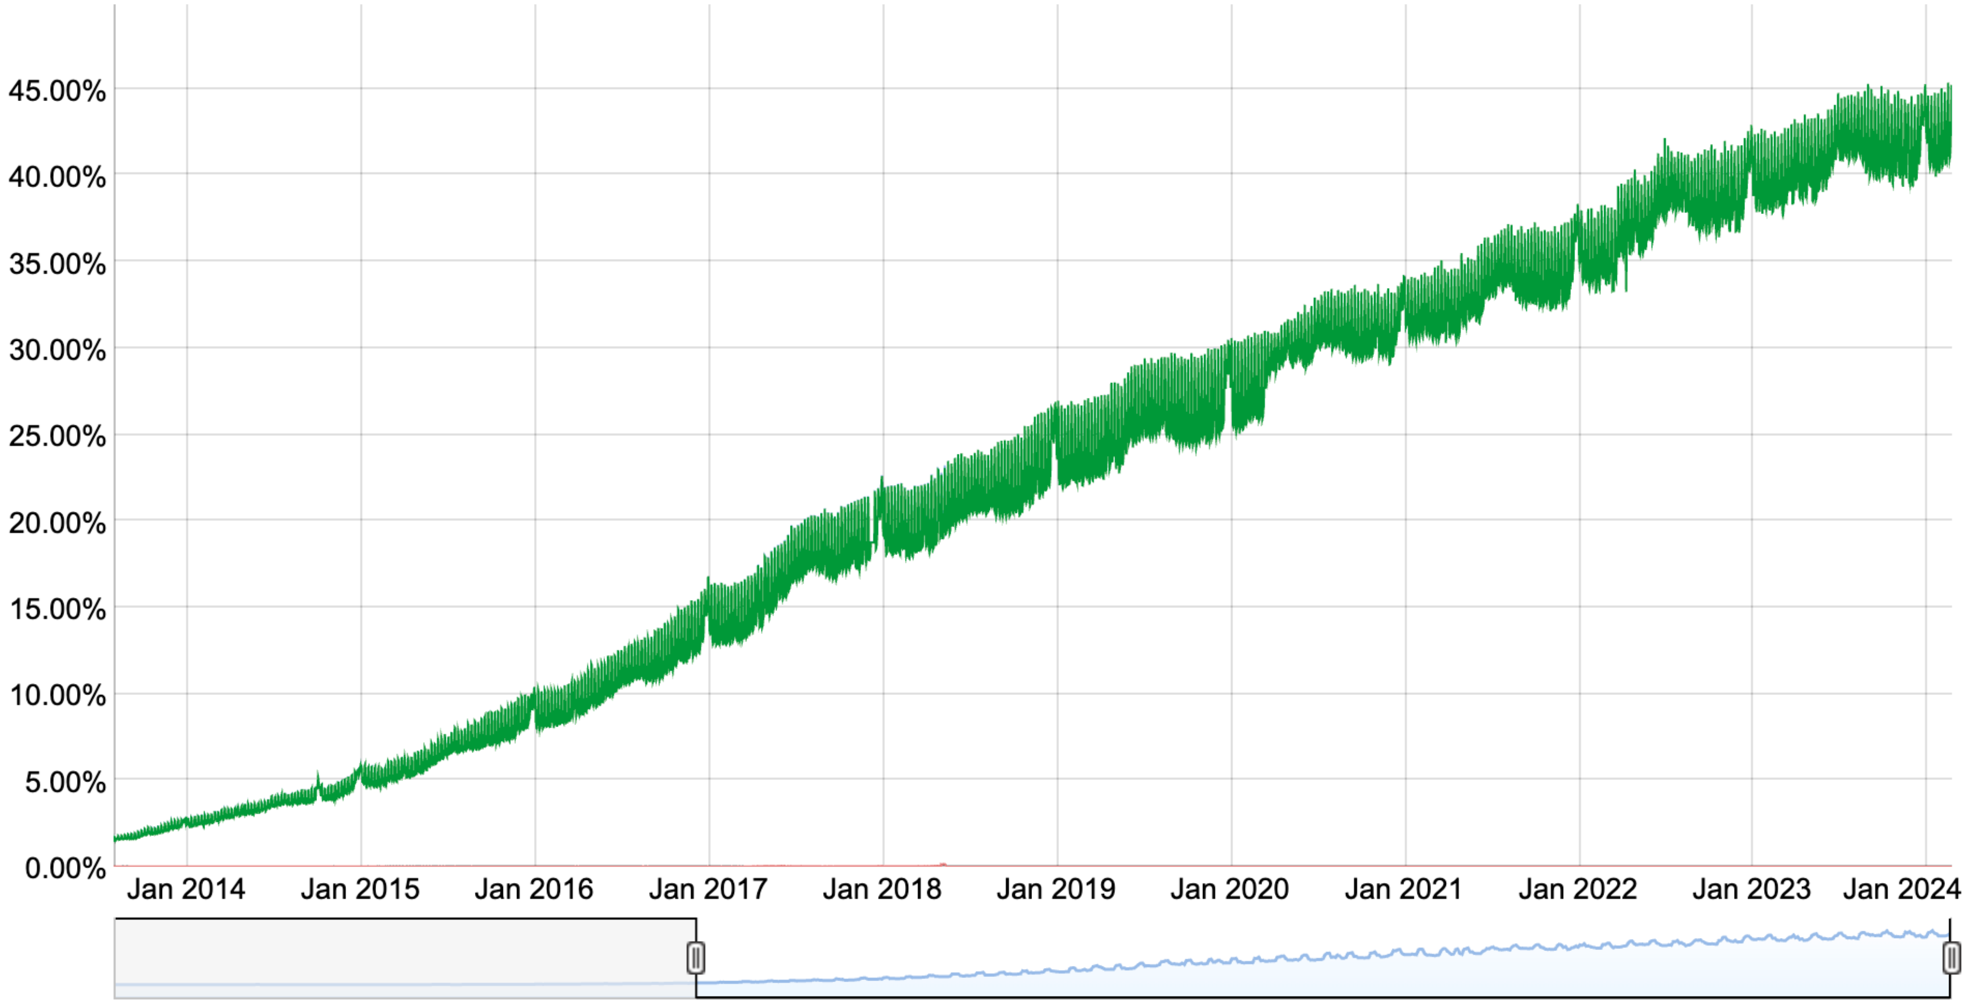

Source: https://www.google.com/intl/en/ipv6/statistics.html#tab=ipv6-adoption

Source: https://www.youtube.com/watch?v=tcae4TSSMo8

What is the problem?

- Less and less IPv4 address space available, not just from RIRs but also on the market
- Number of new internet users, connections, (IOT) devices still increasing
- Many websites/services/etc still aren’t reachable over IPv6
- Many (network) devices not IPv6 capable (yet)
- IPv4 isn’t expensive enough (yet) to make a compelling business case for mass IPv6 adoption

And there are also technical challenges

- ASICs and FPGAs are excellent in forwarding traffic (x86 does a much better job in services)
- Scaling and capacity planning is harder when multiple security services are deployed on the same box (performance impact)
- Scale up has limits to how much you can fit in a single box/rack

