IEEE 2014 CQR
Emerging Technologies
Reliability Roundtable

Readout to CQR

May 13, 2014
Westward Look Wyndham Grand Resort and Spa
Tucson, AZ

The Speakers at the IEEE ETR Round Table

From Left to Right: Chunming Qiao, Gengui Xie, Brian Levy, David Lu, Spilios Makris, Mehmet Ulema, Michael Tortorella, Ying Chin (Bob) Yeh , Chi-Ming Chen

ETR-RT Speakers

(In Alphabetical Order)

- Brian Levy – Juniper Networks (CTO, SP Sector EMEA)
- David Lu – AT&T (VP, Business Ntwk & Corp. Solution IT)
- Spilios Makris – Palindrome Technologies (Director)
- Chunming Qiao – SUNY (Professor)
- Michael Tortorella – Rutgers University (Research Professor)
- Mehmet Ulema – Manhattan College (Professor)
- Gengui Xie – Huawei (VP of R&D Competence Center, China)
- Ying Chin (Bob) Yeh – Boeing Commercial Airplanes (Fellow)

ETR-RT Advisory Board

- Chi-Ming Chen – AT&T (Principal Member of Technical Staff)

ETR-RT Scope

- Identify the RAS (Reliability, Availability, and Serviceability) challenges, requirements, and methodologies in emerging technologies like:
  - ➢ Wireless/Mobility
  - ➢ Cloud Computing
  - ➢ NFV (Network Functions Virtualization)
  - ➢ SDN (Software Defined Networking)
  - ➢ High Availability for Mission-critical Industries (e.g., Airborne Systems)
- Discuss the RAS requirements and technologies, with the goal to promote the inter-industry sharing of related ideas and experiences
- Identify potential directions for resolving identified issues and propose possible solutions

ETR-RT Agenda

7:30 – 8:30 am	Breakfast
8:30 – 8:40 am	Welcome by Chair
8:40 – 10:15 am	Talks & Discussion
	➢ A Long Road to ETR Standardization (Makris)
	➢ Wireless / Mobility (Lu)
	➢ High Availability (Yeh)
10:15 – 10:30 am	Coffee Break
10:30 – 12 noon	Talks & Discussion
	➢ SDN/NFV (Xie, Levy, Ulema)
12 noon – 1:00 pm	Lunch
1:00 – 3:00 pm	Talks & Discussion
	➢ Cloud Computing (Qiao, Tortorella)
3:00 – 3:15 pm	Coffee Break
3:15 – 4:30 pm	Identify potential directions for resolving identified issues and propose possible solutions (All)
4:30 – 5:00 pm	Wrap –up & Summary for Presentation to CQR (All)

The presentation files are posted on the ETR-Roundtable website:
http://www.ieee-cqr.org/ETR-RT.htm

Old Reliability Roadmap

Example 2: User-to-User Connection on an IP Network (ITU-T Rec. Y.1541)

