Network Functions Virtualisation

Bob Briscoe
Chief Researcher
BT

+ Don Clarke, Pete Willis, Andy Reid, Paul Veitch (BT)
+ further acknowledgements within slides

Message Router

Session Border Controller

WAN Acceleration

CDN

Carrier Grade NAT

DPI

Firewall

Tester/QoE monitor

SGSN/GGSN

PE Router

BRAS

Radio Network Controller

Classical Network Appliance Approach

Network Functions Virtualisation Approach

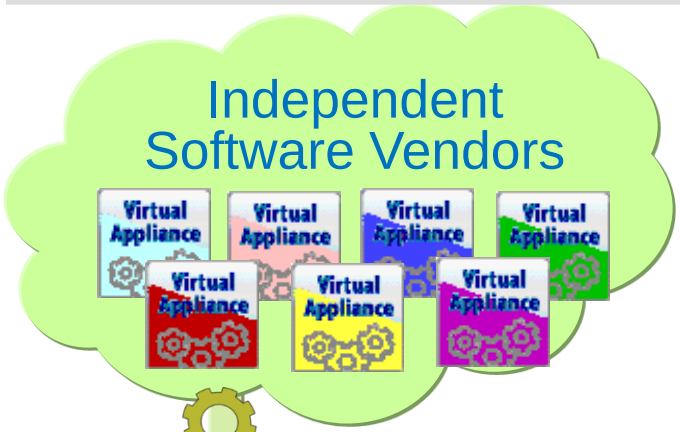

Orchestrated, automatic remote install

hypervisors

Generic High Volume Servers

Generic High Volume Storage

Generic High Volume Ethernet Switches

If price-performance is good enough, rapid deployment gains come free

Mar'12: Proof of Concept testing

- 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

acknowledge:

BT

hp

intel

WIND RIVER

verivue

tail-f

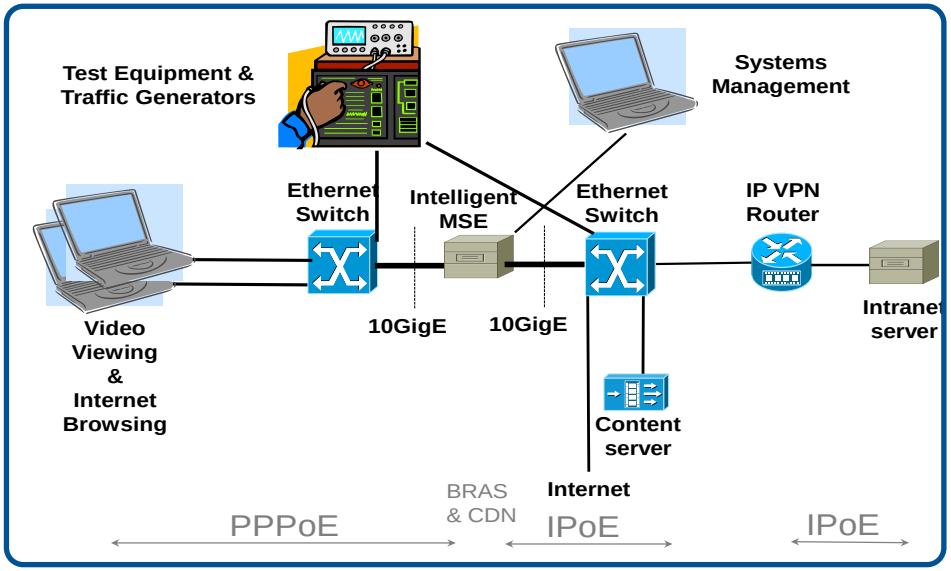

Significant management stack :

1. Instantiation of BRAS & CDN modules on bare server
2. Configuration of BRAS & Ethernet switches via Tail-F
3. Configuration of CDN via VVue mgt. sys.
4. Trouble2Resolve via HP mgmt system

Mar'12: Proof of Concept Performance Test Results

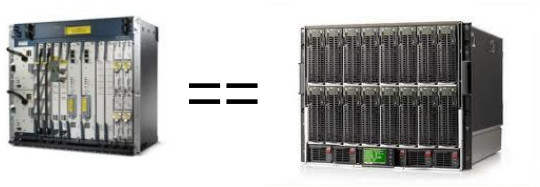

Test Id	Description	Result
1.1.1	Management access	**Pass**
1.2.1	Command line configuration: add_sp_small	**Pass**
1.2.2	Command line configuration: add_sub_small	**Pass**
1.2.3	Command line configuration: del_sub_small	**Pass**
1.2.4	Command line configuration: del_sp_small	**Pass**
1.3.1	Establish 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.
  - Can support 3 vBRAS per server.
- Subsequent research:
  - implemented & testing software Hierarchical QoS
  - results so far show processing is still not the bottleneck
  - (also tested vCDN performance & video quality)

very useful 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.