CGNAT

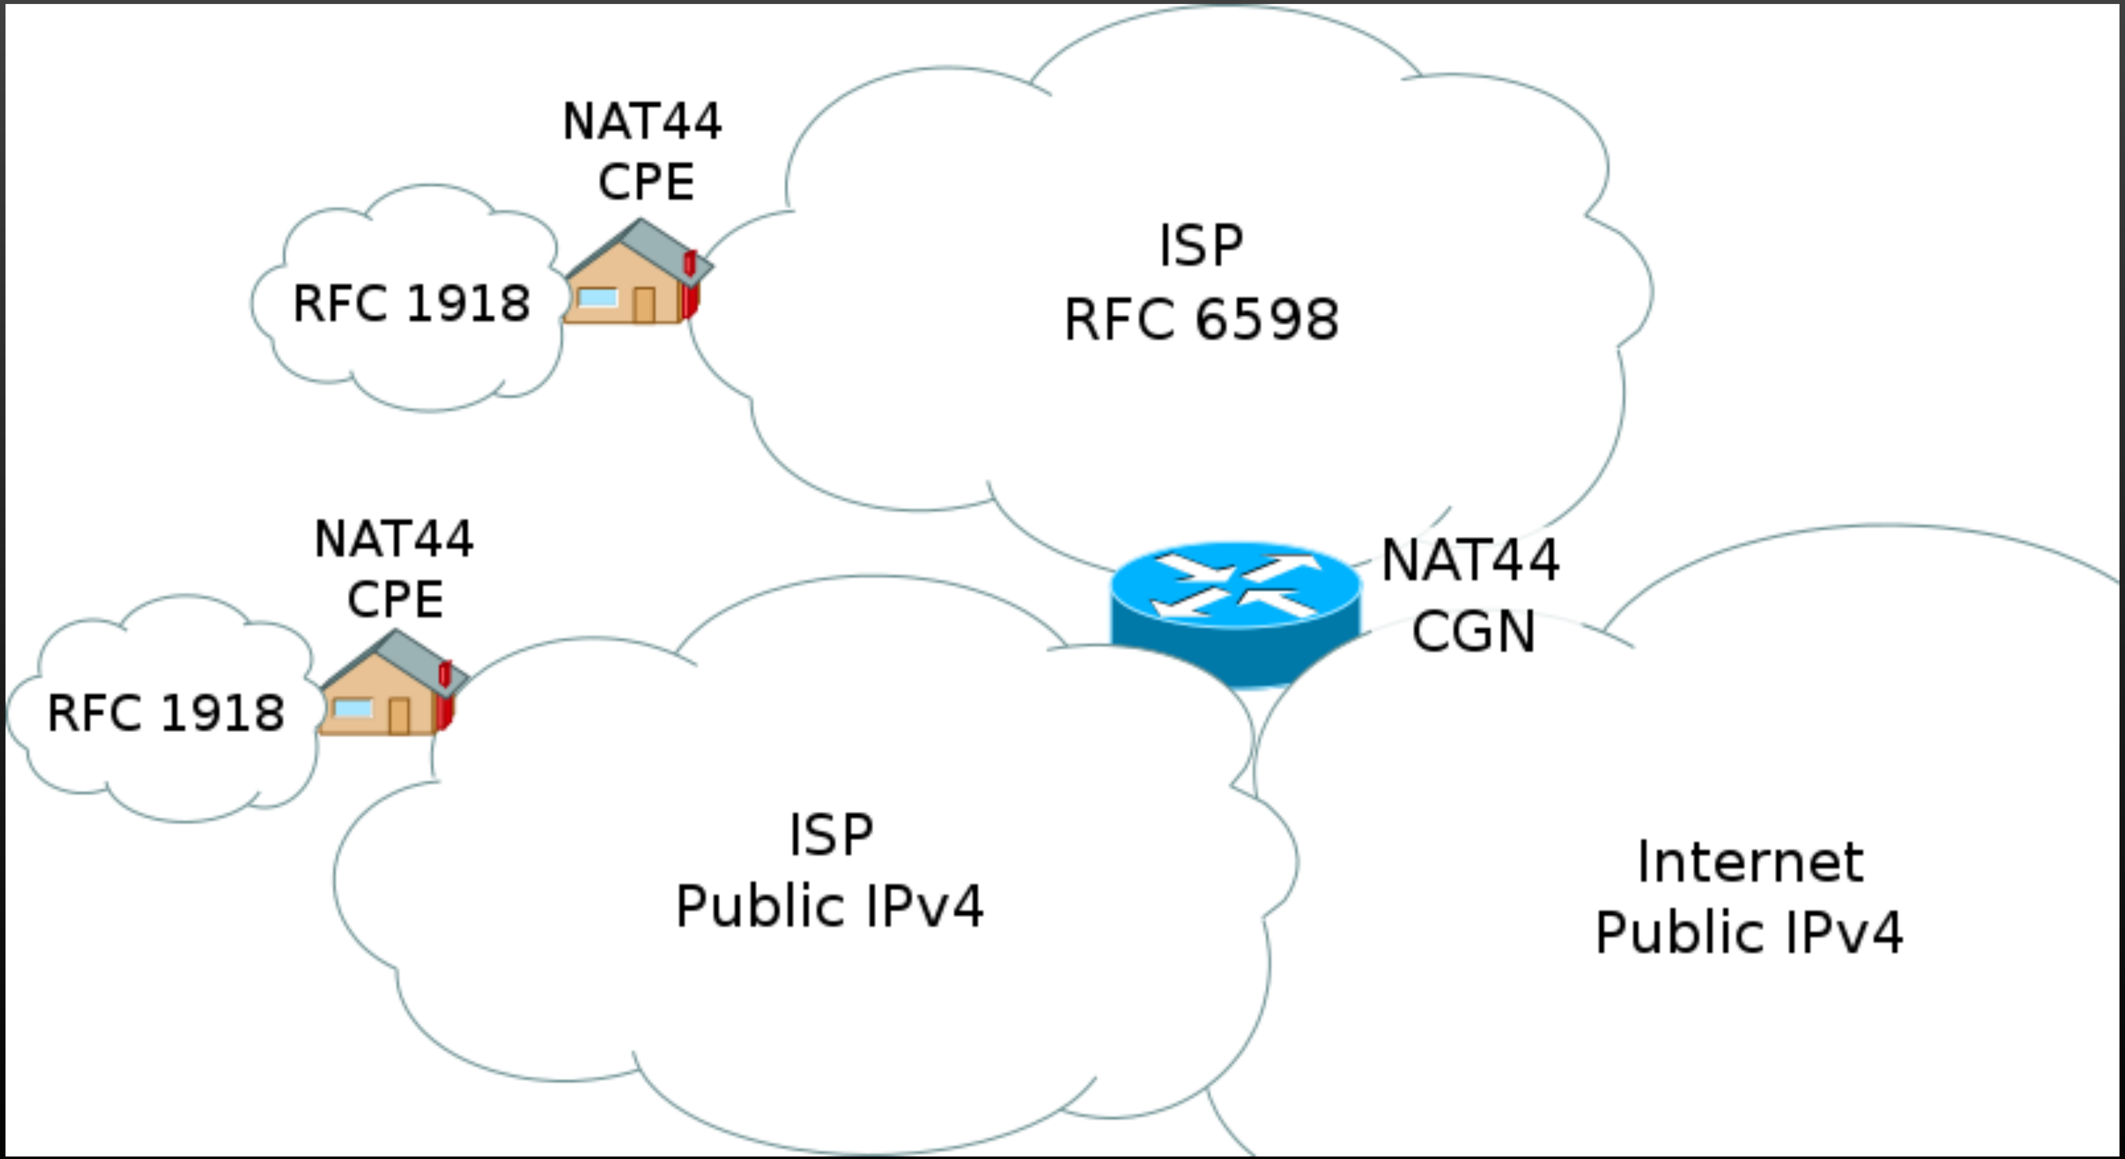

Source: https://en.wikipedia.org/wiki/Carrier-grade_NAT

DSLITE

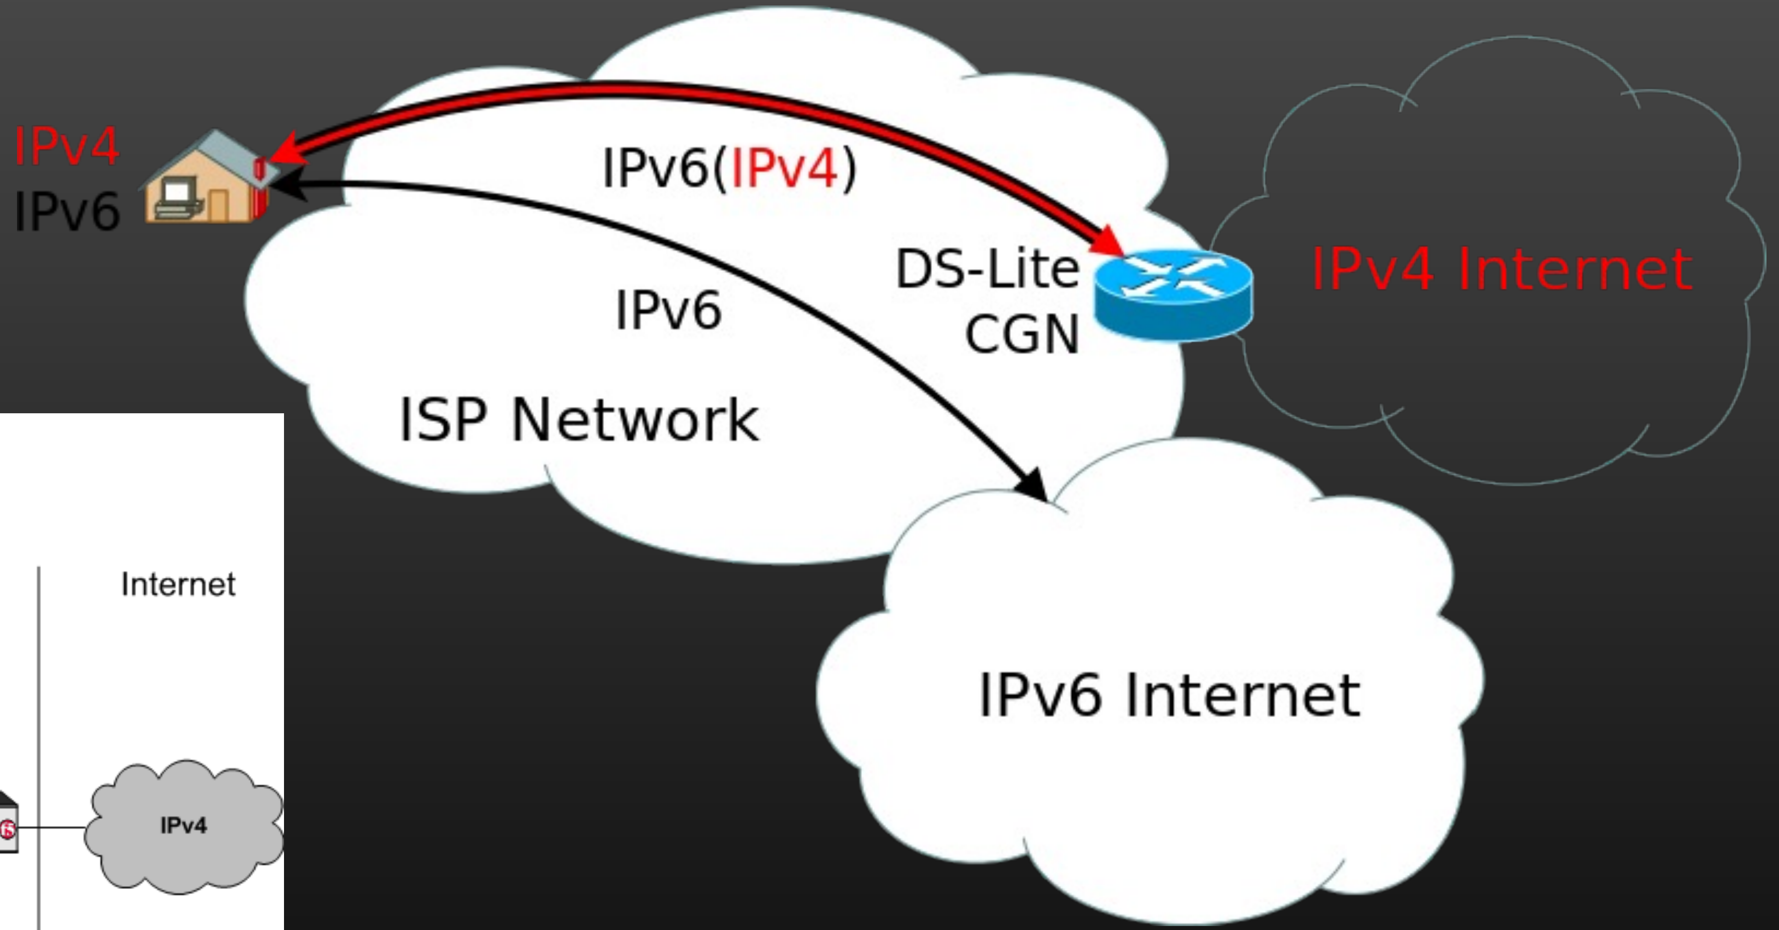

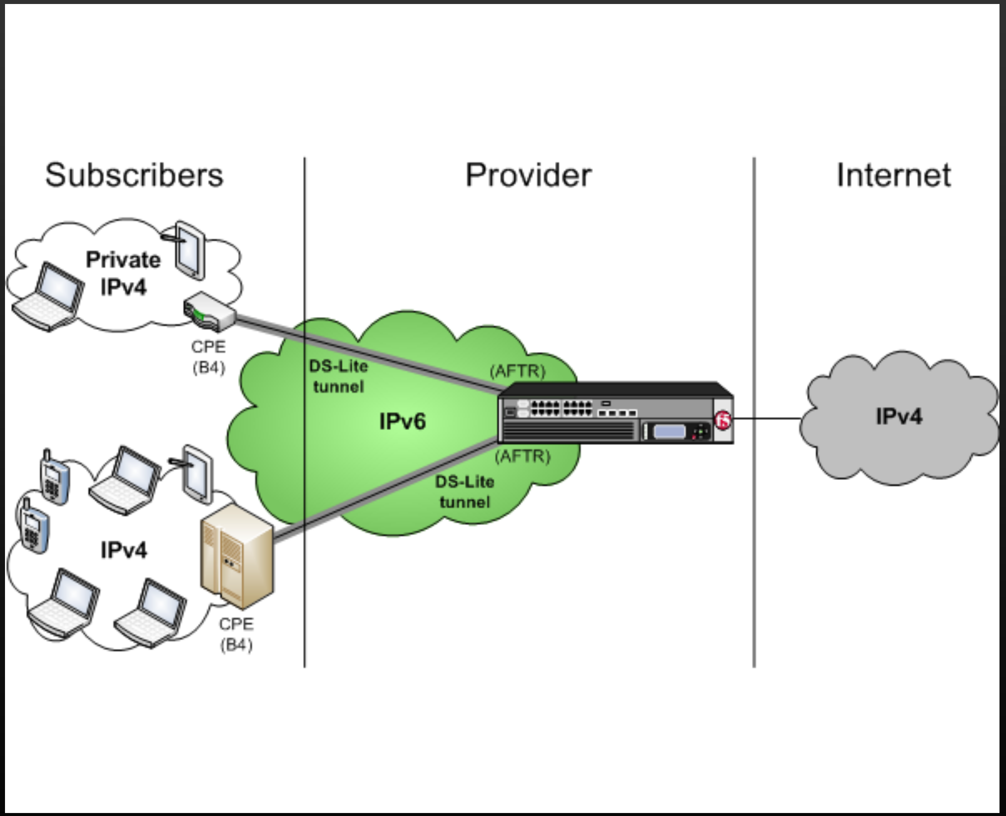

Sources: https://en.wikipedia.org/wiki/IPv6_transition_mechanism
https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-cgnat-implementations-12-1-0/4.html