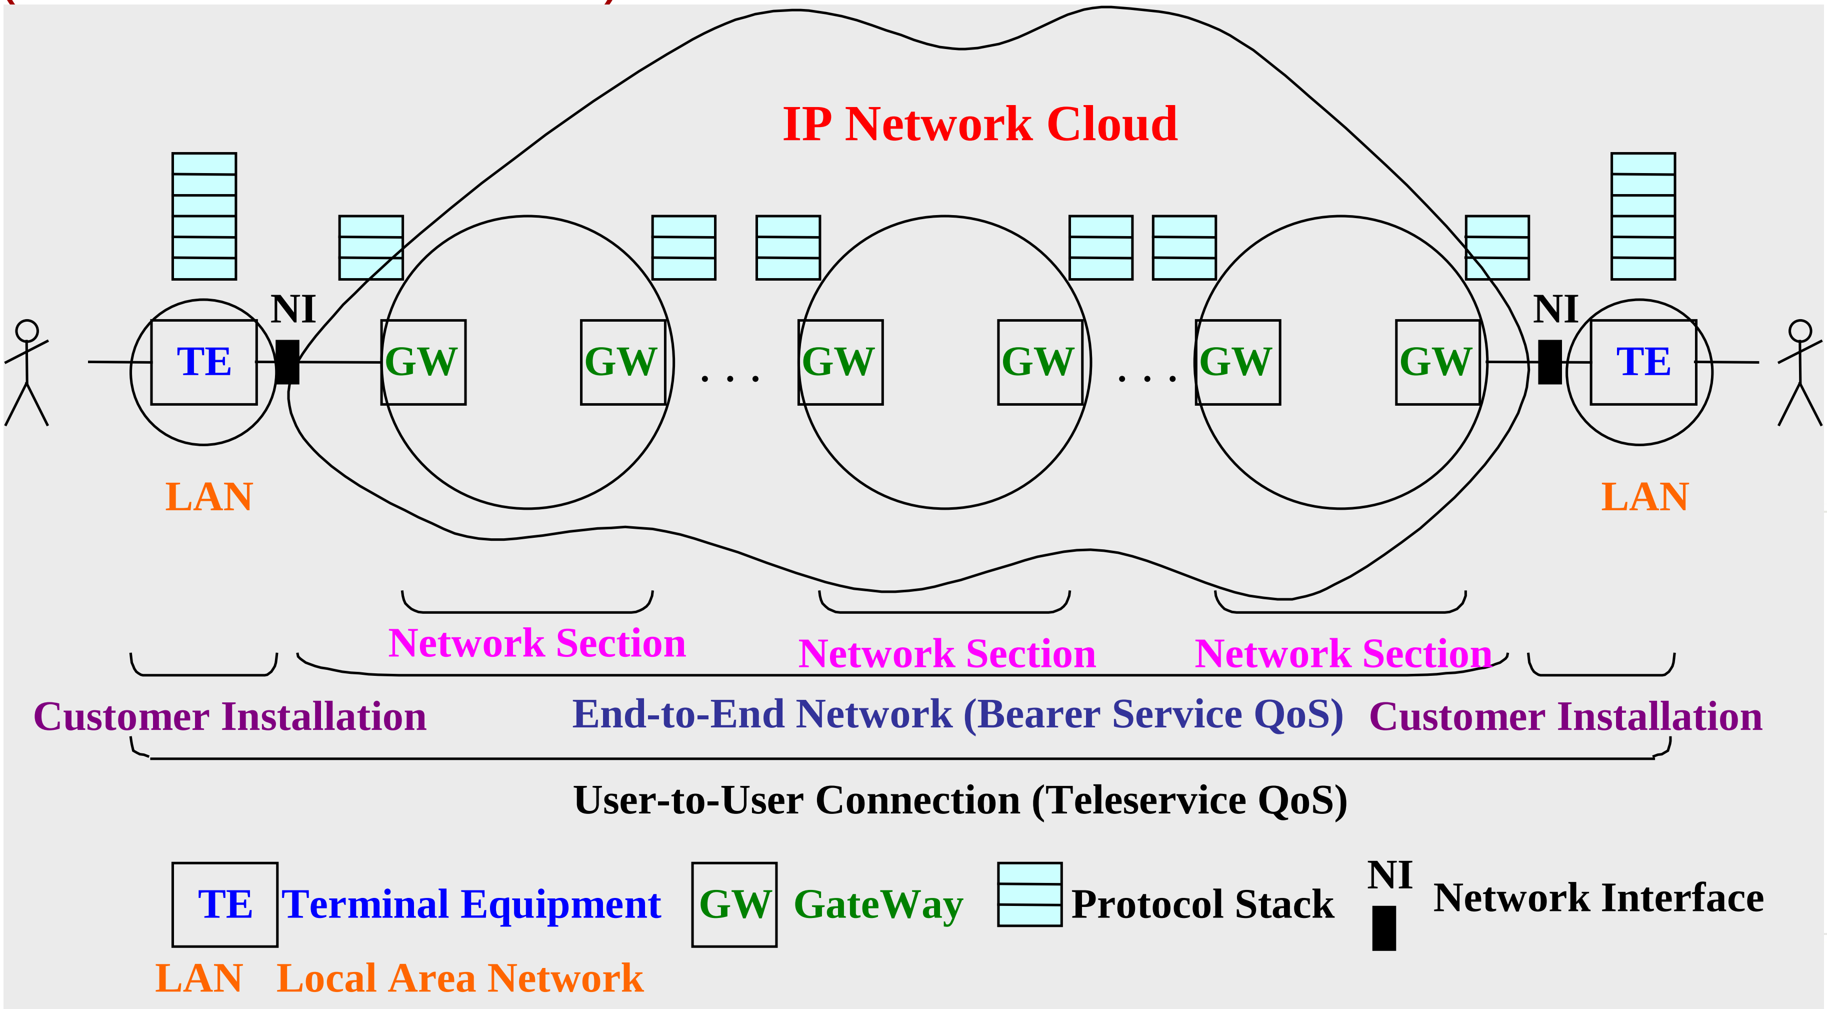

The “IP Network Cloud” will take care of any resilience issues!
Use a “cut-through” path to estimate end-to-end Network Reliability!