3 Complementary but Independent Networking Developments

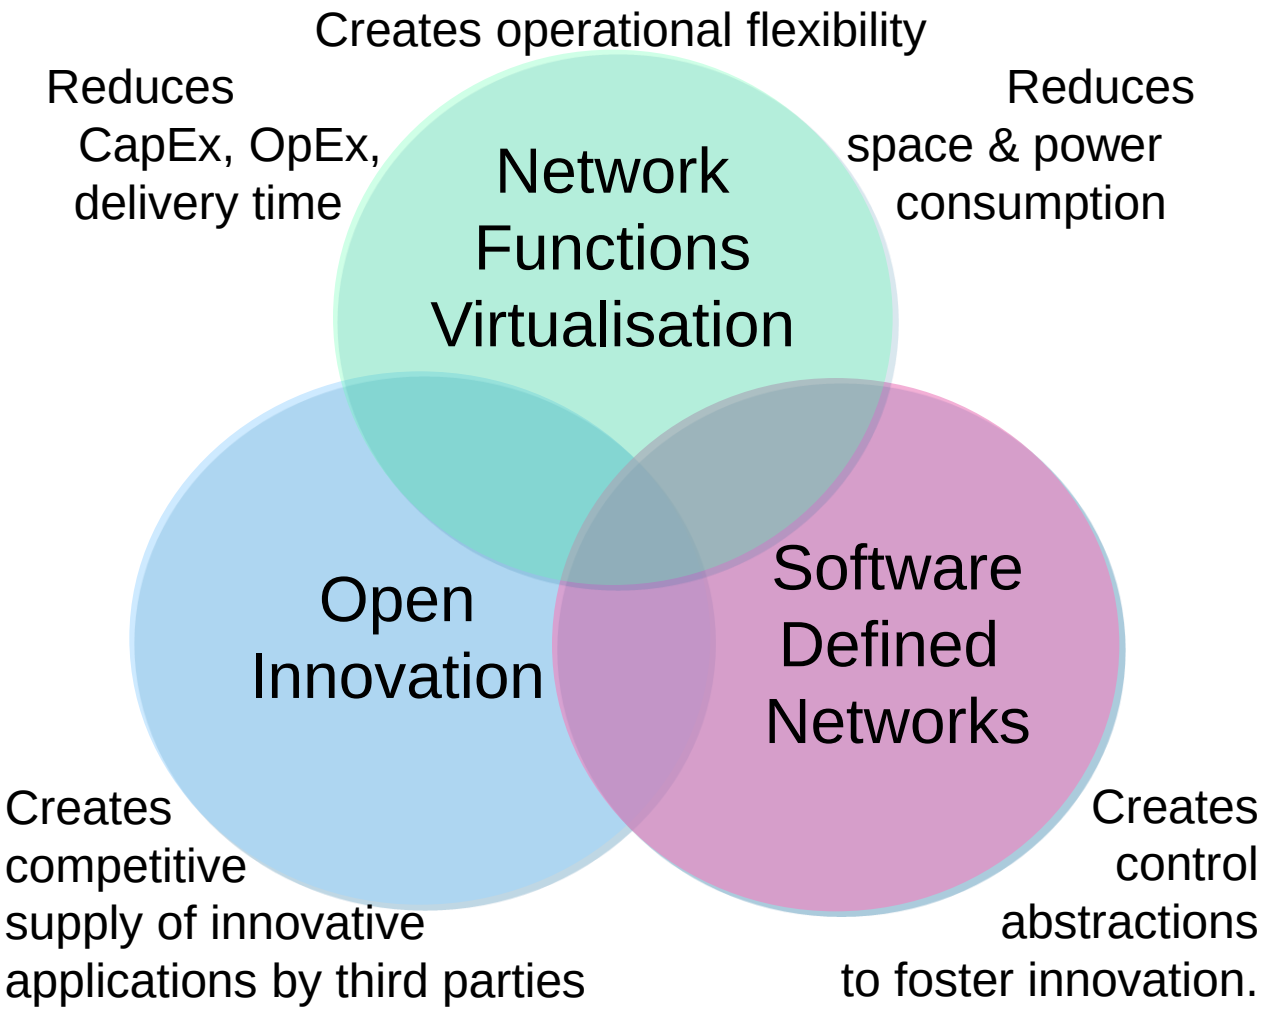

New NfV Industry Specification Group (ISG)

