Metaswitch Networks

THE BRAINS OF THE NEW GLOBAL NETWORK

IEEE CQR

#IEEEComSoc International Communications Quality and Reliability Workshop

9-12 May 2016
Stevenson, WA • USA

TELCO-GRADE SERVICE AVAILABILITY FROM AN IT-GRADE CLOUD

MARTIN TAYLOR
CHIEF TECHNICAL OFFICER

MAY 9, 2016

Context: NFV

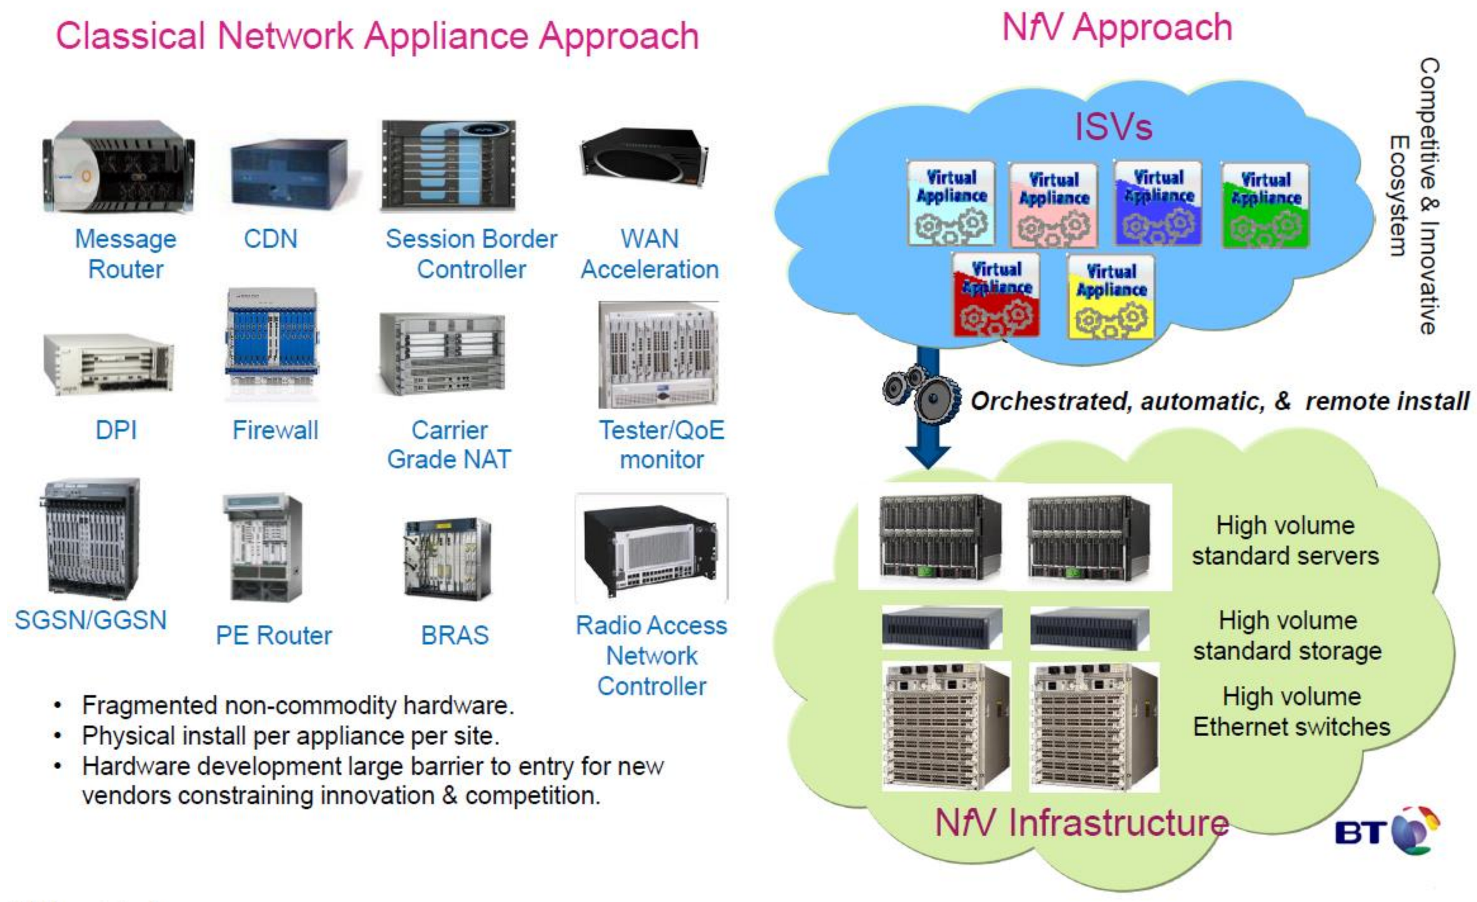

Classical Network Appliance Approach
Message Router
CDN
Session Border Controller
WAN Acceleration
DPI
Firewall
Carrier Grade NAT
Tester/QoE monitor
SGSN/GGSN
PE Router
BRAS
Radio Access Network Controller
• Fragmented non-commodity hardware.
• Physical install per appliance per site.
• Hardware development large barrier to entry for new vendors constraining innovation & competition.
NfV Approach
ISVs
Virtual Appliance
Virtual Appliance
Virtual Appliance
Virtual Appliance
Virtual Appliance
Virtual Appliance
Competitive & Innovative Ecosystem
Orchestrated, automatic, & remote install
High volume standard servers
High volume standard storage
High volume Ethernet switches
NfV Infrastructure
BT

Telco Availability Requirements

Telcordia Technologies Generic Requirements
GR-512-CORE
Issue 2
January 1998

Requirements for Local and Tandem Switching Systems

… **Table 3.** Total Capability Downtime Requirements

Capability Type	Requirement
Analog Line	0.4 minutes/year
ISDN Circuit Switching	0.4 minutes/year
ISDN Packet Switching	0.4 minutes/year

… **Table 4.** Hardware Cutoff Call Rate Requirements

Termination Type	Requirement
Analog Line	15,000 cutoffs per 10^9 hours of call duration
Analog Trunk	15,000 cutoffs per 10^9 hours of call duration
Digital Trunk	10,000 cutoffs per 10^9 hours of call duration