BEFORE Software Defined Networking

Status Quo: Closed Boxes, Fully-Distributed Protocols

- Fragmented non-commodity hardware
- Physical install per system per site
- Hardware development large barrier to entry for new vendors, constraining innovation & competition

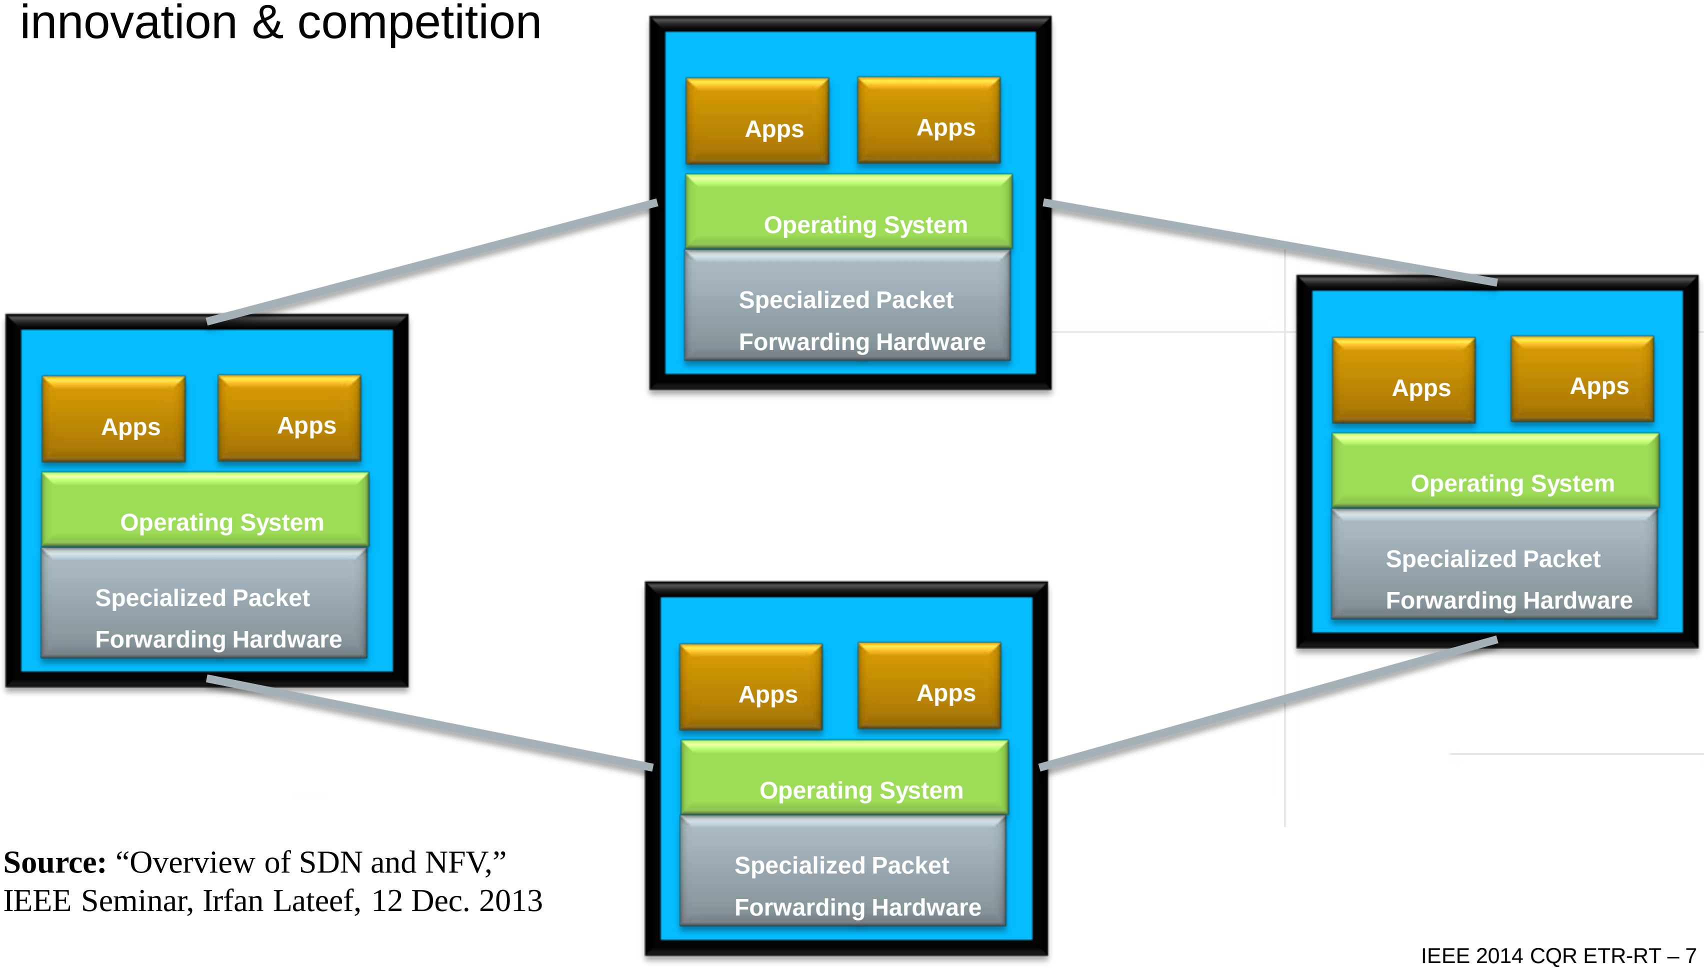
Apps
Apps
Operating System
Specialized Packet Forwarding Hardware
Apps
Apps
Operating System
Specialized Packet Forwarding Hardware
Apps
Apps
Operating System
Specialized Packet Forwarding Hardware
Apps
Apps
Operating System
Specialized Packet Forwarding Hardware

