Network Functions Virtualisation:
The New Frontier of Telecoms Innovation

Daniel King
PhD Student – Lancaster University
d.king@lancaster.ac.uk

Research – Team

- LU Team
  - PhD Supervisors:
    - Professor David Hutchinson
    - Dr Christopher Edwards
    - Dr Nicholas Race
  - Research Partner
    - Chris Ford, Lancaster University Management School
- Academic Rationale
  - Opportunity to investigate an emerging area in computer science and telecommunications research.
  - Provide useful data and evidence to industry and standards development organisations.
- My Industry Experience
  - Bell Labs, Cisco Systems, Redback Networks, Movaz (ADVA), Aria Networks
  - IETF WG Secretary of ROLL, L3VPN, CCAMP and PCE.
    - Author: RFC4687, RFC5557, RFC6006, RFC6007, RFC6163, RFC6639, RFC6805.
    - Currently progressing 7 WG documents and 7 individual drafts.

Research – Network and Function Virtualisation

Classical Network Appliance Approach


BRAS


Firewall


DPI


Access Router


Carrier Grade NAT


Session Border Controller


WAN Acceleration

Network Function Virtualisation

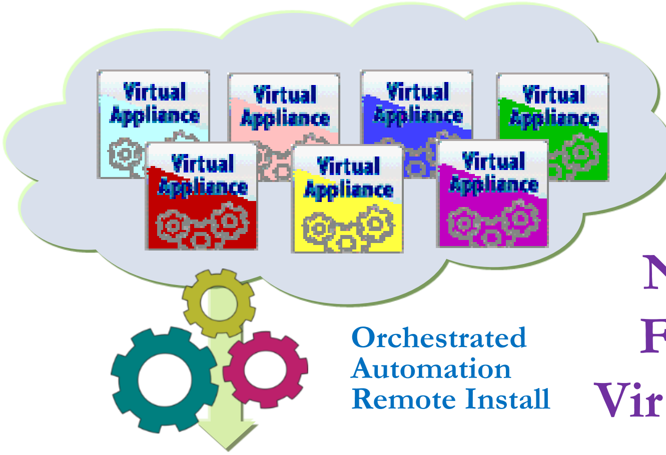

Orchestrated Automation Remote Install

Hypervisors


Generic High Volume Servers


Generic High Volume Storage


Generic High Volume Ethernet Switches

Research – Investigating the Problem Space

- Evidence gathering
- – "**A Critical Survey of Network Functions Virtualization**" to help define the problem space
    - – Qualitative and exploratory study (Eisenhardt 1989, Yin 2009, Thomas 2011)
    - – Inductive, hypothesis-generating approach
    - – Guided by tenets of Grounded Theory (Glaser and Strauss 1967, Charmaz 2006, Corbin and Strauss 2008, Suddaby 2006)
- Analysis (Miles and Huberman 1994)
- – Detailed coding of interview transcripts (nVivo).
    - Development of concepts and their dimensions.
    - Intensive review around each concept.
- Interpretation
- – Combining memos & concepts into cohesive whole.
    - Establishing cross-user connections.
    - Identifying industry comparatives to inform analysis (e.g., Human Genome Mapping)
- Writing up
- – Develop substantive model and frameworks.
- – Construct authentic & plausible arguments (economic and technical) based on evidence.
- – Publishing findings and conclusions documents (including IETF informational I-Ds and ETSI contributions).

Research – NFV Concept Development

- European Telecommunication Standards Institute (ETSI)
  - Role has been to provide an environment to develop the problem space.
  - Responsibility to publish problem statements, requirements and recommendations.
- ETSI NFV History
  - Whitepaper "Network Functions Virtualisation - An Introduction, Benefits, Enablers, Challenges & Call for Action", October 2012.
  - Initial concepts discussed at the end of 2012 in ETSI Future Networks Workshop.
  - Formal Industry Specification Group (ISG) session in January, 2013.
  - NFV ISG has met twice in 2013, with a third session planned for Bonn in July 2013.