Equivalent to > 99.9999% service availability

Equivalent to mean interval between trunk failure events > 11 years

Availability of IT-grade Servers

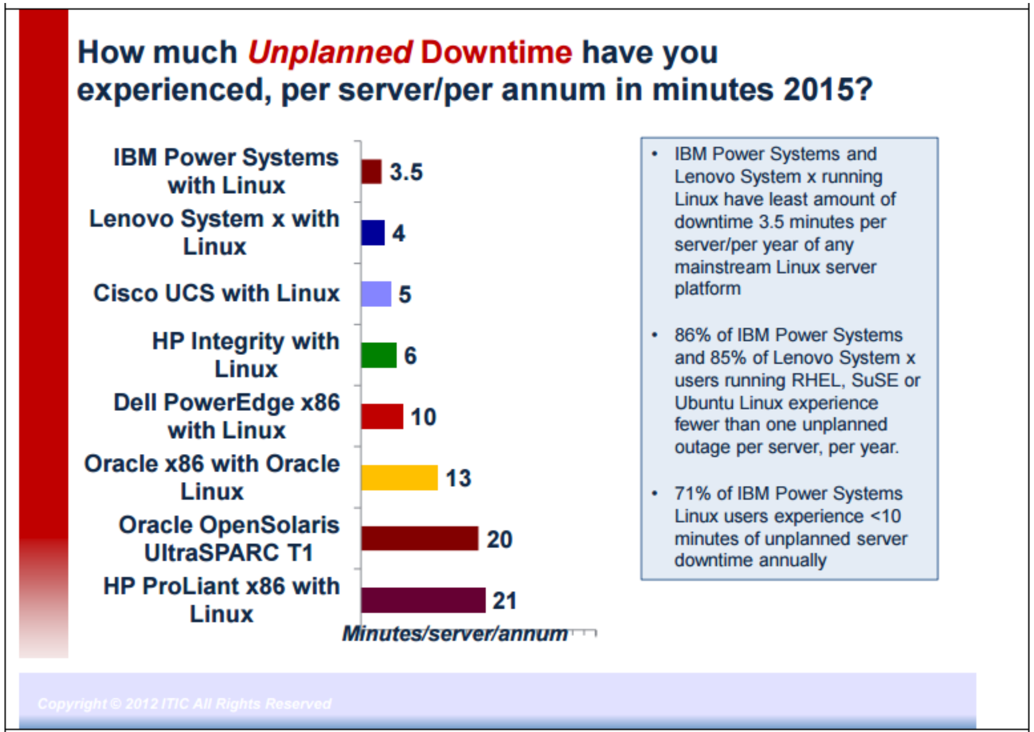

Source: ITIC 2015

Server + Linux OS achieving in the range 99.995 – 99.999 % availability

Availability of DC Infrastructure

UptimeInstitute®

Tier 1	Basic site, no redundant infrastructure	99.671%
Tier 2	Redundant common equipment	99.741%
Tier 3	Redundant power and cooling delivery	99.982%
Tier 4	Cooling equipment redundantly powered	99.995%

Availability of Public Cloud Services

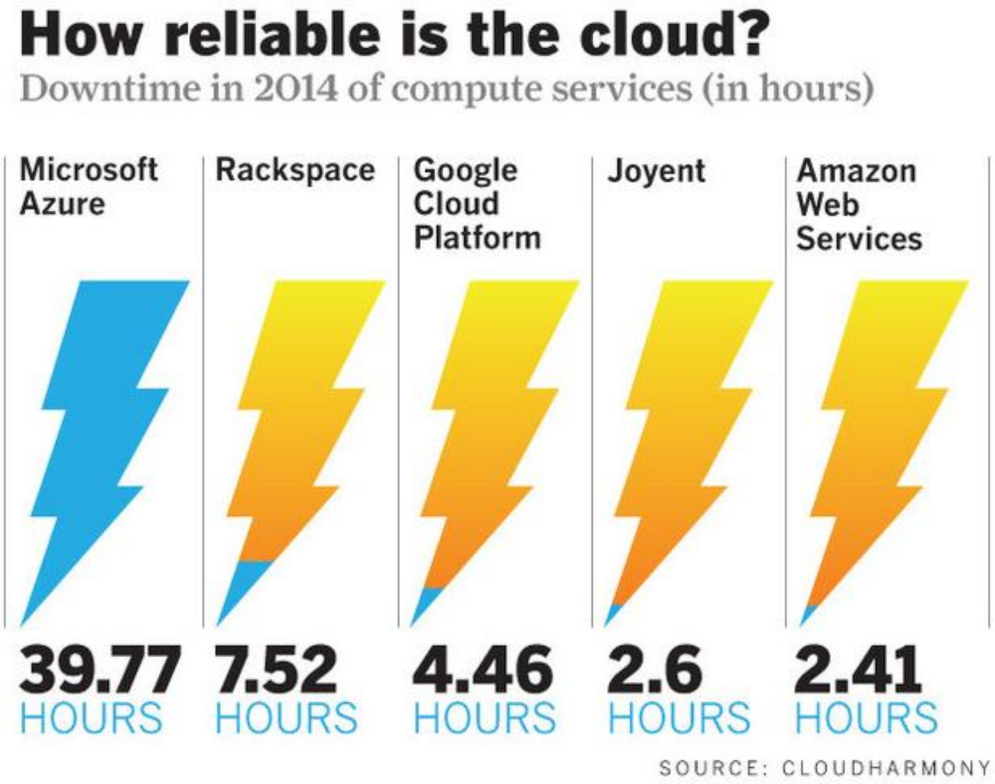

Best was equivalent to 99.97%

Availability of OpenStack

- Not much detailed analysis in the public domain
- Anecdotal evidence (e.g. presentations at OpenStack Summits) suggests ~ 99.95%
- Question is complex because there are different modes of failure
  - Control plane can go down without impacting user plane
- We have seen detailed analysis suggesting user plane availability of 99.96 – 99.97%