Source: "Overview of SDN and NFV,"
IEEE Seminar, Irfan Lateef, 12 Dec. 2013

Software Defined Networking (SDN) High-Level View

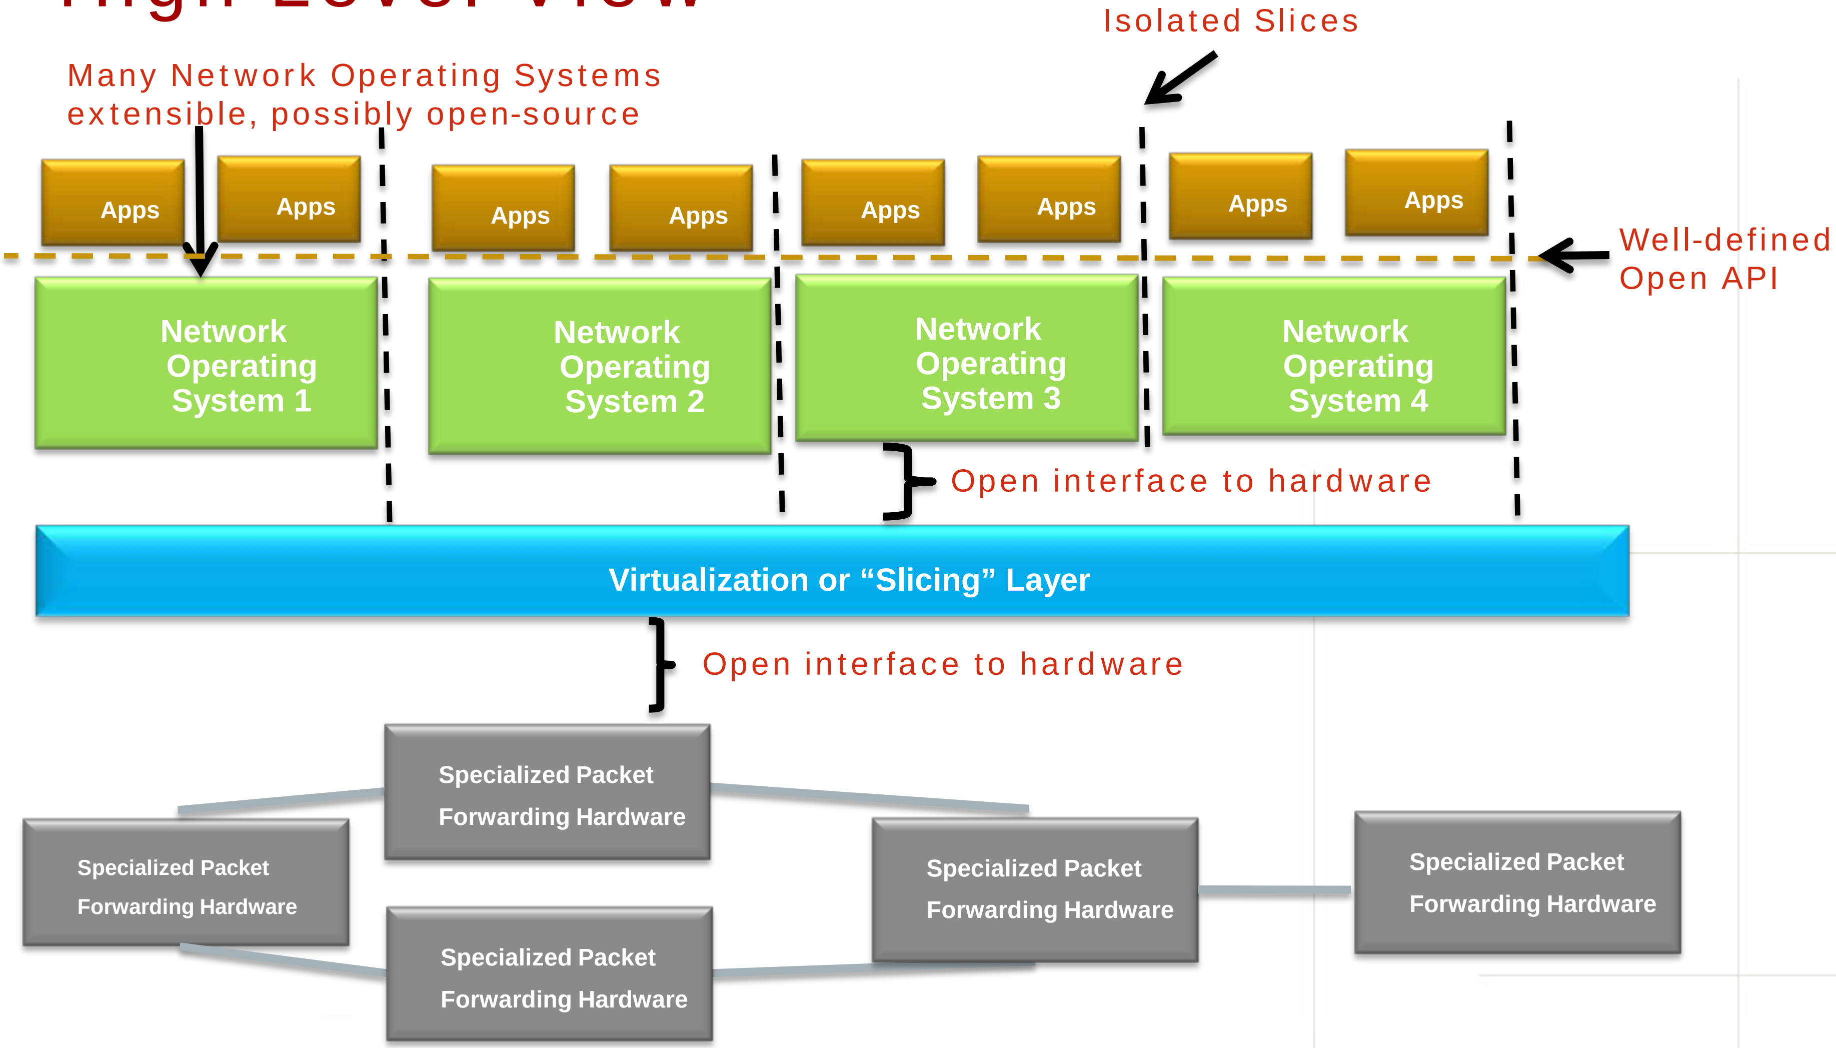
Isolated Slices
Many Network Operating Systems extensible, possibly open-source
Apps
Apps
Apps
Apps
Apps
Apps
Apps
Apps
Well-defined Open API
Network Operating System 1
Network Operating System 2
Network Operating System 3
Network Operating System 4
Open interface to hardware
Virtualization or “Slicing” Layer
Open interface to hardware
Specialized Packet Forwarding Hardware
Specialized Packet Forwarding Hardware
Specialized Packet Forwarding Hardware
Specialized Packet Forwarding Hardware
Specialized Packet Forwarding Hardware

The “Old Reliability Roadmap” needs to be adjusted!

