GSMA Anti-Trust Policy Notice

- Anti-trust law prohibits (i) agreements (written or implicit) between competitors which may negatively impact consumers or competitors and (ii) sharing of confidential information
- All GSMA participants **must** abide by the following rules:
  - ✓**DO** clearly identify the positive purpose of each project and follow it
  - ✓**DO** consult with legal in areas where you are unsure
  - X**DON'T** enter into agreements that restrict other parties' actions or creates barriers to market entry
  - X**DON'T** discuss or exchange information on pricing, business plans, or any other confidential or commercially sensitive data

Linux Foundation - Anti-Trust Policy Notice

- Linux Foundation meetings involve participation by industry competitors, and it is the intention of the Linux Foundation to conduct all of its activities in accordance with applicable antitrust and competition laws. It is therefore extremely important that attendees adhere to meeting agendas, and be aware of, and not participate in, any activities that are prohibited under applicable US state, federal or foreign antitrust and competition laws.

- Examples of types of actions that are prohibited at Linux Foundation meetings and in connection with Linux Foundation activities are described in the Linux Foundation Antitrust Policy available at http://www.linuxfoundation.org/antitrust-policy. If you have questions about these matters, please contact your company counsel, or if you are a member of the Linux Foundation, feel free to contact Andrew Updegrove of the firm of Gesmer Updegrove LLP, which provides legal counsel to the Linux Foundation.

Common NFVI Telco Taskforce

RM Core Deep Dive

Mark Shostak, Kelvin Edmison

January 16, 2020

THE LINUX FOUNDATION

GSMA

RM Core Deep Dive Topics

- CNTT Target Workloads
- Non-Conforming Technology Proposal
- VNF Profile Generations and Evolution
- H/W Profiles and Performance
- H/W Selection Guidelines

CNTT Target Workloads

- Target workload classes for CNTT NFVI
- Priority of workload classes
- Relevance to & alignment with Public Cloud

Non-Conforming Technology Proposal (Petar Torre)

Decoupling applications from Infrastructure and PaaS, other application components, and application management/control

Telco (on-prem) Cloud:

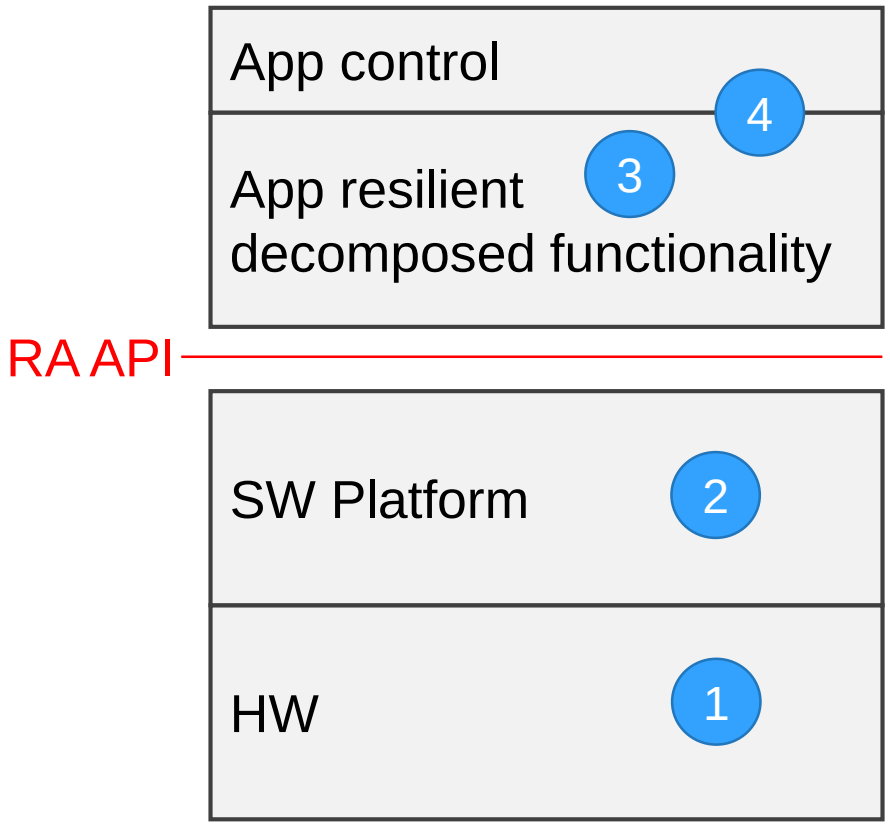

Service agility and easier operations will depend on levels of Decoupling between application and:

1. Infrastructure
2. Platform as a Service
3. Application Resiliency

Relevant for sizing infrastructure and operations:

4. Other app functionality (decomposition, manageability)

Non-Conforming Framework

- Handling technology that pierces the virtualization abstraction (e.g. SR-IOV, GPUs, FPGAs, SmartNICs), s/w and other areas
  - Decouple offending component from aggregate feature. Ex.:
    - Feature: SR-IOV
    - Dependent on: PCI-PassThrough/Direct Assignment
    - Issue: Requires h/w-specific code in workloads
    - Issue impact: Violates CNTT principle; Requires VNF customization; Restricts portability
- Policy identifies
  - Steps to mitigate/resolve, timeline to resolve, treatment by RC, etc.
  - Template format to ensure consistency across policies

VNF Profile Generations

- Drivers:
  - Generations in host hardware
    - Micro-architecture as ABI
    - ABI changes in new micro-architectures trigger need for specification of a per profile generation that can trace changes in ABI
  - Technology exceptions
  - NFVI need to simultaneously support VNFs that have migrated away from technology exceptions, and those that have not yet done so

VNF Profile Generations: Technology Deprecation straw-man

- Scenario: Profile B1 contains a permitted exception
- RM 3.0 puts a usage note on the technology.
- RM 4.0 introduces the alternative technology and marks the use of B1 with a stronger deprecation warning.
- RM 5.0 removes support for that profile, and VNFs will fail validation if they use it.