Meeting Telco-grade Objectives

- The NFV "stack" comprises many elements, none of which achieves > 99.999% availability
- A telco-grade service must not be vulnerable to the failure of a single instance of any element in the stack
- We would obviously expect to deploy the service across redundant compute nodes and redundant data centers
- The cloud environment is almost certainly the weakest link

We cannot escape the conclusion that a telco-grade service must be deployed across multiple independent and redundant cloud instances

Two Approaches to Telco-grade

Entire service in one HA cloud

Service deployed across redundant clouds

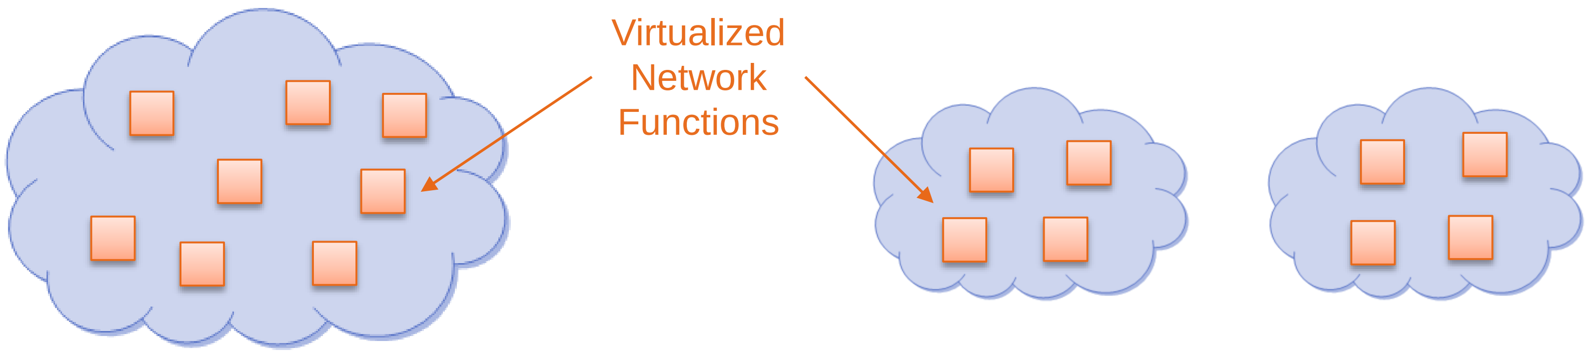

For a five-nines service, we need a <u>six-nines</u> cloud

For a five-nines service, we only need a <u>three-nines</u> cloud

Redundant Clouds: Shared-Nothing

Sharing **anything** between cloud instances introduces a form of coupling that can propagate failures

Google Cloud Platform

Google Compute Engine Incident #16007

SUMMARY:

On Monday, 11 April, 2016, Google Compute Engine instances in all regions lost external connectivity for a total of 18 minutes, from 19:09 to 19:27 Pacific Time.