- First meeting mid-Jan 2013
  - > 150 participants
  - > 100 attendees from > 50 firms
- Engagement terms
  - under ETSI, but open to non-members
  - non-members sign participation agreement
    - essentially, must declare relevant IPR and offer it under fair & reasonable terms
  - only per-meeting fees to cover costs
- Deliverables
  - White papers identifying gaps and challenges
  - **as input to relevant standardisation bodies**
- ETSI NfV collaboration portal
  - white paper, published deliverables
  - how to sign up, join mail lists, etc

  http://portal.etsi.org/portal/server.pt/community/NFV/367

- Network-operator-driven ISG
  - Initiated by 13 carriers shown
  - Consensus in white paper
  - Network Operator Council offers requirements
  - grown to 23 members so far

gaps & challenges
examples

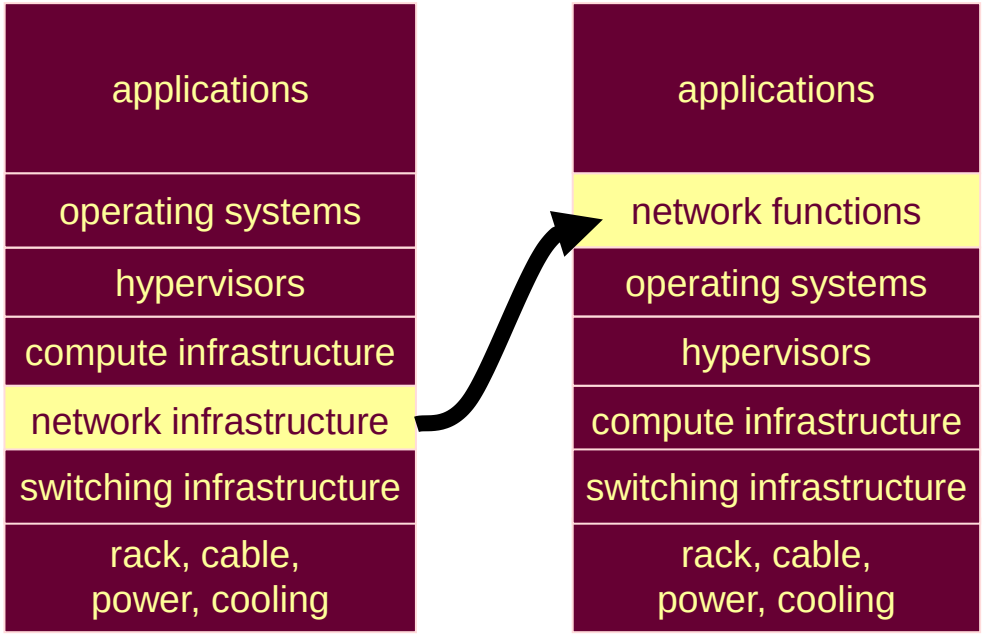

- management & orchestration
  - infrastructure management standards
  - multi-level identity standard
  - resource description language

- security
  - Topology Validation & Enforcement
  - Availability of Management Support Infrastructure
  - Secure Boot
  - Secure Crash
  - Performance Isolation
  - Tenant Service Accounting