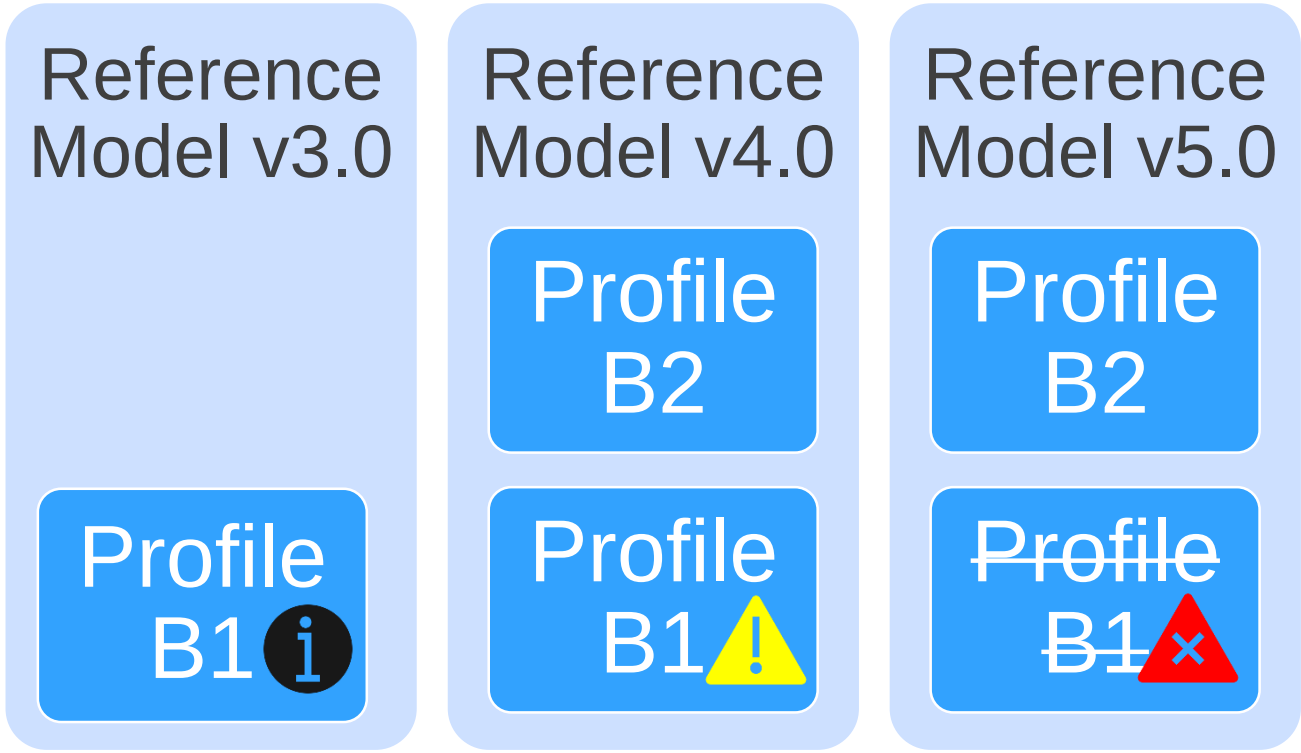

VNF Profile Generations: Profile template straw-man proposal

Profile Attributes	
Profile Name	B1
Profile Family	Basic
CPU Micro-architecture	Haswell/Broadwell
CPU Speed	2.2/2.1 GHz
Overbooked	N

Profile Name	vCPUs	RAM	Disk
B1.tiny	1	0.5	1
B1.small	1	2	20
B1.medium	2	4	40
B1.large	4	8	80
B1.2xlarge	8	16	160
B1.4xlarge	16	32	320

Permitted Exceptions
SR-IOV

Note: Actual values are not being proposed; only examples of how attributes relate with subsequent generation of profile

VNF Profile Generations: Profile template straw-man proposal

Profile Attributes	
Profile Name	**B2**
Profile Family	Basic
CPU Micro-architecture	**Cascade Lake**
CPU Speed	2.1 GHz
Overbooked	N

Profile Name	vCPUs	RAM	Disk
B2.tiny	1	0.5	1
B2.small	1	2	20
B2.medium	2	4	40
B2.large	4	8	80
B2.2xlarge	8	16	160
B2.4xlarge	16	32	320

Permitted Exceptions
~~**SR-IOV**~~

Note: Actual values are not being proposed; only examples of how attributes relate with previous generation of profile

H/W Profiles and Performance

Problem Statement:

Need to be able to assure reasonable level of deterministic performance

- Question: To what level of granularity should CNTT specify h/w
- Objectives and drivers for specifying h/w config
- Question: Can you combine normalized performance coefficients w/ high-level h/w specification to address problem?

H/W Profiles and Performance

- CNTT documents strive for abstraction of hardware from VNF workloads
- However…CPU microarchitectures periodically add new instructions & features that are not supported on previous versions
  - Correctness impact as well as performance impact
  - How to balance this with CNTT abstraction goals?

H/W Profiles and Performance: comparing architectures

- Following slide is from Intel's forecast on Goldmont Plus & Tremont performance comparison
- Both are x86 Low Power processor (think Atom)
  - Comparison is informative only
  - Wide range of performance gain based on test type

ISO-Frequency Single Thread Performance Improvement

(Relative to Goldmont Plus)