Caused by propagation of corrupted route configuration between regions

Key Lesson

"Coordination's friend is contagion"

A costly investment in redundant Tier 4 data centers can be completely undone by failures that propagate through the cloud

Shared-Nothing OpenStack Example

2016 OpenStack Summit Austin

Alan Meadows – Scaling OpenStack with a shared nothing architecture

https://www.youtube.com/watch?v=FHwxbIOX_Iw

Improves resiliency with additional advantages for

- Flexibility
- Upgrades and updates
- Performance and scaling
- Design complexity

ETSI NFV Architecture

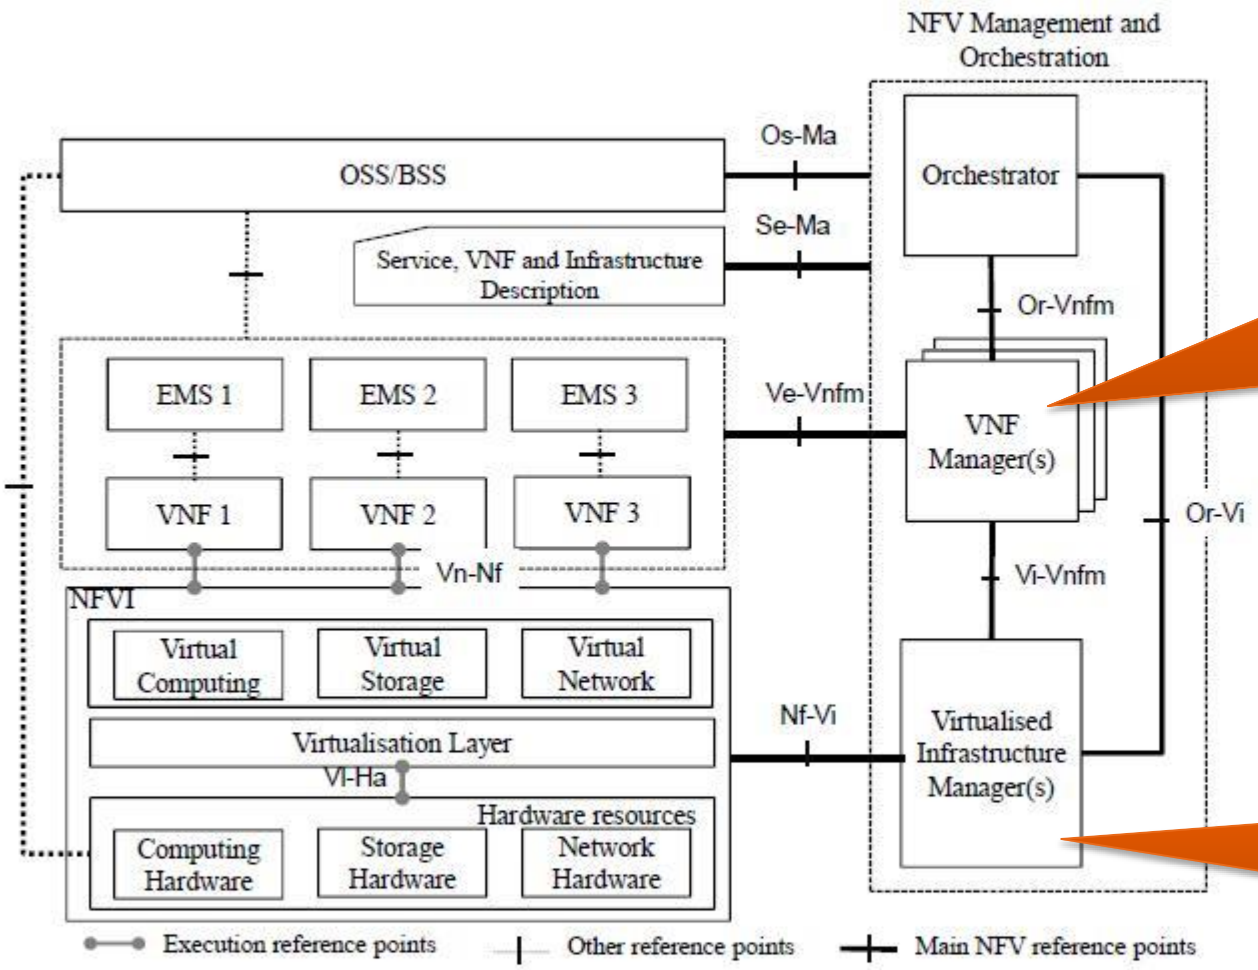

Life-cycle management of VNFs (deploy, scale, heal, upgrade)