Research – ETSI NFV ISG Structure

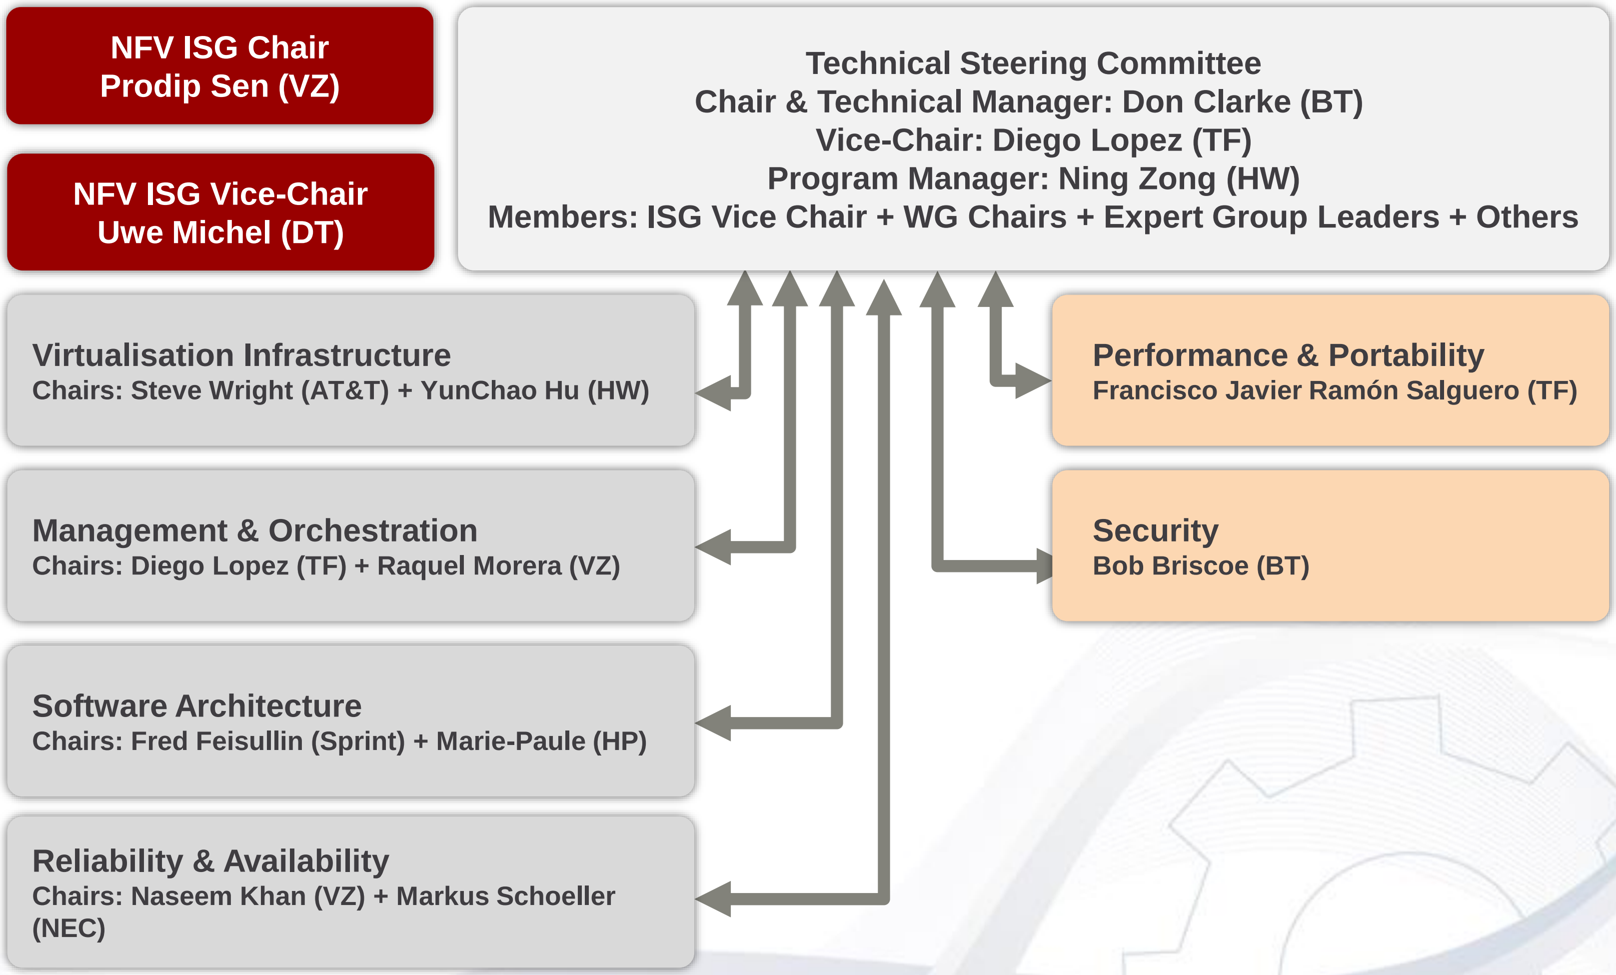

Research –NFV ISG Work Contributions

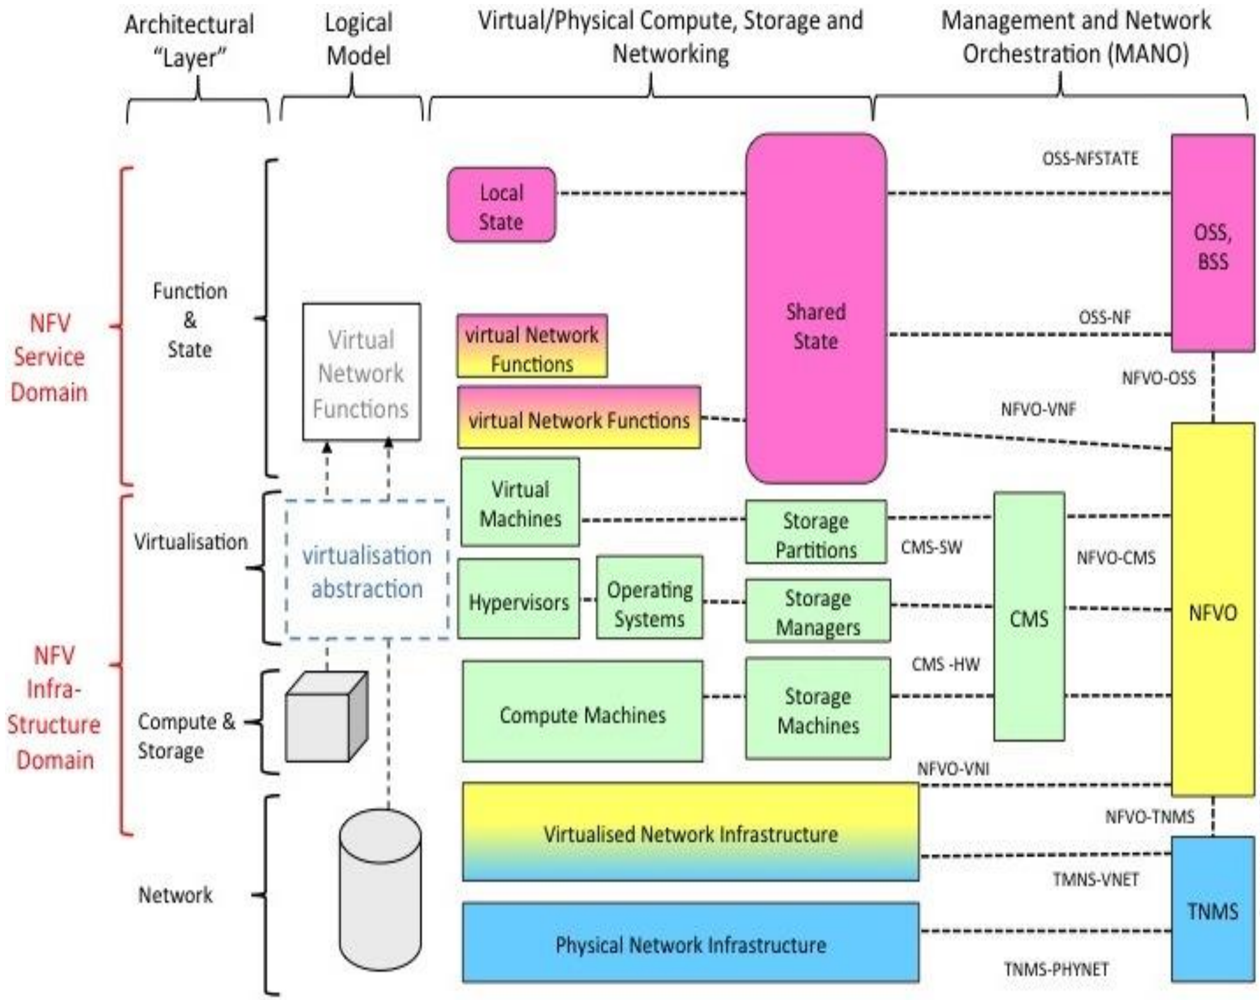

Research – NFV Interviewees

- A total of Twenty (20) CSPs have been identified and targeted.
- Discussions and interviews to date:
  – British Telecom
  – Verizon
  – KDDI
  – AT&T
  – Telefonica
  – Telstra
  – NTT docomo
  – France Telecom
  – Deutsche Telekom
- Initial focus on CSPs to gain rich data and develop initial concepts.
- Second round includes vendors and other stakeholders.