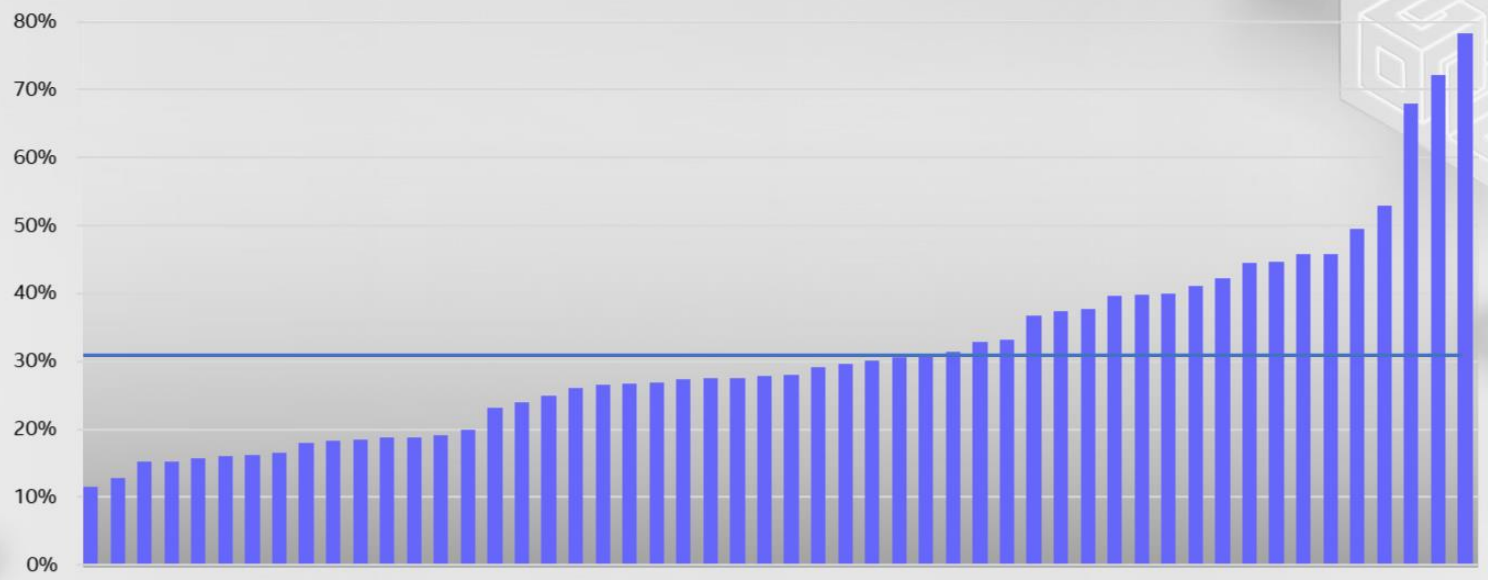

Single Thread Performance – Components of SPECint* Rate Base 2006/2017 & SPECfp* Rate Base 2006/2017

Some tests show ~10% lift, and some show ~80% lift

Based on performance projections as of October 2019 and subject to change, with SOCs at ISO frequencies
For more information about performance and benchmark results go to www.intel.com/benchmarks

Source: https://newsroom.intel.com/wp-content/uploads/sites/11/2019/10/introducing-intel-tremont-microarchiture.pdf

H/W Profiles and Performance: calls for way forward

- Can you combine normalized performance coefficients w/ high-level h/w specification to address problem?
- Is anyone aware of a better way?

Guidelines in host hardware selection: Straw-man proposal

- Intent is to create expectation that VNF workload on profile generation n-1 will work correctly, and perform with similar or better speed, on the next generation.

- Guidelines for selecting computes to match flavours
  - Select desired micro-architecture (e.g. Haswell), + system features (e.g. SmartNICs)
    - Select memory based on optimal memory controller performance
    - Select CPU model based on
      - desired # of cores to RAM ratio
      - Clockspeed + IPC gain equivalent to, or better than, previous generation
    - If present, select local disk options to provide sufficient storage for cores-to-disk ratio
  - For indivisible units (e.g GPUs, smartNICs, determine how many per host can be supported, and assign per flavour?

Why introduce memory into the selection so early?

- Lenovo paper on memory fill rules https://lenovopress.com/lp0742.pdf
  - Measures performance when under-filling or unbalancing DIMM slots
  - Sample findings for 2P system (12 DIMM slots)
    - Using just 8 slots in balanced mode results in 68% memory throughput
    - Using just 8 sots in unbalanced mode results in 34% memory throughput

- Potential for memory-induced performance problems, or over-provisioning on memory, if memory is not considered early in cpu/memory specification