API: Application Programming Interface

NFV Challenges: Far more about Function

RAS challenges caused by

✓New architecture: Decoupled HW&SW, COTS, more layers,…..

✓New Characteristics: Scalable, Dynamic/Elastic ,……

✓New system: Normally more SW bugs for a completed new system

5. Some HA mechanisms might need to be re-architected because of the virtualization.

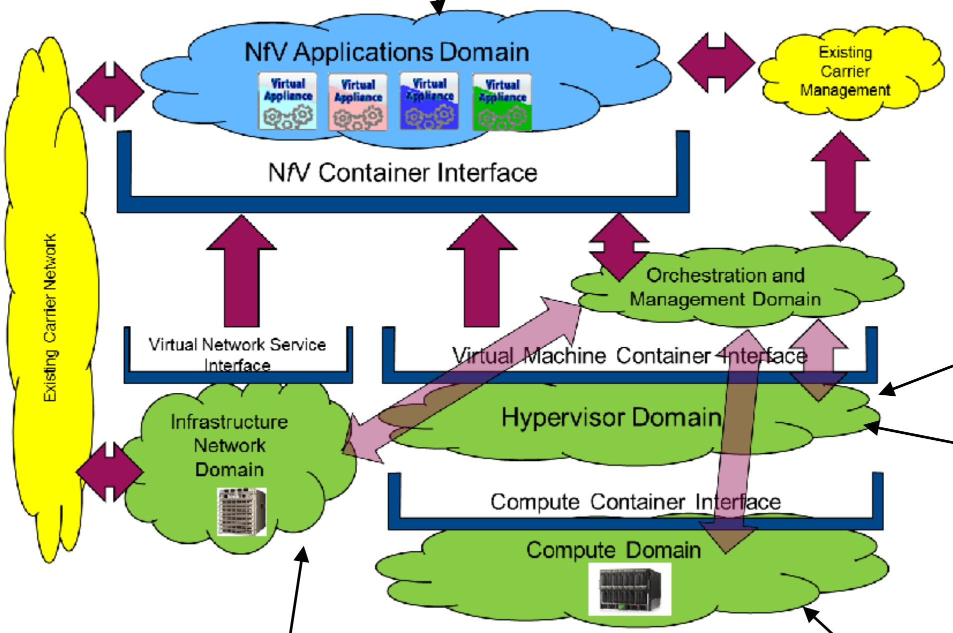
NfV Applications Domain
Virtual Appliance
Virtual Appliance
Virtual Appliance
Virtual Appliance
Existing Carrier Management
NfV Container Interface
Existing Carrier Network
Orchestration and Management Domain
Virtual Network Service Interface
Virtual Machine Container Interface
Hypervisor Domain
Infrastructure Network Domain
Compute Container Interface
Compute Domain

3. New fault modes for the new virtualization layer.

2. The virtualization has separated HA and the HW fault detection system.

4. HW/SW Components from different providers make fault localization more challengeable.

1. Less reliable COTS hardware

Overall Network (or System) Reliability: Key Issues

- In today’s traditional networks, we build in resiliency often by diversity (e.g., path diversity, redundant equipment). We want to ensure that the probability of a total network (or system) failure is minimized.

- With NFV, we have a consistent platform architecture which hosts the various network functions. Faults within the platform architecture could span many (or even all) functions and the chance of a catastrophic failure could be higher than with a traditional network’s physical infrastructure.

- There are things that we can do to make the NFV platform architecture highly resilient. However, it is a more closely coupled system with a greater statistical coherence. Thus, the basic probability of a total NFV platform failure – e.g., during an upgrade or fault condition – could be much higher than in a traditional physical network. This can be minimized but a lot of thought is needed to address it.

VNF Run Time Issues

- When we deploy a new VNF we use the VNF descriptor as the “recipe” telling us what Virtual Machines to deploy, what images to load, etc. Associated with this recipe is a run time descriptor which informers the orchestrator as to the geographic rules for deployment examples of which could be:

  1. All Virtual Network Function Components (VNFCs) may be on the same processor
  2. No Virtual Network Function Components (VNFCs) may be on the same processor
  3. Any processor cluster is allowed
  4. Same processor cluster only
  5. Same node
  6. Anywhere

- Obvious the issues here are not just about reliability but also about performance, some types of VNFC connections (serial port and shared memory) require same processor geography

- When constructing the software within a VNF it is open to the vendor to decide on how re-sync will take place in the case of the partial failure of a function

