Deutsche Telekom

Optimization Tools within a Software Defined Network Environment

Lessons learned from the collaboration between Deutsche Telekom and Intel in the NIMS program

Speakers:

Tarek Elbasyouny

Deutsche Telekom VNF cloudification lead designer, with a focus on Data center efficiency and media plane optimization.

Michael Holzerland

Intel Account CTO for Deutsche Telekom / T-Systems, with focus on NFV Workloads and Data center efficiency.

Agenda

- Background to NIMS
- Intel beyond CPU
- Open Source Tool for perf. measurement - Prox -
- Cross-Numa Media Plane Optimization
- VNF High Availability Models S-DC Efficiency
- Learnings from working in an layered stack
- NIMS Co-Creation Team: Data Center Efficiency
- Questions?

IMS LIFE CYCLE ISSUE DRIVES THE CLOUD TRANSFORMATION

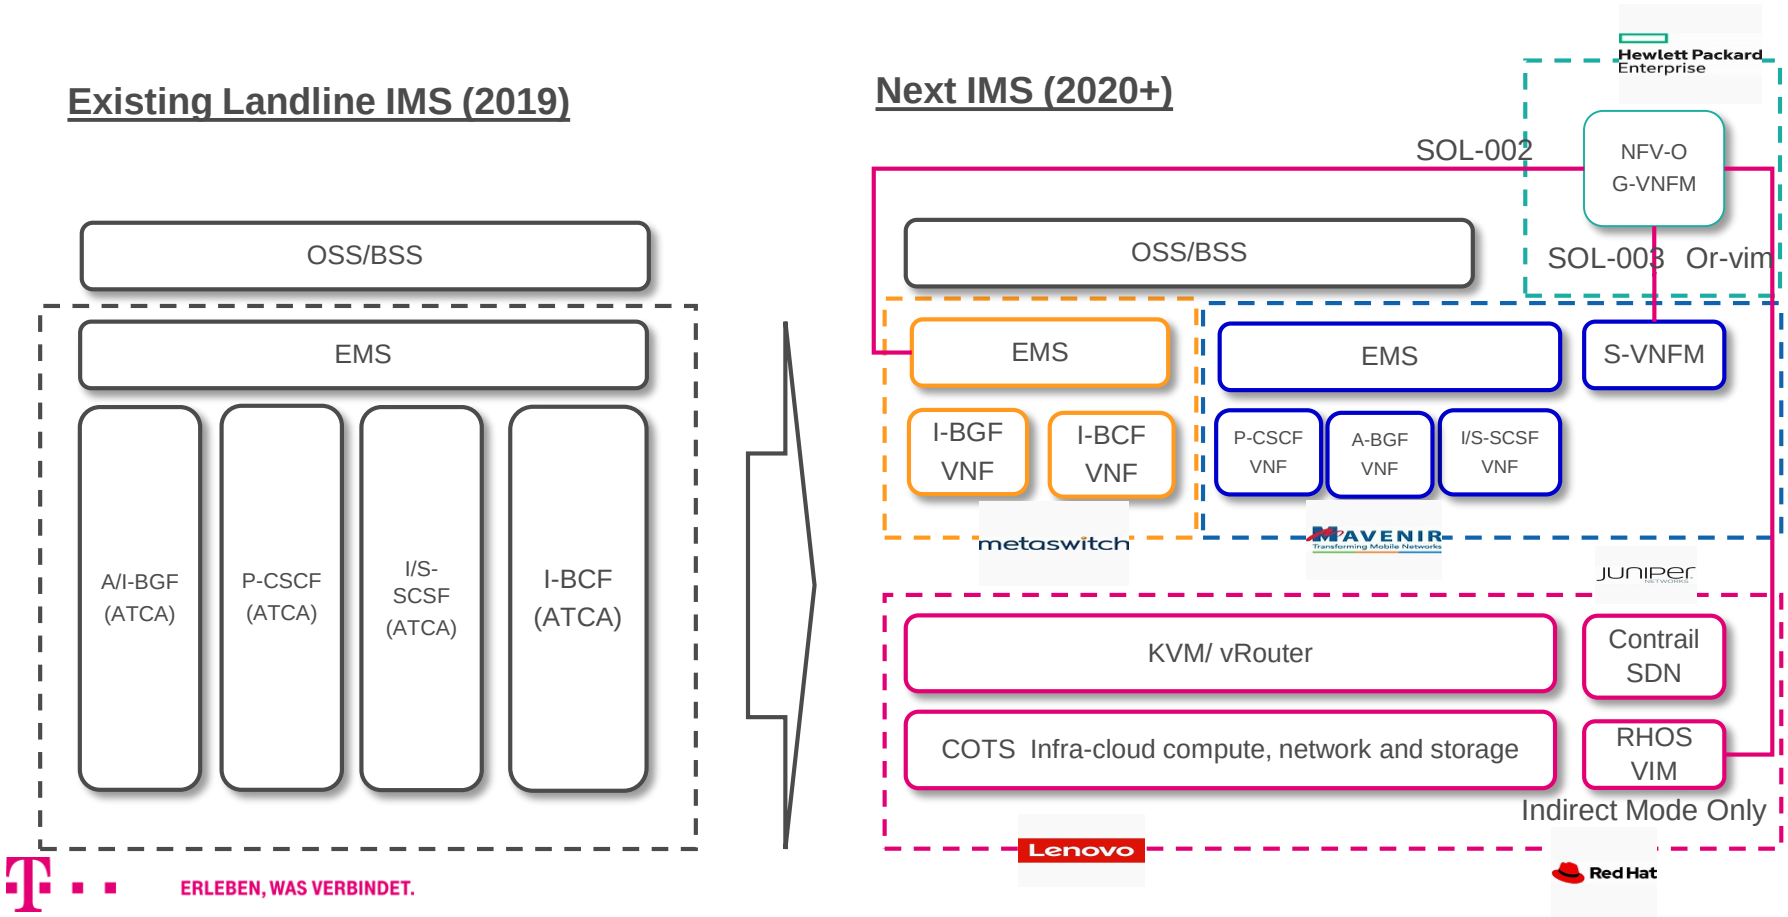

Intel: Beyond the CPU

In the next IMS project we contributed the following things to the project in addition to optimal CPUs for NFVi workloads:

- Security Hardening
- Test Lab support
- NFVi Performance Tests
- Comparison of achieved Results
- Optimization of the NFVi Architecture

Intel: PROX - Open Source Tool for Perf. measurement

PROX - is an reconfigurable DPDK Engine

- Prox can run on Bare Metal, as VM, as Docker container and is Open sourced into OPNFV
- allows us to measure: Latency, Dataplane Performance, Packet Loss, Multicast SRIOV