Q&A

and spare slides

Domain Architecture

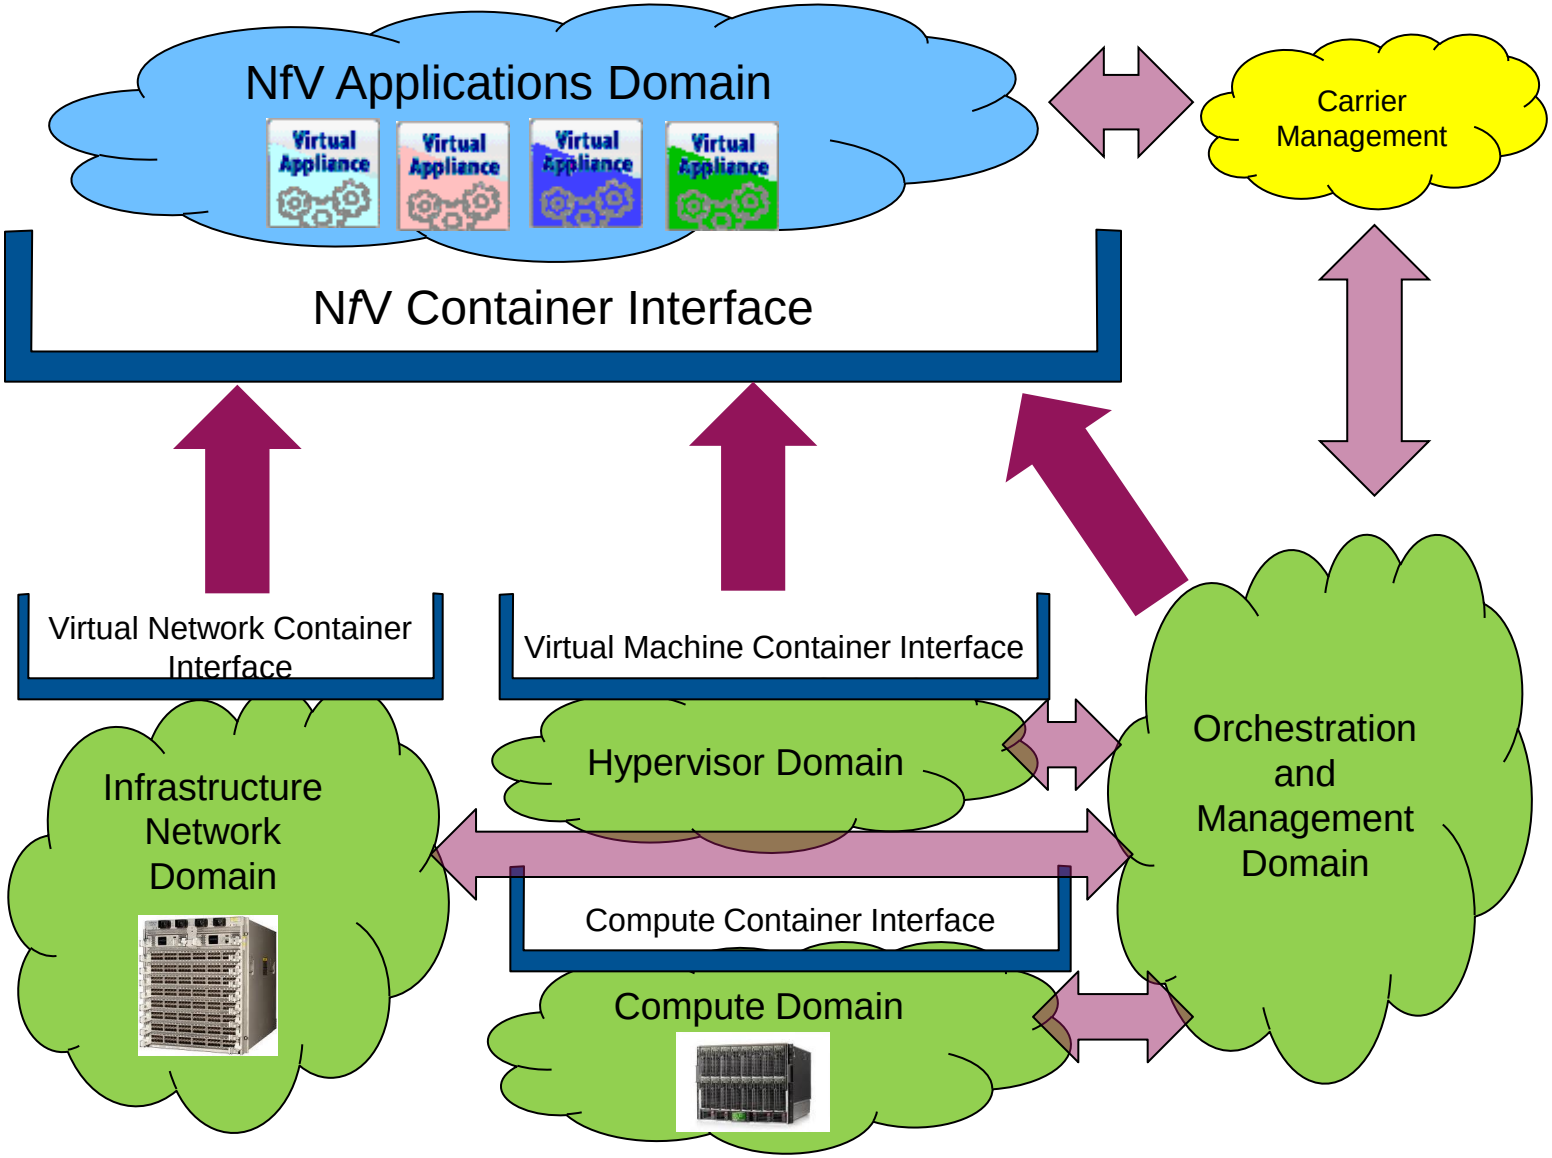

NVF ISG Organisation Structure…

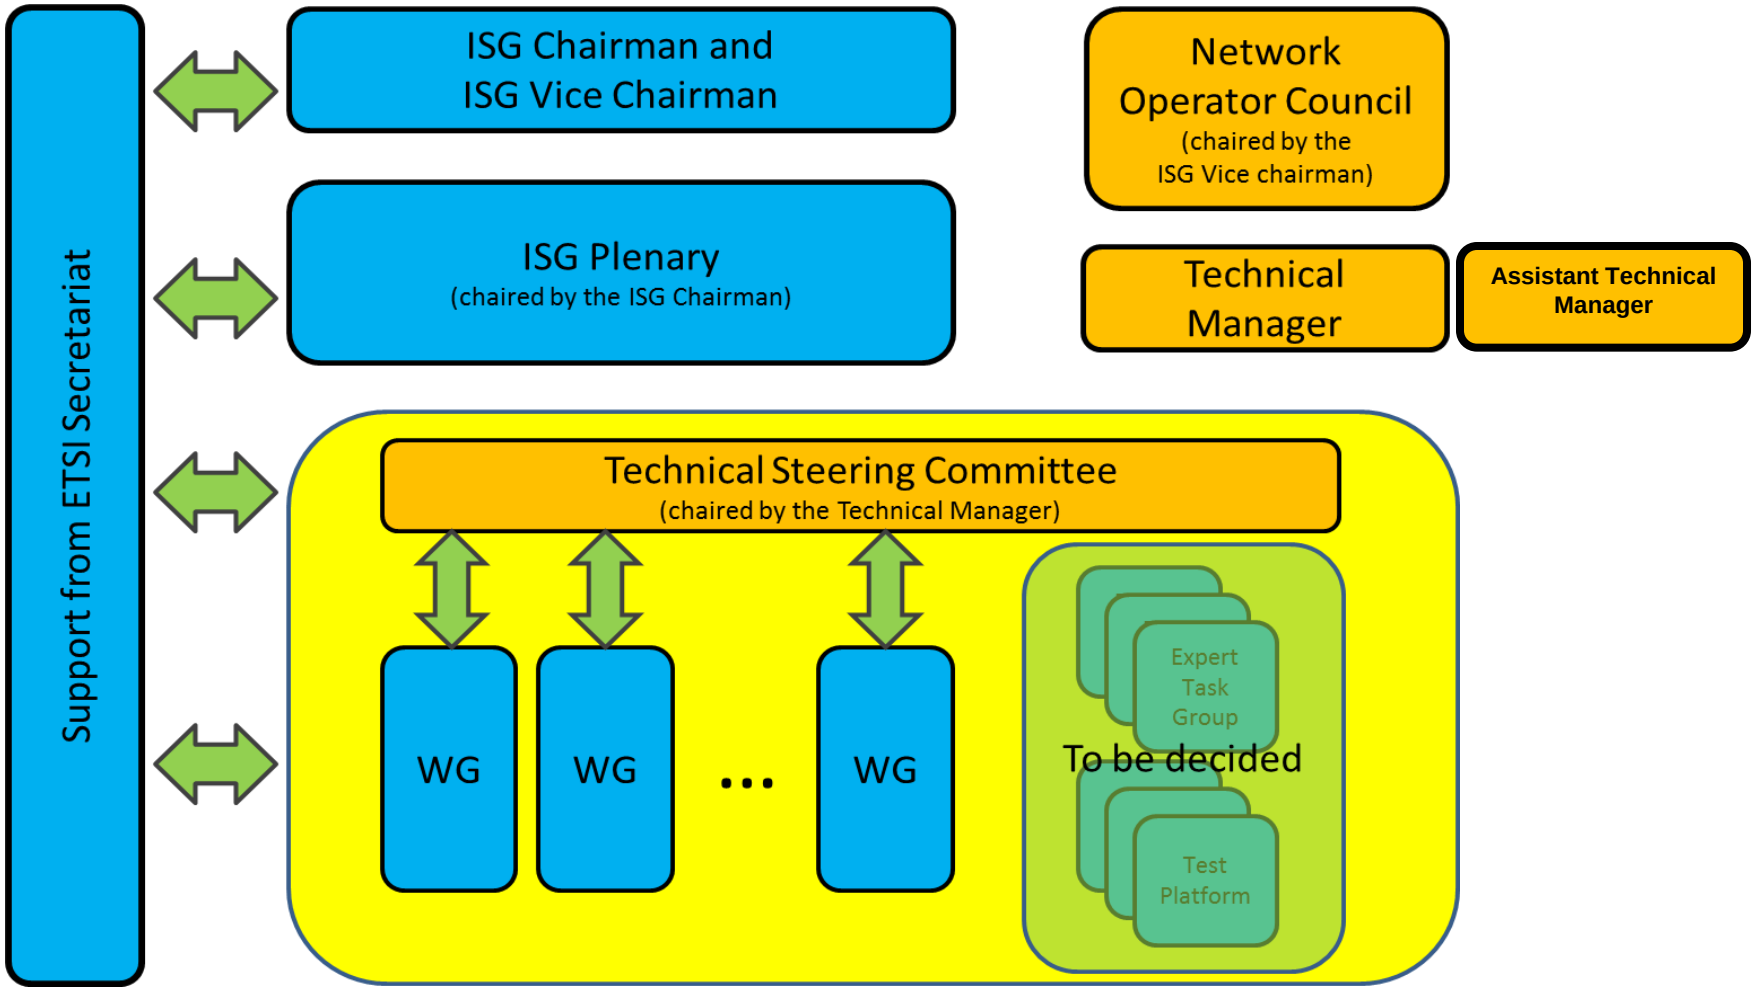

ISG Working Group Structure