Resiliency at Different Levels of Abstraction

- As well as building residency into a VNF function we can build resiliency into service chains (service chains being groups of VNFs joined together to support a service
- We can build a service chain with parallel branches say with a load balancer between them. In this case if a total function fails in one branch, the parallel branch can take over
- Load balancers can be used within network functions, between network functions or even in the infrastructure level.
- The dynamic nature of NFV offers new possibilities for protection such as the ability to create a new service chain on demand maybe over a difference physical infrastructure and to off-load to this. We are truly in a new world and are now longer bound by a fixed a siloed infrastructure
- Dynamic Scale-Out of functions can be used to mitigate unplanned failures

Infrastructure Resiliency

- The NFV platform of course runs on a physical infrastructure of servers storage and hypervisors virtual functions can dynamically move across this infrastructure which can be diverse.

- Root cause analysis of failures in the virtual layer of the architecture can be difficult to correlate to the physical layer and there is a danger that new functions could be mapped constantly to a failed piece of hardware and then remapped until they work.

- We need to be able to identify hardware failures with correlation to the virtualised functions which run on them. Virtualised resiliency models at the software layer do not obviate the need for hardware reliability but can mask failures at the service level.

- We have to consider infrastructure resiliency all the way from the physical right up to the hypervisor layer. A serious bug in the hypervisor layer could cause multiple failures of functions across the infrastructure for example.

SDN/NFV: ETR-RT Conclusions - 1

- The ETSI NFV architecture represents a true paradigm shift in the way network functions are actualized
- It offers new residency capabilities and flexibility; however with that comes an infrastructure that is more tightly bound architecturally and we need to taking account for this as we design for reliability
- We need to build in appropriate resilience at all levels of abstraction from the physical to the virtual
- We have far more flexibility that ever before and we need to be careful not to increase the operational complexity here
- The key will be automation; it just simply will not be possible to manage such an infrastructure any other way

Cloud Technologies

- Basic infrastructure components
  - Physical servers (and virtual machines, a.k.a VMs), racks, clusters
  - Power distribution units (PDUs) and cooling infrastructures
  - Switches, routers, and data center networks
- Increasing adoption/reliance
  - Providers: Amazon, Google, Microsoft, Rackspace, SaleForce, …
  - Clients: Individuals, and small-to-large companies/institutions
- Availability/reliability is a top concern
  - Cited by 67%, followed by device based security (66%), and cloud application performance (60%).

Cisco Global Cloud Networking Survey, 2012.

Cloud Technologies: Failures Are All Too Common

- Frequent small-scale failures and infrequent large-scale failures

- Typical first year for a new cluster (Jeff Dean, Google)
  - ~0.5 overheating (power down most machines in <5 mins, ~1-2 days to recover)
  - ~1 PDU failure (~500-1000 machines suddenly disappear, ~6 hours to come back)
  - ~20 rack failures (40-80 machines instantly disappear, 1-6 hours to get back)
  - ~5 racks go wonky (40-80 machines see 50% packet loss)
  - ~3 router failures (have to immediately pull traffic for an hour)
  - ~dozens of minor 30-second blips for DNS
  - ~1000 individual machine failures
  - ~thousands of hard drive failures

Why Current Cloud Services are Flawed

- Current Service Level Agreement (SLA) is loosely defined in terms of availability/reliability measurements

- Penalty term is not user-friendly. The refund is usually issued in the form of credit with a lot of exclusions:
  - Amazon EC2 will refund the user in the form of credit if fail to meet the SLA
  - Rackspace will credit the user 5% month fee for each 30 mins network/infrastructure downtime, up to 100% monthly fee of the affected server

- Lack of high availability/reliability guarantee for critical services
  - Cannot guarantee 3-9's, let alone 4-9's as in Telco networks

Cloud Technologies: Key Challenges and Solutions

- A user/app may request:
  - # of VMs for response-time performance: *n (e.g., 100)*
  - Desirable availability (possibly a range): *α (e.g., 99.9%)*
  - Desirable contract duration: *t* (e.g., 3 months)

The Cloud SP performs the following:

- Downtime prediction based on failure models
  - Model component failures
  - Determine downtime distributions
- Availability-aware cloud resource provisioning and allocation
  - Determine the optimal (minimal) # of backup VMs, *k*, to be allocated
  - Both risk and energy minimizing placement of *n*+*k* VMs
- SLA contract design
  - Determine its costs: CAPEX (~h(*n; k*)) and OPEX (~ energy consumption)
  - A price list (schedule) for <duration, availability-guarantee, penalty>