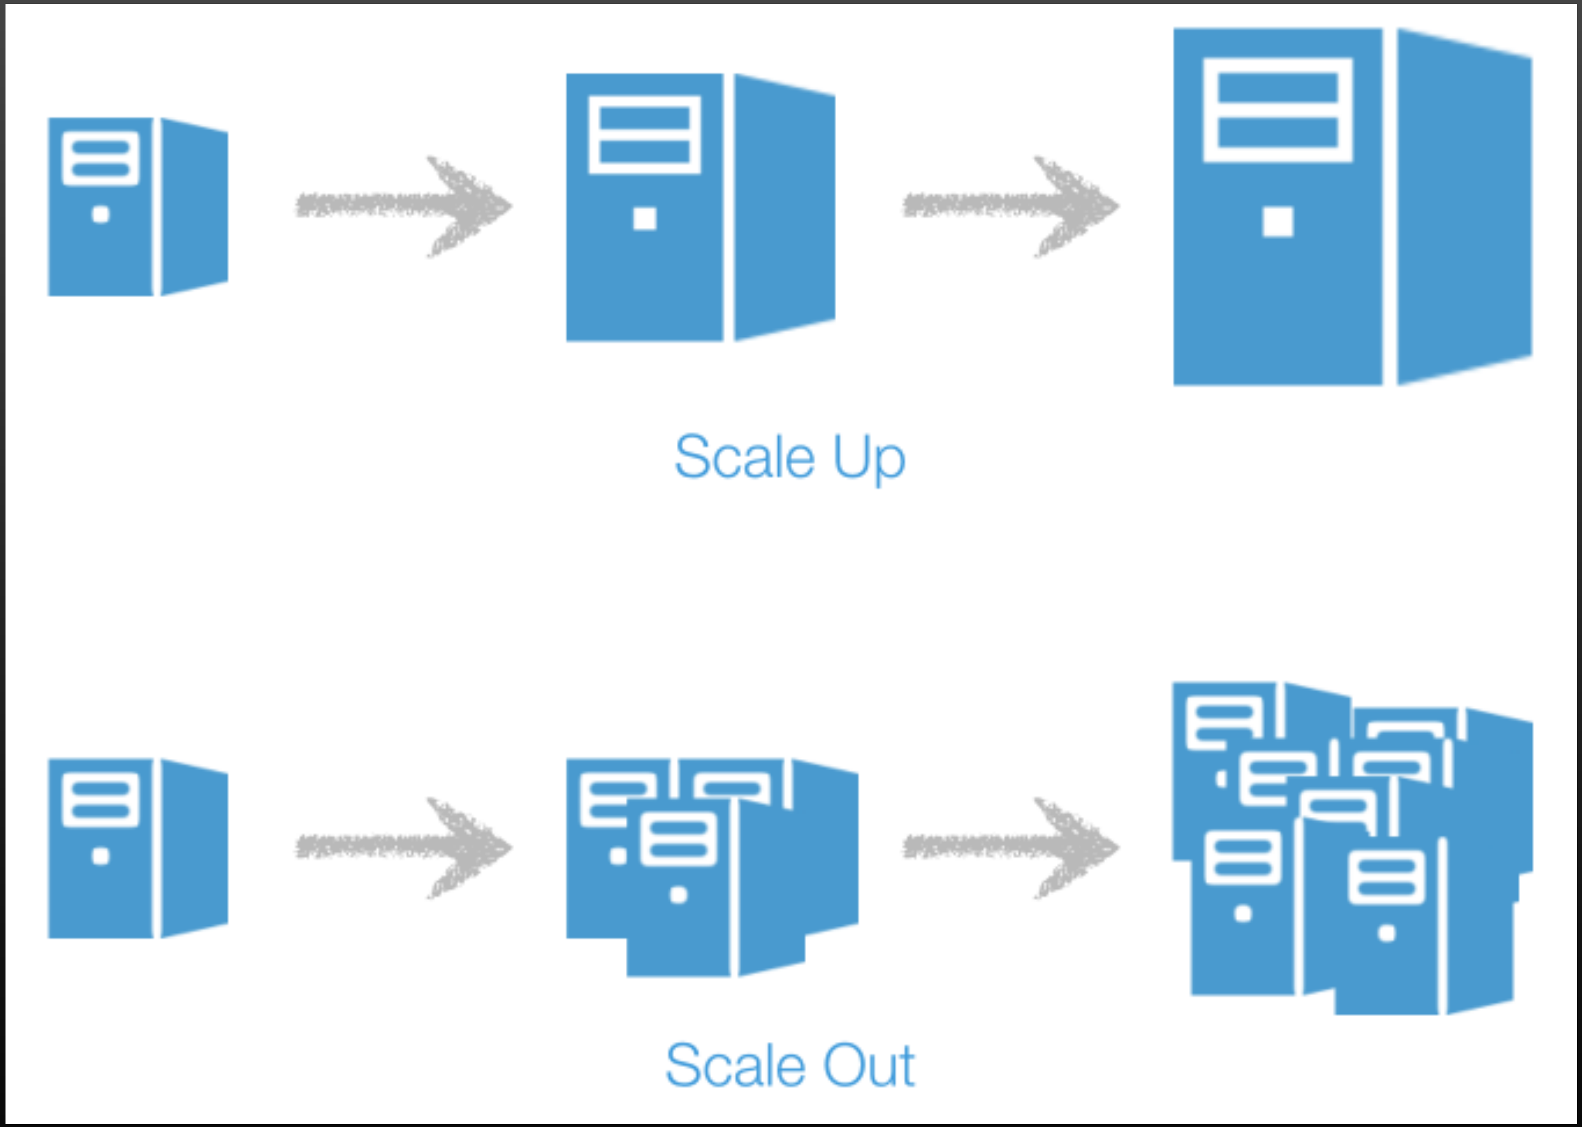

Source: https://yourbutlerspantry.com/?w=scale-up-vs-scale-out-portworx-ff-OGJAm1l6

Source: https://portworx.com/blog/scale-up-vs-scale-out/

Service Providers Need Cloud Outcomes

Simple to Build, Operate & Consume

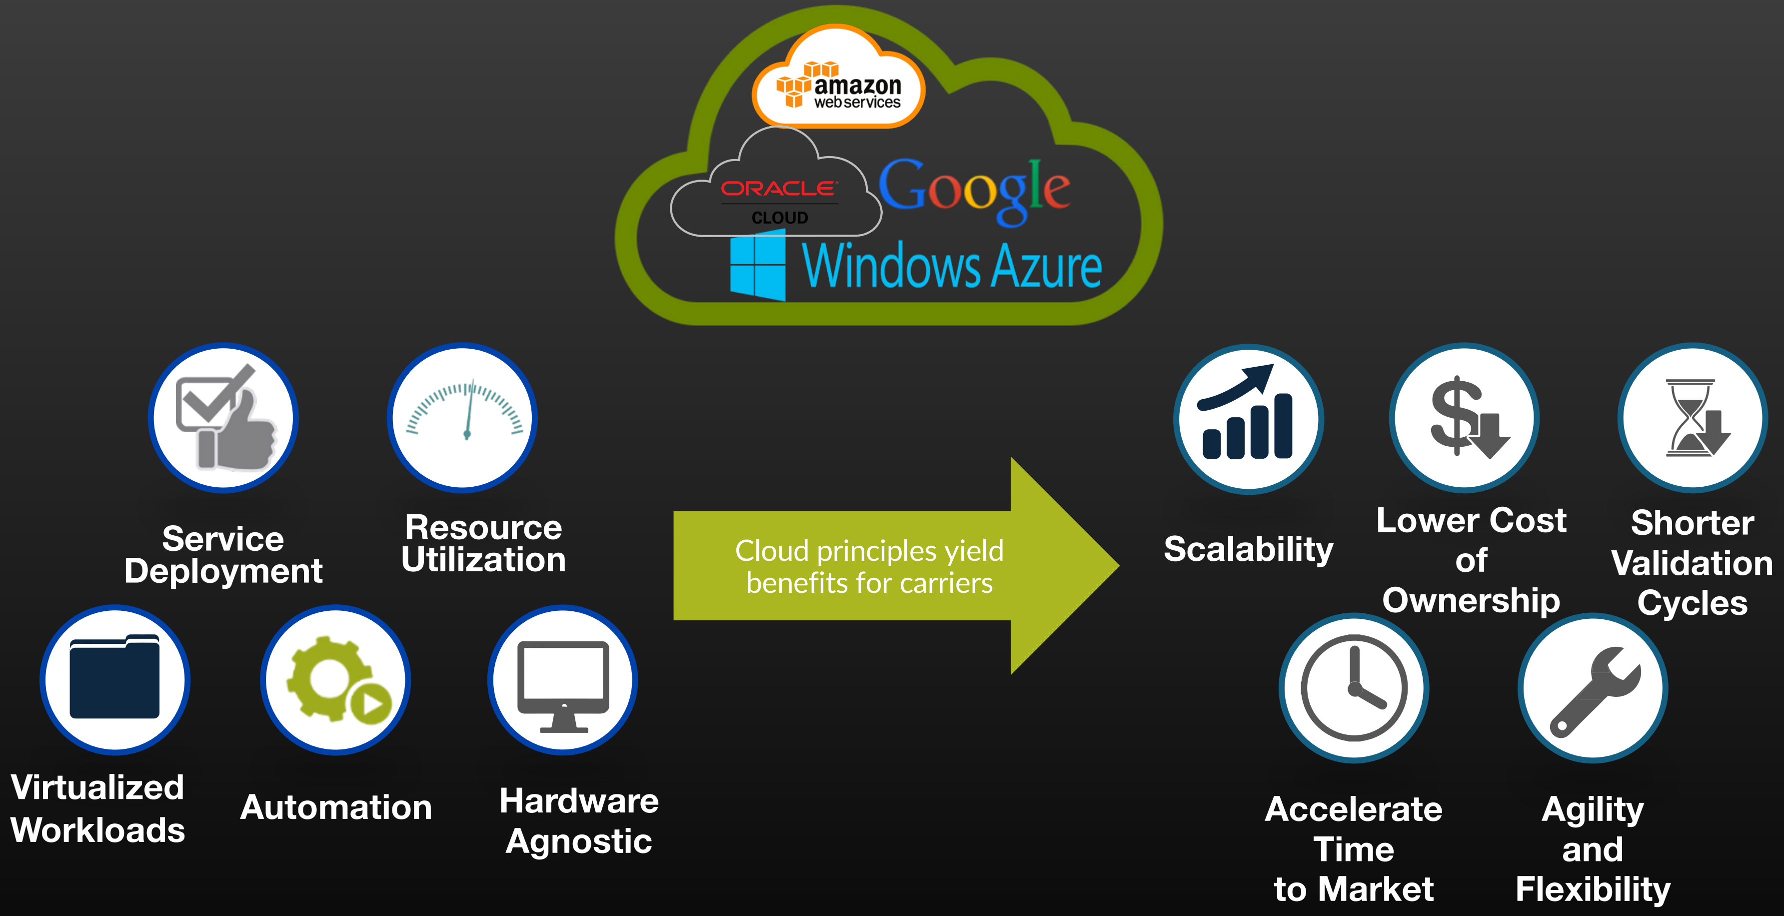

Cloud principles bring webscale capability to service providers

VNF's are not static, have a lifecycle, an evolution

EXAMPLE: The virtualization Journey of the Service Provider Edge

	1	2	3	4	5
	Disaggregation	**Integration into NFVI and SDN**	**Separation of UP and CP**	**Decomposition into micro-services ("Cloud Native")**	**Convergence between wireless and wireline**
	• Basic Virtualization of PNF to VNF • 1 VM per line card	• Horizontal NFVI	• CUPS • Edge Cloud	• Containers • Stateless design • Service mesh	• NFVI for wireline and wireless access and services
Business Driver	• Vendor unlock • Cost • Agility	• telco transformation • virtualization skills	• leverage key assets • build the Edge Cloud for IoT and 5G	• agility to deploy new services • distribute them • life-cycle management • lower TTR.	• rationalize assets • reduce cost • Convergence of services
Role of the NFVI & SDN	• BMS	• high performance • multi-VNF • multi-tenancy	• scale out, • network slicing	• Networked containers • Mixed VMs and containers • Isolation, security of containerized applications.	• high performance virtualization • BMS and PNF control via SDN