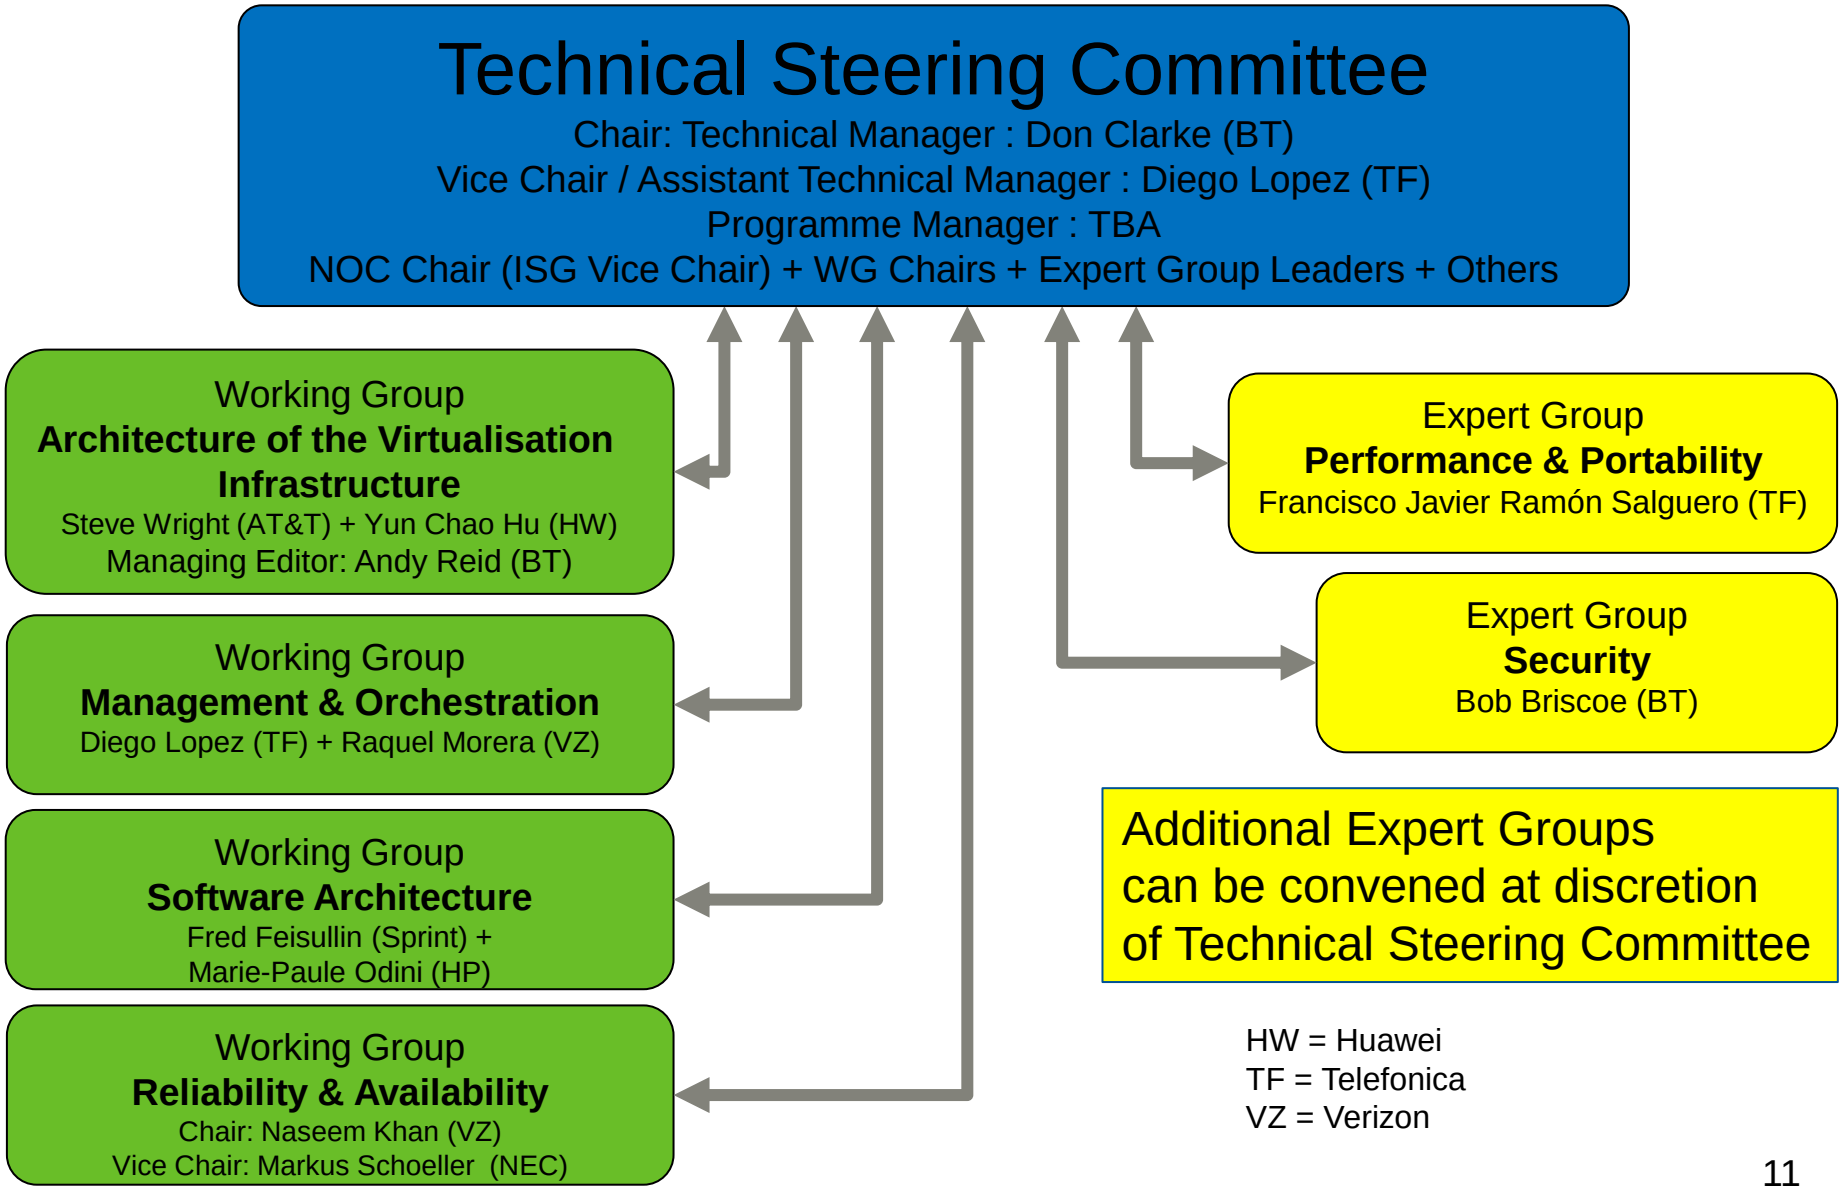

Hypervisor Domain

- General cloud hypervisor is designed for maximum application portability
  - Hypervisor creates
    - Virtual CPUs
    - Virtual NICs
  - Hypervisor provides
    - Virtual Ethernet switch
  - Hypervisor fully hides real CPUs and NICs
- N*f*V Hypervisor is aimed at removing packet bottlenecks
  - Direct binding of VM to core
  - Direct communication between VMs and between VMs and NIC
    - User mode polled drivers
    - DMA remapping
    - SR-IOV
- Many features already emerging in hypervisors