Fulfilling the Promise of Service Reliability in Emerging Technology based Networks

- End-users of services don’t know and don’t care what technologies are being used to deliver these services
  - Even when technologies change (*e. g.*, SDN)
  - Enterprises and individuals may have different attitudes
- Service Providers are responsible for ensuring that their infrastructure is capable of delivering reliable services
  - SLA

Cloud Services: SLA Risk Analysis

- SLAs for enterprise services are common
- SLAs for cloud services are riskier
  - Reliability has been less well characterized

Cloud Computing: ETR-RT Conclusions - 2

- Conceptually straightforward
  - Key is understanding customer reliability requirements for the cloud service
    - Accessibility
    - Continuity / Fulfillment
    - Release

- Execution can be challenging
  - Infrastructure models
  - Networks with unreliable elements

Cloud Computing: ETR-RT Conclusions - 3

- Availability in cloud computing is very important
  - Has not received sufficient attention
  - Existing approaches not effective and need overhaul
  - Impedes many applications / business opportunities
- Key challenges and promising solutions
  - Downtime prediction based on failure models
  - Availability-aware VM provisioning and placement
  - SLA contract design for pricing, availability guarantee, penalty and duration
- Need multi-disciplinary and university-industry collaboration

IEEE Software Defined Ecosystem Standards
Highlights of the April 25, 2014 Mtg in Newark, NJ USA

- Objective
  - Identify primary standards development opportunities in SDN/ NFV & related areas
- Approach
  - Demonstrate and document the steps necessary to establish an early industry presence in key areas of opportunity
  - Provide a "gap" analysis to determine standardization opportunities in SDN/NFV and and related areas
  - Explore launching new standards activities in these areas under ComSoc sponsorship,
- Outcome - The participants:
  - Performed a gap analysis for each topic (i.e., security, reliability, performance) to determine standardization opportunities in SDN, NFV, and related areas
  - Discussed possible future steps to make further progress on each such group
  - Agreed to start the process to form a formal Study Group with the objective of launching new standards activities in these areas under IEEE ComSoc sponsorship
  - Agreed to recruit contributors

Source: Meeting Report of IEEE Software Defined Ecosystem Standards Working Meeting (M. Ulema, May 2, 2014)

ETR-RT Conclusions - 4

- Identified key issues that need to be addressed:
  - Lack of SDN/NVF Standardization in many areas including:
    - Reliability Framework
    - Reliability Metrics
    - Reliability/Availability SLAs
  - Potential overlap of Standards activities in different organizations (e.g., ETSI NFV, ONF, 3GPP, ATIS, IETF, etc.)
- Consider the end-user experience (who will care about it?)

Maintain momentum on Emerging Technologies (SDN, NFV, Cloud Computing, etc.)

by avoiding a protracted Reliability Standards effort fragmented in different SDOs

ETR-RT Conclusions - 5

- Lack of industry consistency for reliability analysis or benchmarks of Emerging Technologies-based networks
- Need for network providers to know if what they are building with Emerging Technologies will deliver the service availability/performance required by end-users
- New fault modes for the SDN/NFV virtualization layer (e.g. total platform failure)
- New SDN/NFV self-healing capabilities need to be evaluated
- Be realistic about the “x-nine” availability claims
  - Examples:

PSTN Network	= 99.94% (not even 4-nines!)
SONET/SDH Ring	= 99.9992% (5-nines)
SDN/NFV-based networks	= ???? (some talk about 6-nines!)

ETR-RT Proposed Next Steps

- Potential directions for resolving identified issues
  - Review (gap analysis) efforts in different Standards Development Organizations – SDOs (e.g., ATIS, ITU-T, IEEE, ETSI, IETF/IRTF, 3GPP, OMG, ONF, OMA-DM) and government-sponsored organizations/councils (e.g., NIST, CSRIC)
  - Support IEEE ComSoc Standardization efforts
    - Study Group on Security, Reliability, & Performance for Software Defined and Virtualized Ecosystems
      - Need to identify overlaps
      - Need to perform gap analysis for reliability-related issues
      - Prioritize potential tasks based on the gap analysis
    - Align with the requirements provided by ETSI NVF and WG Reliability/Availability
    - Liaise with ETSI TSC to agree on-going relationships
  - Organize more IEEE roundtables!