VIM = an OpenStack instance

Mapping to ETSI NFV Architecture

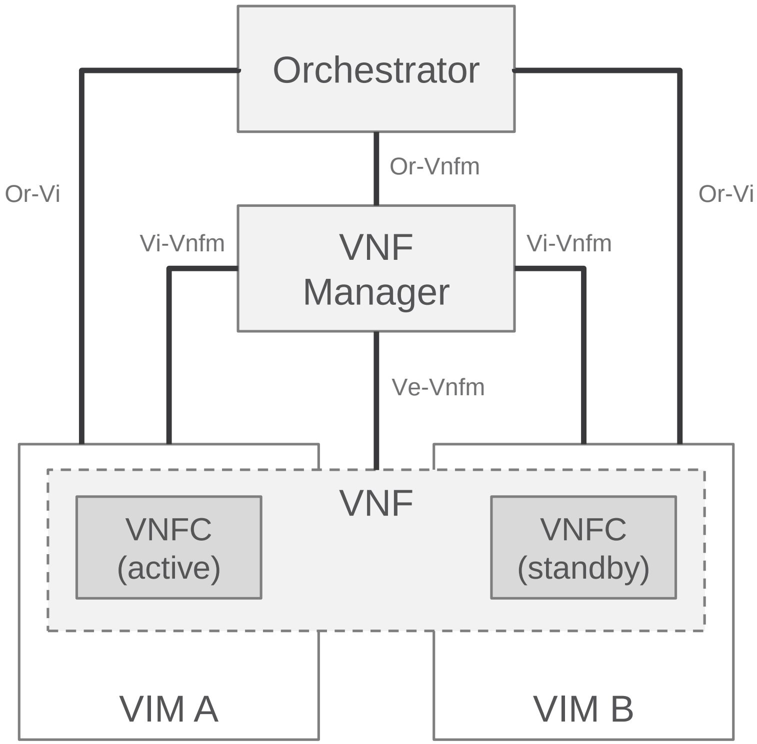

This is the “obvious” way to interpret the ETSI architecture for multi-VIM redundancy

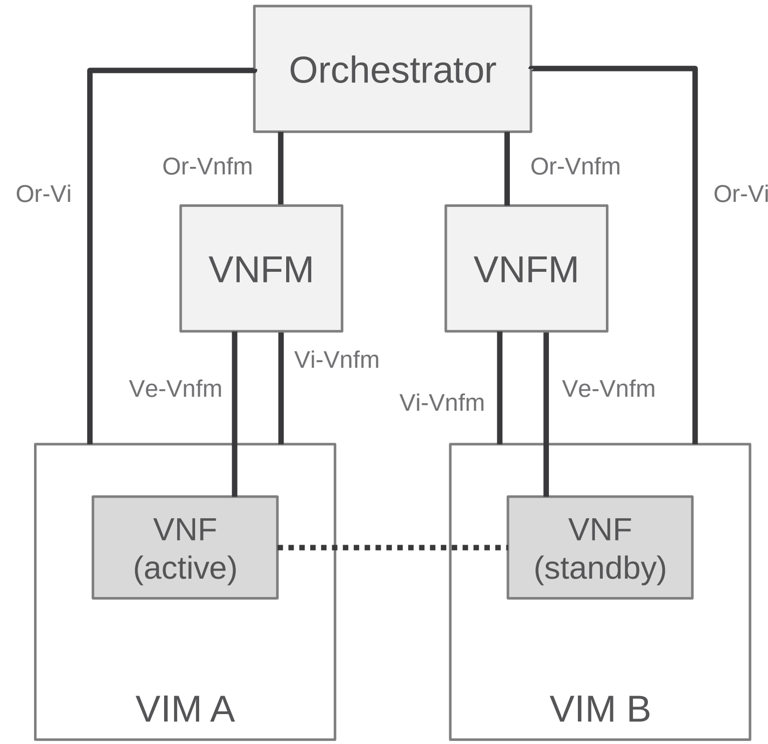

But this approach simplifies the VNFM and reduces the coupling between VIM instances