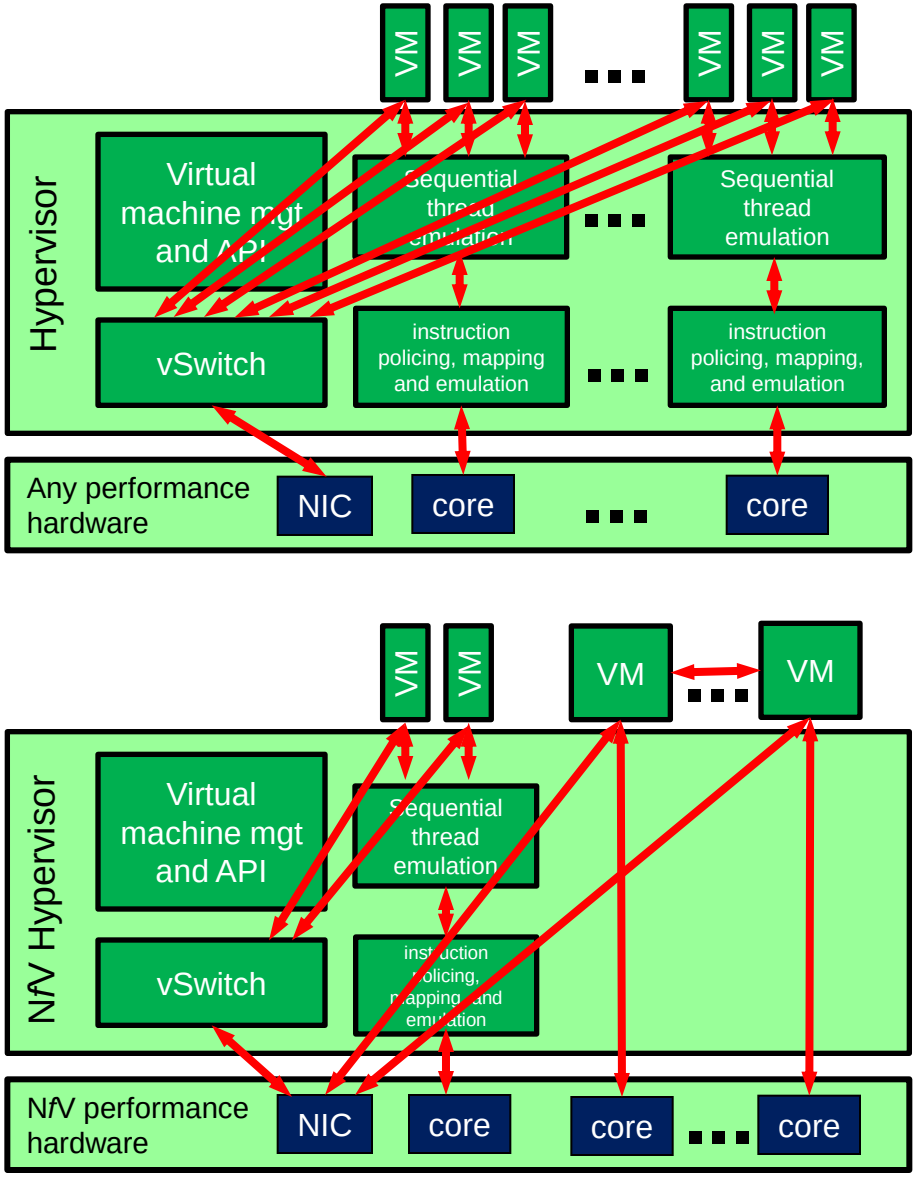

Orchestration and Infrastructure Ops Domain

- Automated deployment of N*f*V applications
  - Orchestration console
  - Higher level carrier OSS
- Tools exist for automated cloud deployment
  - vSphere
  - Openstack
  - Cloudstack
- NfV infrastructure profile for N*f*V application to
  - Select host
  - Configure host
  - Start VM(s)
- Application profile to specify
  - Service address assignment (mechanism)
  - Location specific configuration