- allows us to find issue with non IO balanced HW, do internal testing of switching solutions.

Generator is sending UDP (1 flow) packets (64 bytes) to SUT. SUT sends packets back

Test	Speed requested	Sent to NIC	Sent by Gen	Forward by SUT	Rec. by Gen	Avg. Latency	Max. Latency	Packets Lost	Loss Ratio	Result
1	100.0% 14.881 Mpps	2.106 Mpps	2.106 Mpps	1.421 Mpps	1.421 Mpps	1787 us	1827 us	2957028	32.63%	FAILED
2	50.0% 7.440 Mpps	2.133 Mpps	2.133 Mpps	1.421 Mpps	1.421 Mpps	1794 us	1819 us	3096515	33.65%	FAILED
3	25.0% 3.720 Mpps	2.140 Mpps	2.140 Mpps	1.421 Mpps	1.421 Mpps	1789 us	1812 us	3092732	33.63%	FAILED
4	12.5% 1.860 Mpps	1.860 Mpps	1.860 Mpps	1.421 Mpps	1.421 Mpps	1675 us	1697 us	1885790	23.60%	FAILED
5	6.2% 0.930 Mpps	0.930 Mpps	0.930 Mpps	0.930 Mpps	0.930 Mpps	35 us	102 us	5	0.00%	SUCCESS
6	9.4% 1.395 Mpps	1.395 Mpps	1.395 Mpps	1.395 Mpps	1.395 Mpps	44 us	110 us	5	0.00%	SUCCESS
7	10.9% 1.628 Mpps	1.628 Mpps	1.628 Mpps	1.421 Mpps	1.420 Mpps	1663 us	1703 us	887060	12.69%	FAILED
8	10.2% 1.511 Mpps	1.511 Mpps	1.511 Mpps	1.421 Mpps	1.421 Mpps	1691 us	1719 us	387690	5.97%	FAILED
9	9.8% 1.453 Mpps	1.453 Mpps	1.453 Mpps	1.421 Mpps	1.421 Mpps	1673 us	1703 us	138028	2.21%	FAILED
10	9.6% 1.424 Mpps	1.424 Mpps	1.424 Mpps	1.421 Mpps	1.421 Mpps	1669 us	1708 us	13181	0.22%	FAILED
11	9.5% 1.410 Mpps	1.410 Mpps	1.410 Mpps	1.410 Mpps	1.410 Mpps	46 us	159 us	5	0.00%	SUCCESS
END	9.5% 1.410 Mpps	1.410 Mpps	1.410 Mpps	1.410 Mpps	1.410 Mpps	46 us	159 us	5	0.00%	SUCCESS