Key Trends & Lessons Learned

Each VNF vendor has its own history, approach

- Chassis to fat VMs approach is common
- There is experience now on how these VNF vendors perform (vEPC, vIMS…)

Efficient containerization and micro-services will make the difference between VNF vendors

- NFVI and SDN stack become even more critical: requires fully featured virtual networking, high performance, high availability, elasticity with high scale; requires support of both VM, containers and BMS.
- Stateless design means better high availability, easier scale-out, better alignment with CI/CD

NFVI and SDN stacks need to be versatile

- integrate all requirements of the main use-cases (vEPC, vIMS, vGiLAN, vRAN) from the main VNF vendors
- Lots of new discoveries and unknowns

NFVI + SDN + Testing of stack : integration has an infinite number of permutations

- Risk of falling back into vertical stack
- Building blocks + CI/CD is fundamental to success

Scale Out Services Solution – several options

- ✓ Many compact & modular routers and VNFs supported
- ✓ Security Services: Carrier Grade NAT, IPsec, Stateful Firewall Services, etc.
- ✓ Offers a stateless implementation of Security Services
- ✓ Positioning: Enhancing Provider Edge infrastructure with Security Services
- ✓ Combines router with virtual firewall VNF to enable complex services
- ✓ Automation tools for VNF onboarding/configuration
  - ✓ Automation is provided as part of the solution and is available for customization to meet customer specific deployment scenarios.
- ✓ Bring Your Own Server, Cloud or packed solution with x86
- ✓ Scalable service clusters

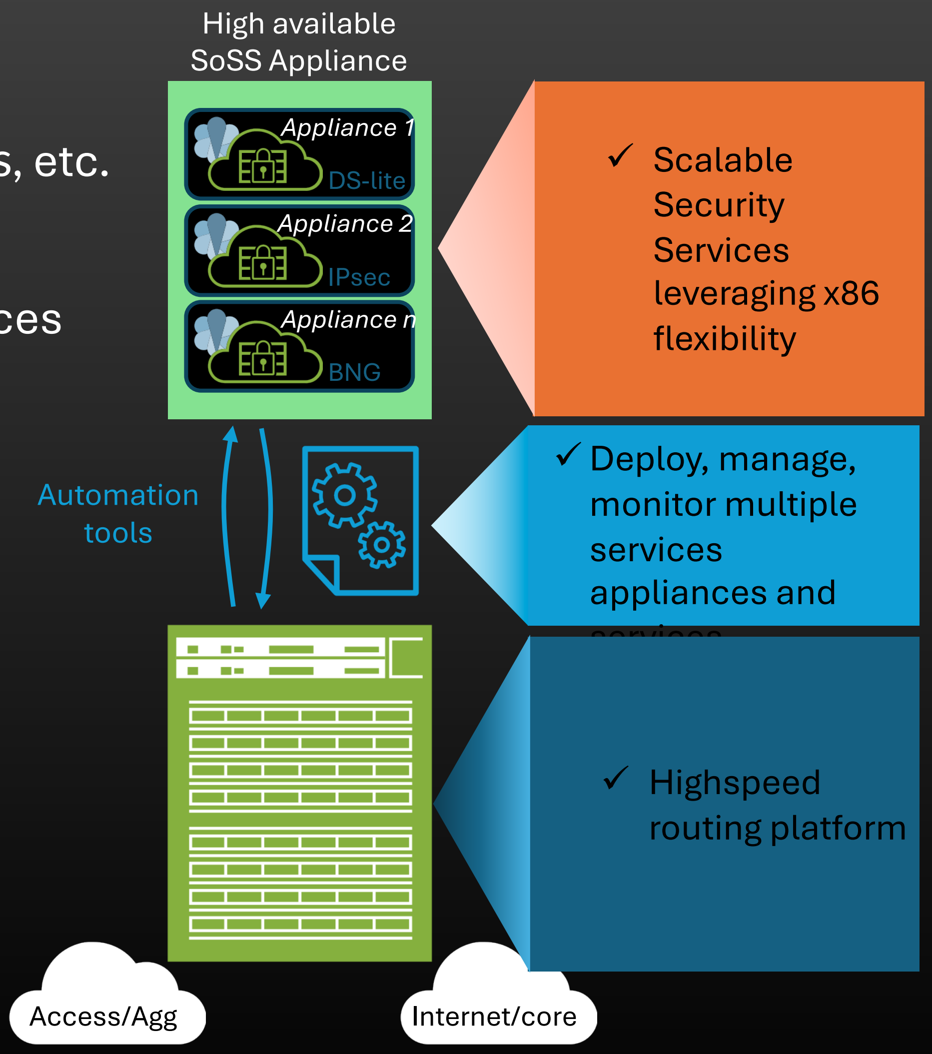

Data Plane “Integration”

- Data plane integration leveraging ECMP, BGP, BFD, MPLS
  - Load balancing using ECMP
  - ECMP Consistent hashing (needed for >1 VNF
  - Hashing at ingress card (use source-IP hash)
  - Stateless

- Option for load distribution possible, for example per vAFTR FQDN based DS-Lite tunnel distribution

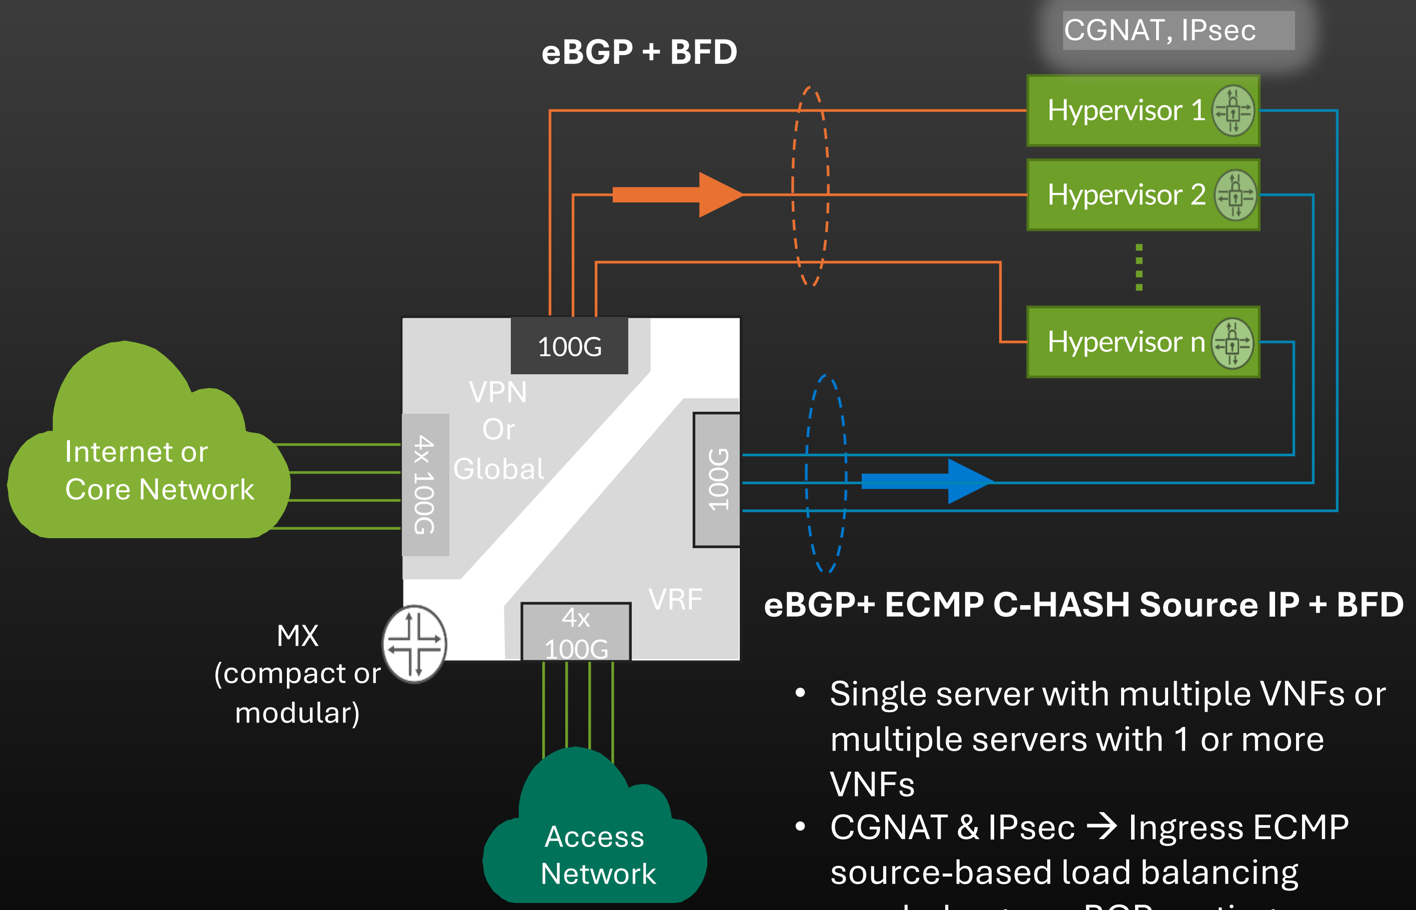

- Single server with multiple VNFs or multiple servers with 1 or more VNFs
- CGNAT & IPsec → Ingress ECMP source-based load balancing needed, egress BGP routing
- Both sides ECMP required only for FW service with no NAT