More Precise Mapping to ETSI NFV

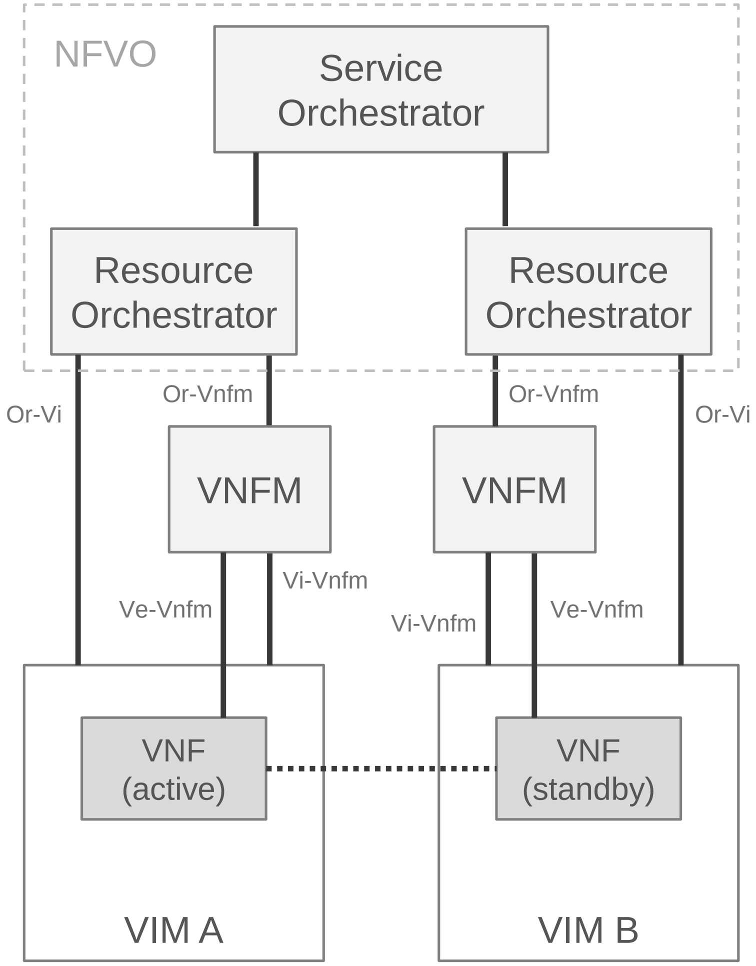

Service orchestrator – global scope

Resource orchestrator – VIM scope

Coordination between VNFs requires a shared data store that spans multiple VIMs

Service Availability vs Call Cutoff

Service Availability

99.999% means
“dial-tone” unavailable
< 6 minutes / year

99.9999% means
“dial-tone” unavailable
< 36 seconds / year

Call Cutoff

GR-512
Interval between
digital trunk drops
> 11 years

GR-511
Overall probability of
cutting off an
established call
< 1/8000

Call Continuity Across Failover?

Use Case: a VNF that performs media processing on voice calls (e.g. Interconnect SBC) deployed redundantly across two shared-nothing VIM instances

Assume VIM availability is 99.97%
→ 158 minutes / year downtime

What is the frequency of VIM failover?
Assume 15 minutes Mean Time to Repair
Failover events per year = ~ **10**

Probably not acceptable to drop tens or hundreds of thousands of calls at each failover event