Findings – So Far (1)

- Operators have been independently researching network and function virtualisation with hardware and software vendors for years.
- "Enablers for NFV?"
  - Open Innovation during early stages of process and technology development
  - Performance of commodity hardware
  - Success of previous Hosted and Cloud Services
- Most interviews highlighted that industry cooperation is required to:
  - Sanity check use cases.
  - Apply pressure on vendors.
  - Provide the economy of scale for commercial development, deployment and operation of NFV-enabled services.

Findings – So Far (2)

- Infrastructure Complexity
  - Increasing variety of proprietary hardware and dedicated function.
  - Current nodes are fragmented with disparate operation and management.
- Energy Consumption
  - Sites are expanding while operators and customers are being directed to reduce CO2 emissions.
- Service Deployment
  - The time to specify, procure, integrate and deploy needs to be radically reduced.
  - Increased automation of service deployment.
- Rationalisation of Operation Support Systems
  - Physical presence and consequent operations per component and site.
  - Too many disparate OSS and NMS entities in the network.

Findings – Network Functions Virtualisation

- BT Virtualisation Testing from 2012 [1]
- Combined BRAS & CDN functions on Intel® Xeon® Processor 5600 Series HP c7000 BladeSystem using Intel® 82599 10 Gigabit Ethernet Controller sidecars
  - BRAS chosen as an “acid test”
  - CDN chosen as architecturally complements BRAS
- BRAS created from scratch so minimal functionality:
  - PPPoE; only PTA, priority queuing; no RADIUS, VRFs
  - CDN COTS – fully functioning commercial product

[1] Bob Briscoe, Don Clarke, Pete Willis, Andy Reid, Paul Veitch, “Network Functions Virtualisation”
http://www.ietf.org/proceedings/86/slides/slides-86-sdnrg-1.pdf

Findings – Network Functions Virtualisation

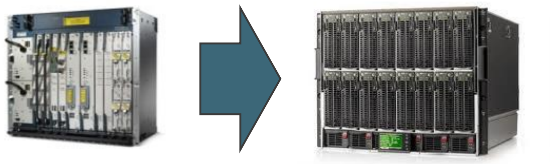

Test Id	Description	Result
1.1.1	Management access	Pass
1.2.1	Command line configuration: add_sp_small	Pass
1.2.2	Command line configuration: add_sub_small	Pass
1.2.3	Command line configuration: del_sub_small	Pass
1.2.4	Command line configuration: del_sp_small	Pass
1.3.1	Establish PPPoE session	Pass
1.4.1	Block unauthorized access attempt: invalid password	Pass
1.4.2	Block unauthorized access attempt: invalid user	Pass
1.4.3	Block unauthorized access attempt: invalid VLAN	Pass
1.5.1	Time to restore 1 PPPoE session after BRAS reboot	Pass
1.6.1	Basic Forwarding	Pass
1.7.1	Basic QoS - Premium subscriber	Pass
1.7.2	Basic QoS - Economy subscriber	Pass
2.1.1	Command line configuration: add_sp_medium	Pass
2.1.2	Command line configuration: add_sub_medium	Pass
2.2.1	Establish 288 PPPoE sessions	Pass
2.3.1	Performance forwarding: downstream to 288 PPPoE clients	Pass
2.3.2	Performance forwarding: upstream from 288 PPPoE clients	Pass
2.3.3	Performance forwarding: upstream and downstream from/to 288 PPPoE clients	Pass
2.4.1	Time to restore 288 PPPoE sessions after BRAS reboot	Pass
2.5.1	Dynamic configuration: add a subscriber	Pass
2.5.2	Dynamic configuration: connect new subscribers to BRAS	Pass
2.5.3	Dynamic configuration: delete a subscriber	Pass
2.5.4	Dynamic configuration: delete service provider	Pass
2.6.1	QoS performance – medium configuration	Pass
3.1.1	Command line configuration: add_sp_large	Pass
3.1.2	Command line configuration: add_sub_large	Pass
3.2.1	Establish 1024 PPPoE sessions	Pass
3.3.1	Performance forwarding: downstream to 1024 PPPoE clients	Pass
3.3.2	Performance forwarding: upstream from 1024	Pass

- Average 3 Million Packets Per Second per Logical Core for PPPoE processing.
  - Equivalent to 94 M PPS/97 Gbps per Blade = 1.5 G PPS/1.5 Tbps per 10 U chassis[1].
  - Test used 1024 PPP sessions & strict priority QoS
  - Test used an Intel® Xeon® E5655 @ 3.0 GHz, 8 physical cores, 16 logical cores (not all used).
- Scaled to 9K PPPoE sessions per vBRAS.
  - Support of 3 vBRAS per server.
- Subsequent BT research:
  - Implemented & testing software Hierarchical QoS.
  - Results so far show processing is still not the bottleneck.
  - Also tested vCDN performance & video quality.