Per Customer based slicing globally

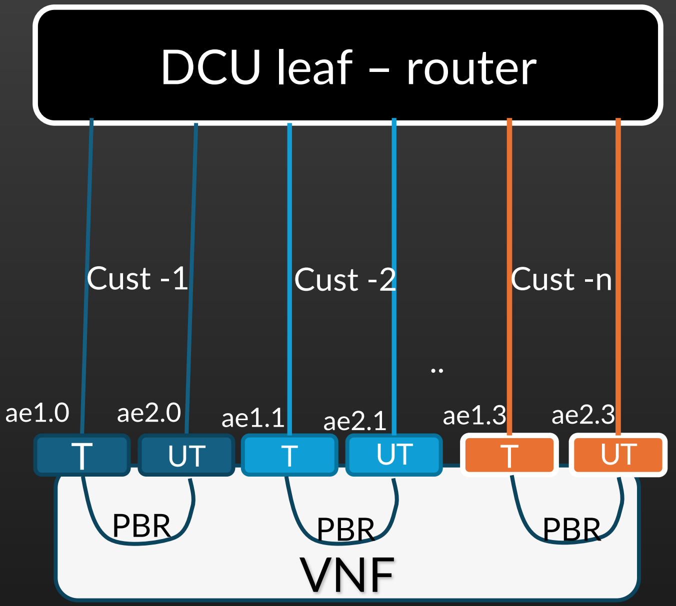

What do we need on VNF ?

- 2 zones
- 2 routing instances
- 2 IFLs each mapped to a routing instance
- 2 firewall filters and bound to interface in RI on ingress with next routing instance of other
- Security policies wrt the zones
- Default route in each RI with NH pointing to DCU leaf
- SNAP would need to take customer related information and translate to the configuration pertaining to vSRX

Advantages

- Simple
- No need for additional license apart from std/adv1/adv2 license
- Can scale well
- Per customer logging based on dataplane. All logs are carried over the same interface to collector
- Per customer slicing with dedicated RI / Zones / IFL / Policies

Disadvantages

- Global view and not per customer view from management
- Lots of configuration parameters to manage with IFL / IRB / vLANS , static routes..etc
- Need Enhancement to support consistent hashing/ symmetric flow for static route design in the MX

Split customers using MP-BGP (Per Customer VRF)

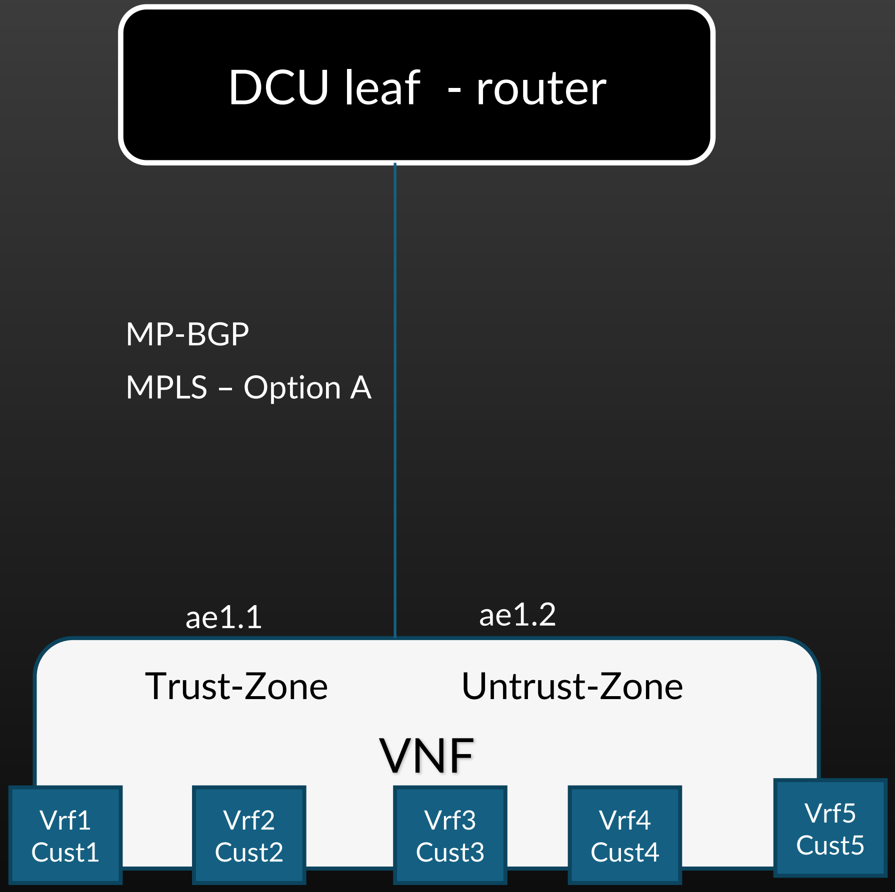

What do we need on the VNF ?

- Configure VRF per customer to hold customer route / Default route
- Configure MP-BGP MPLS
- MP-BGP MPLS Option A connectivity to DCU
- SNAP would need to take customer related information and translate to the configuration pertaining to vSRX

Advantages

- Eliminates the need for multiple IFLs, IRBs, security zones, Static routes, Vlans, BFD
- Simple configuration with MP-BGP
- MP-BGP maintains all the routing dynamically
- MPLS label mapping places customer traffic in the respective customer VRFS
- Can scale well and achieve scale out with already available CHASH/Symmetric flow feature
- Per customer logging based on data plane. All logs are carried over the common interface to a log collector

Limitation

- All customer traffic will pass through the same IFL and security zone
- Customer separation is via VRF from forwarding perspective
- Customer multi tenancy relies on the security policy
- Customer network routes need to be advertised to vSRX

Redundancy with ECMP Next-hops over 2 MP-BGP Connection - Server to Client Route Advertisement