Call Continuity Across Failover

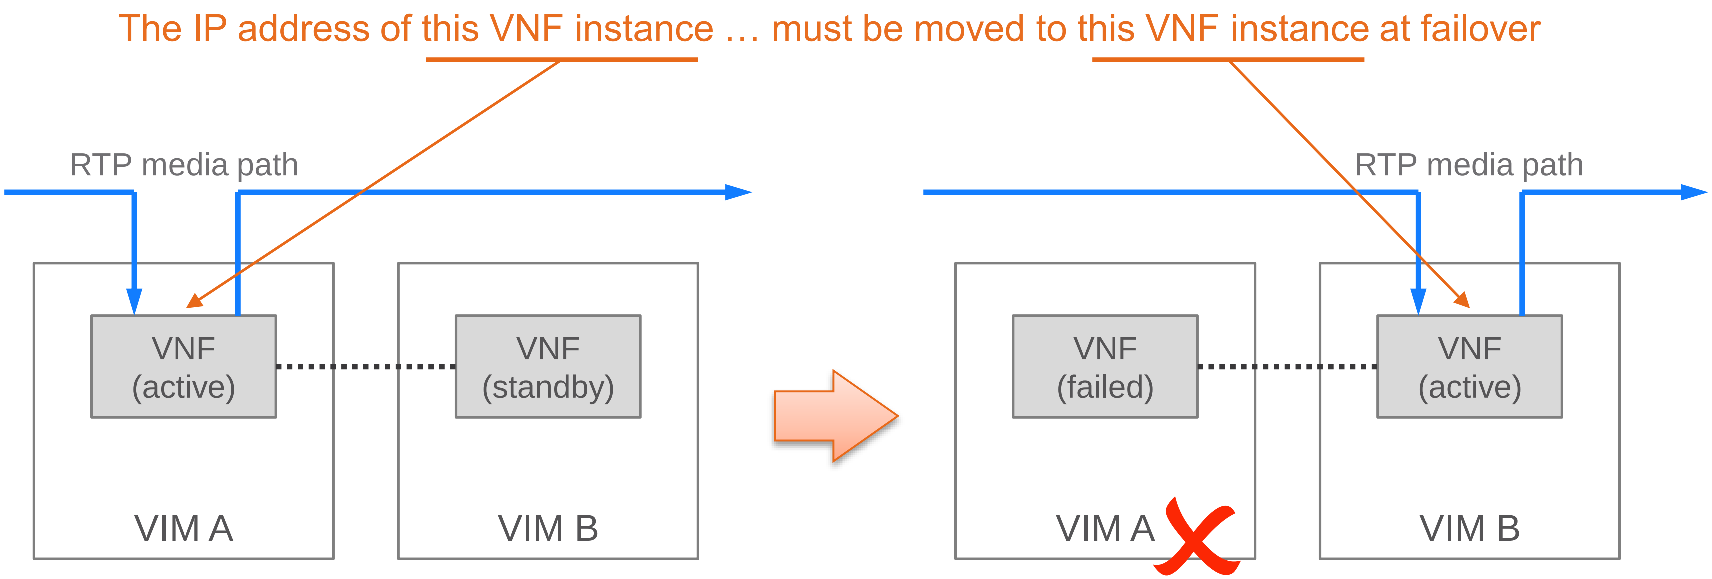

Moving IP addresses is the only way to preserve large numbers of RTP sessions across a failover with sub-second interruption

Moving IP Addresses Between VNFs

Today this is normally accomplished by connecting both VNF instances to the same L2 network, and using GARP

This technique can be extended between VIMs, although it may be painful if VIMs are geographically separated

Moving IP addresses at L3 is more “network-friendly”

Can be done via L3 control plane, e.g. injecting / withdrawing routes via BGP – but critically dependent on routers to respond quickly enough

This is an obvious candidate for interaction between VNFs and the SDN – but the requirement is not widely understood in the SDN community

Current State of Play for NFV

- Multiple shared-nothing VIM instances becoming accepted as the basis for telco-grade services
- Significant open questions remain
  - How to deploy redundant VNFs across multiple VIM instances
  - Scope of VNF Manager function in the overall architecture
  - How to move IP addresses for real-time media failover
- We believe techniques exist that make true telco-grade service availability a realistic goal for NFV
- But these techniques require careful application

martin.taylor@metaswitch.com