"Performance potential to match the performance per footprint of existing BRAS equipment."

[1] Using128 byte packets. A single logical core handles traffic only in one direction so figures quoted are half-duplex.
[2] http://www.btplc.com/Innovation/News/NetworkVirtualization.htm

Next Steps – Management & Orchestration

- Management & Service Orchestration
  - Discovery of network resources.
  - Routing and path computation.
  - Network resource abstraction, and presentation to application layer.
  - Multi-layer coordination and interworking.
  - Multi-domain & multi-vendor network resources provisioning through different control mechanisms (e.g., Optical, OpenFlow, GMPLS, MPLS).
  - Policy Control.
  - OAM and performance monitoring.
- Leveraging existing technologies
  - What is currently available?
  - Integrate with existing and developing standards!

Next Steps – Management & Orchestration

- Application-Based Network Operations
  - A PCE-based Architecture for Application-based Network Operations
  - draft-farrkingel-pce-abno-architecture
- "Standardised" components
  - Policy Management
  - Network Topology
    - LSP-DB
    - TED
  - Path Computation and Traffic Engineering
    - PCE, PCC
      - Stateful & Stateless
      - Online & Offline
      - P2P, P2MP, MP2MP
  - Multi-layer Coordination
    - Virtual Network Topology Manager
  - Network Signaling & Programming
    - RSVP-TE
    - ForCES and OpenFlow
    - Interface to the Routing System (I2RS)

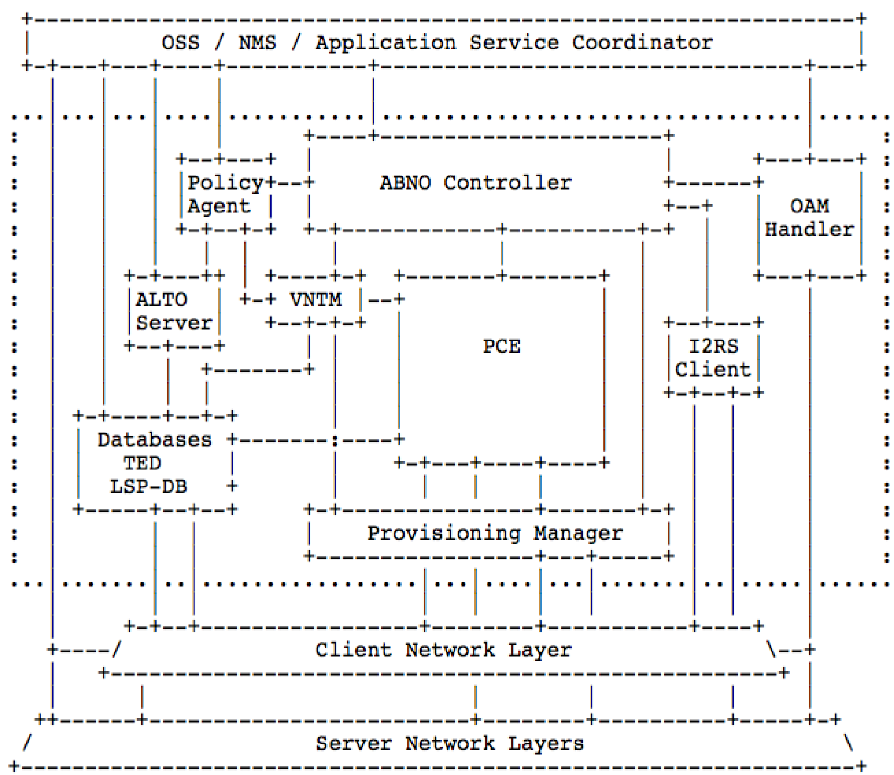

Figure 1: Generic ABNO Architecture

Next Steps – Currently

- Publish “Survey” results and findings.
- Developing orchestration and provisioning architecture and components for NFV applications
  - “Application-Based Network Operations (ABNO)” as an IETF Standard
- Documenting technical gaps for resiliency and restoration across use cases:
  - “Use cases and Requirements for Virtual Service Node Pool Management”
  - “An Overview of Reliable Service Nodes Discovery and Provision Protocols”
- Build Something!

Thank You!

Any comments or questions are welcome.

Daniel King
PhD Student – Lancaster University
d.king@lancaster.ac.uk