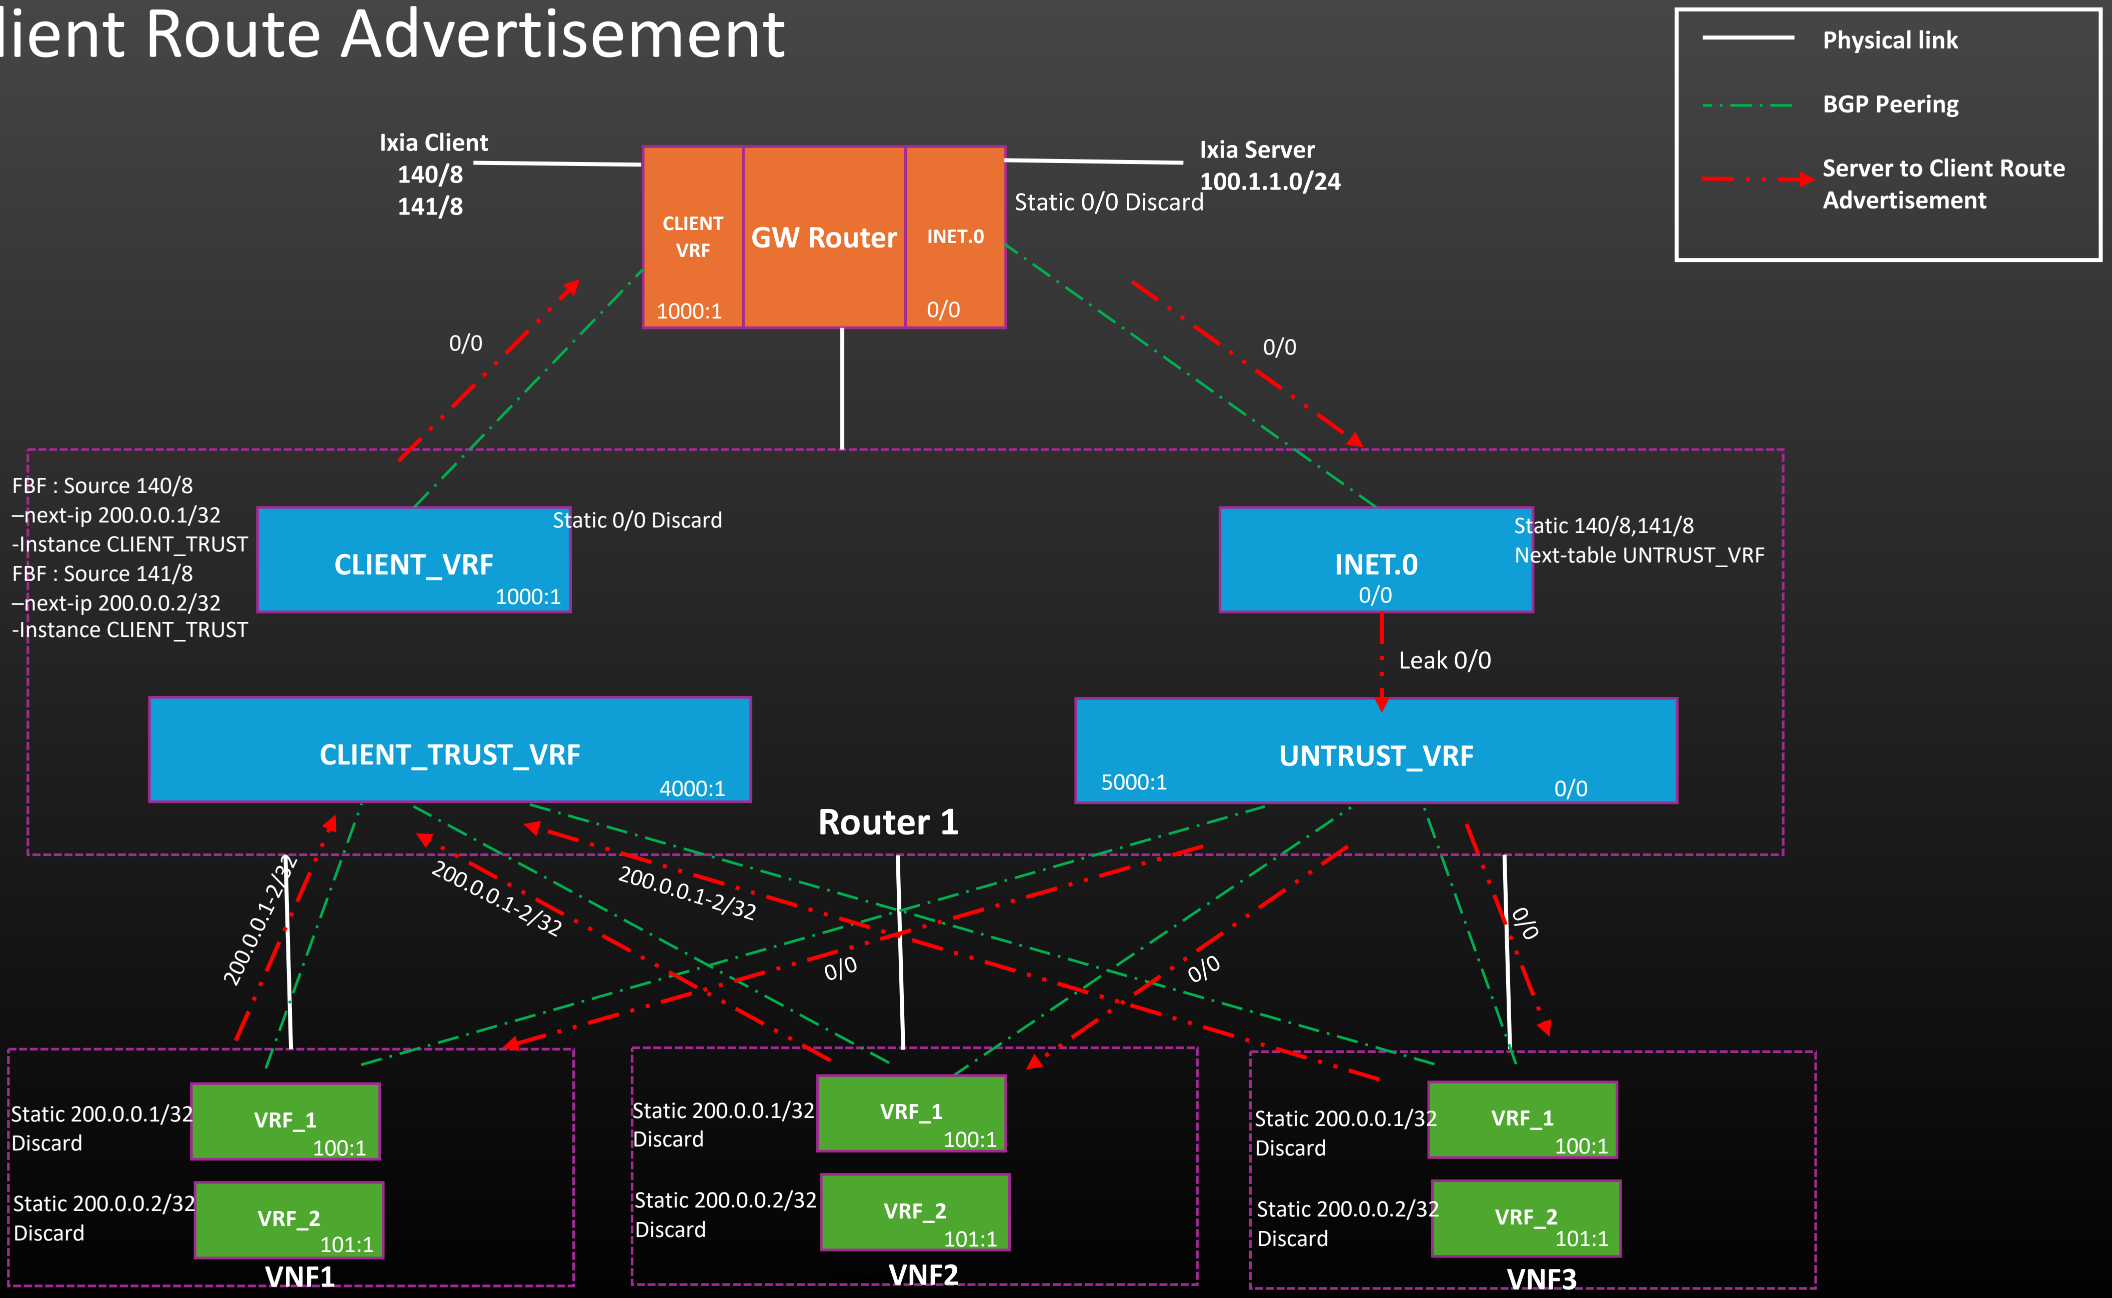

Redundancy with VNF ECMP Next-hops over 2 MP-BGP Connection - Client to Server Route Advertisement

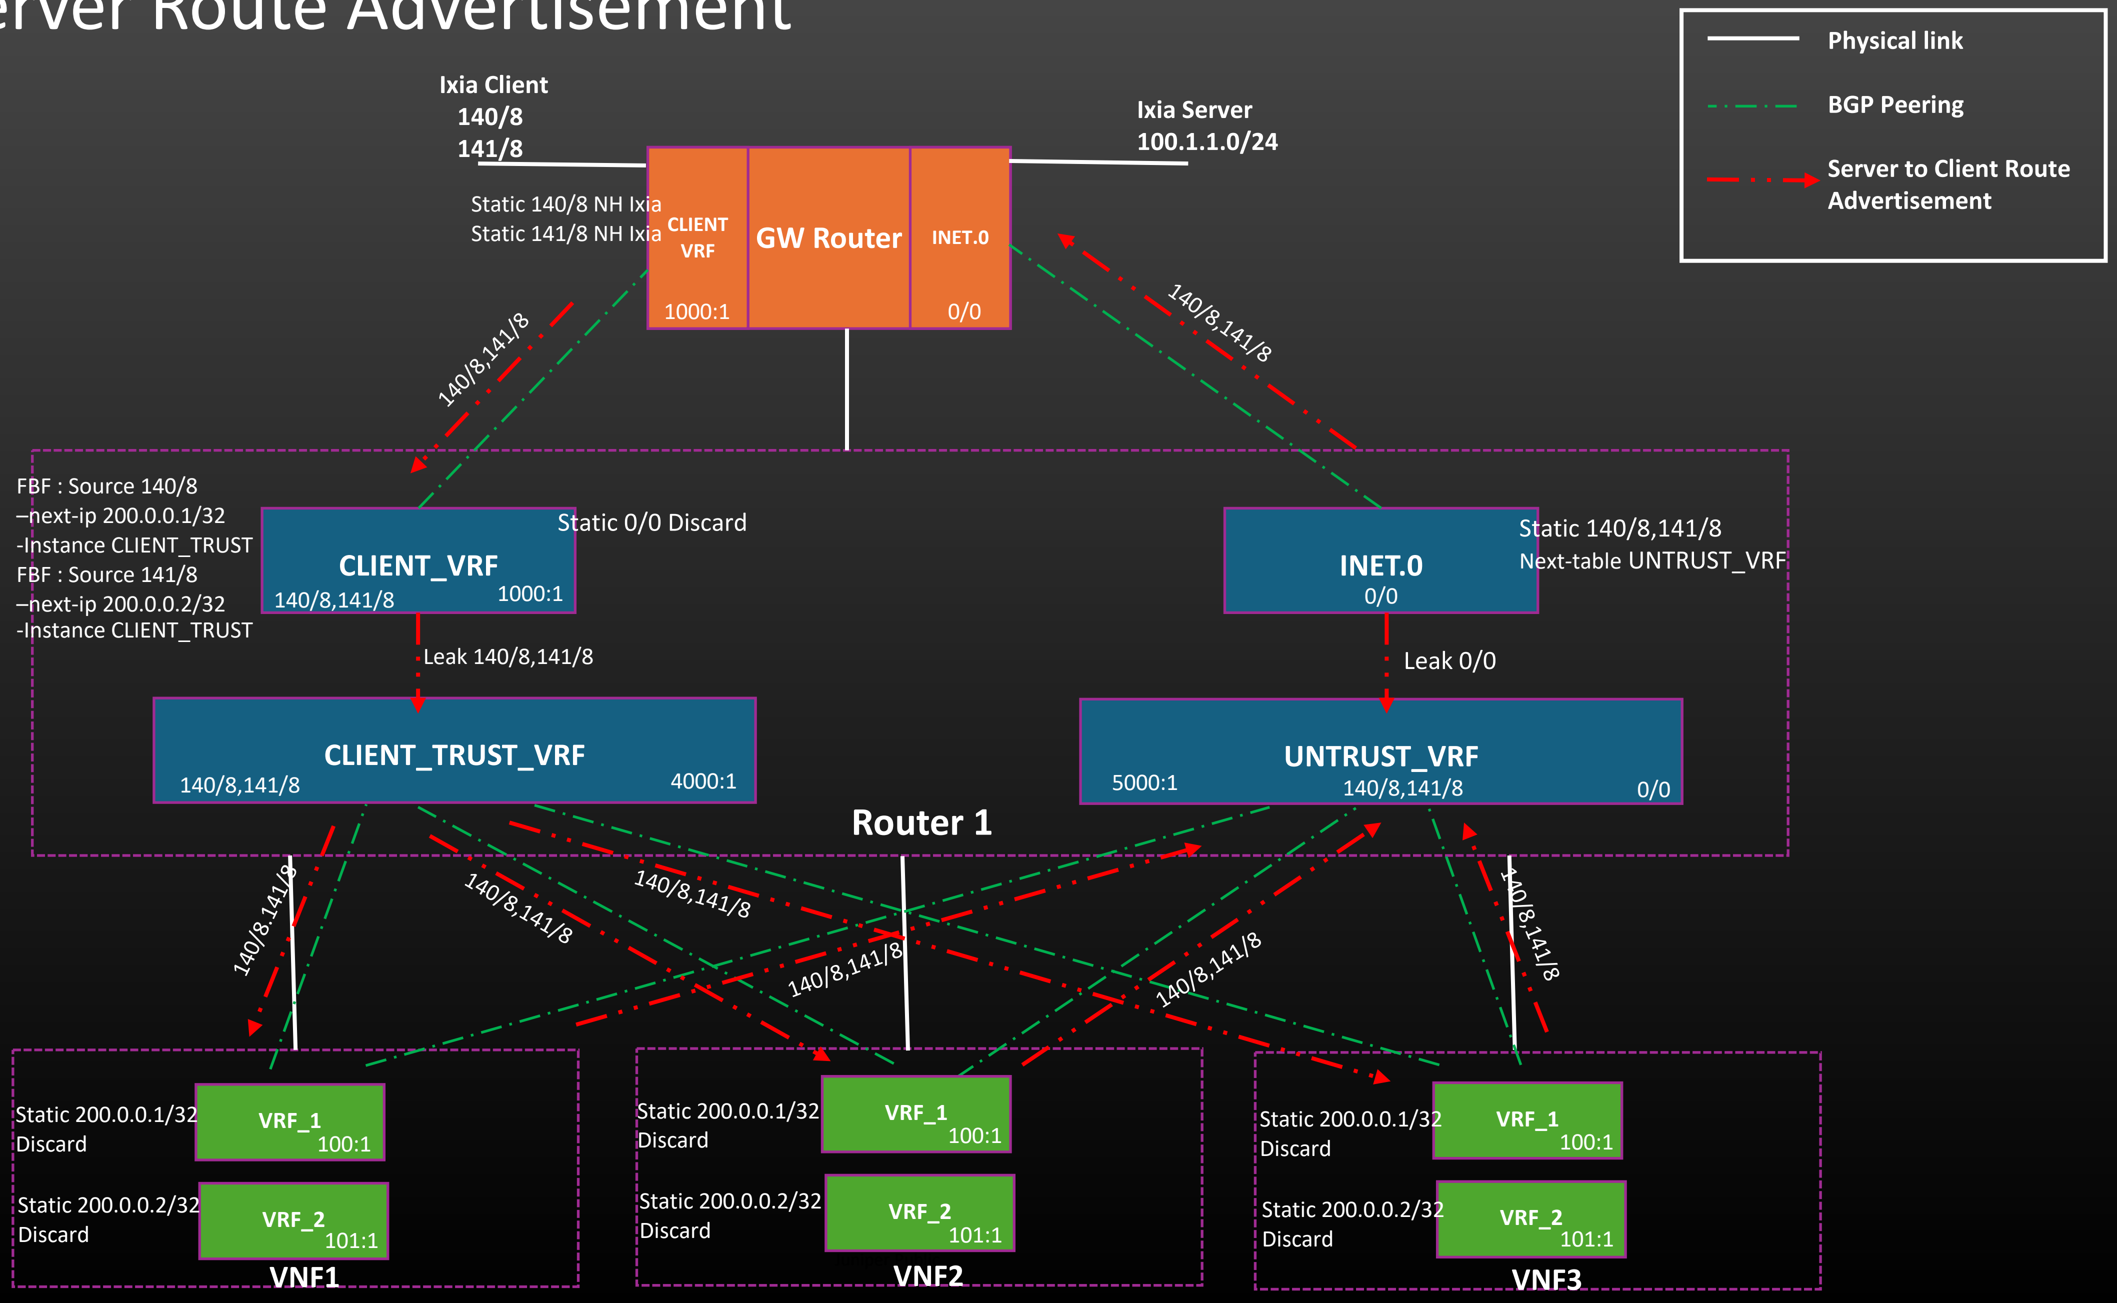

Proof of Concept with AMD CPUs

PowerEdge R7625	
Processor:	AMD EPYC 9654 2.40GHz, 96C/192T, 384M Cache (360W) DDR5-4800
Additional Processor:	AMD EPYC 9654 2.40GHz, 96C/192T, 384M Cache (360W) DDR5-4800
Processor Thermal Configuration:	Heatsink for 2 CPU with GPU configuration
Memory Configuration Type:	Performance Optimized
Memory DIMM Type and Speed:	4800MT/s RDIMMs
Memory Capacity:	24 32GB RDIMM, 4800MT/s Dual Rank
Additional Network Cards:	Broadcom 5720 Dual Port 1GbE Optional LOM 2 Mellanox ConnectX-6 DX Dual Port 100GbE QSFP56 Network Adapter, Low Profile 4 Mellanox ConnectX-6 DX Dual Port 100GbE QSFP56 Network Adapter, Full Height

Example: CGNAT and DS-Lite with vSRX

- Single instance of vSRX 17 and 31 vCPUs
- 28M Sessions concurrently active in all tests
- With PBA – 1 block of 500 Ports utilized
- With DetNAT – 1 block of 500 Ports utilized
- With DSLITE – 28K tunnels each with 1000 sessions

vSRX Model	Memory	Junos Version
vSRX 17 vCPU	64G	22.2R2.2
vSRX 31 vCPU	96G	23.1R1.1

vSRX Size and Packet Size	Sessions	Throughput in Gbps				
		NAT44	NAT64	DetNAT	PBA	DSLITE
17 vCPUs + 64G Memory - IMIX 908 Bytes	**28M**	**82.5**	**71.8**	**81.5**	**81.6**	**51.2**
31 vCPUs + 96G Memory - IMIX 908 Bytes	**28M**	**100.2**	**87.2**	**97.4**	**96.8**	**82.6**

DS-Lite Perf in Gbps	vSRX with 17 vCPUs			vSRX with 31 vCPUs		
	# of vSRX's	# of Single Socket Servers	# Dual Socket Servers	# of vSRX's	# of Single Socket Servers	# Dual Socket Servers
100 Gbps	2	2	1	2	2	1
200 Gbps	4	4	2	3	3	2
300 Gbps	6	6	3	4	4	2
400 Gbps	8	8	4	5	5	3

THANK YOU

Melchior Aelmans
maelmans@juniper.net