⟶ PROX helps to compare results from Measurements across Telcos on a Worldwide base

CROSS-NUMA MEDIA PLANE OPTIMIZATION

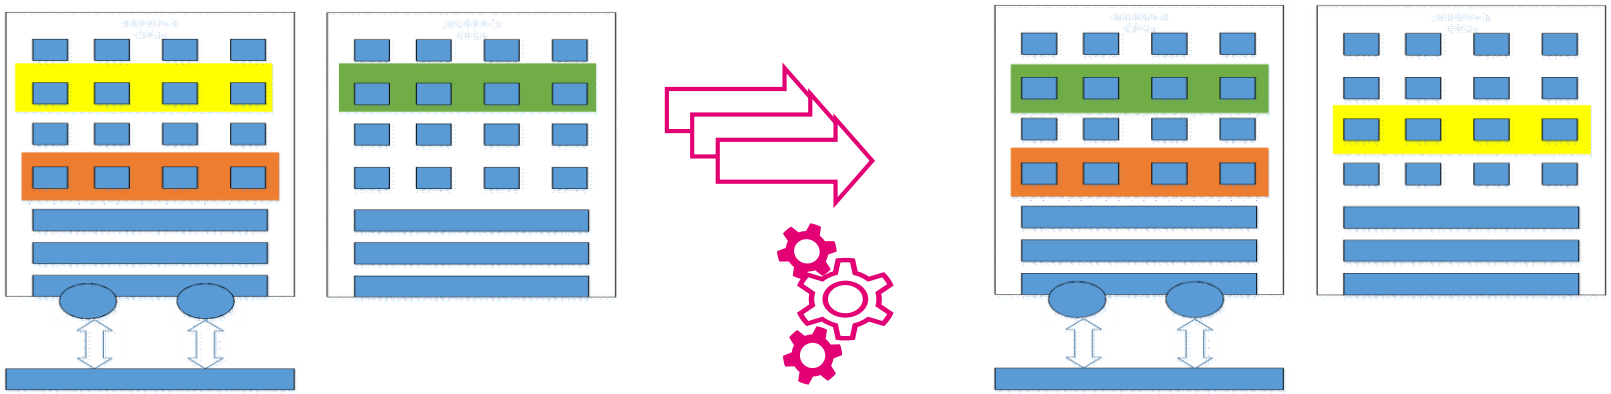

- Throughput of small size VoIP packets (~128 bytes) could degrade considerably, when the VNF and the SDN forwarder/NIC are not co-located on the same NUMA node.
- Control Plane (sip and diameter signaling) use bigger size packets and are less sensitive to cross NUMA
- Testing is essential to measure the exact throughput loss ratio in the specific VNF and NFV-I setup.

- In the Next IMS project, various options are under consideration to enforce scheduling of media-handling VMs (i.e. A/I-BGFs) on NUMA-0, where the vRouter runs:
  - A: Physnet*: Bind external RTP media plane networks to NUMA-0 (based on Openstack Rocky enhancement)
  - B: Pinning on NUMA-0 Vcpu ranges (Less scalable solution)

*https://specs.openstack.org/openstack/nova-specs/specs/rocky/implemented/numa-aware-vswitches.html

VNF HIGH AVAILABILITY MODELS- DC EFFECIENCY

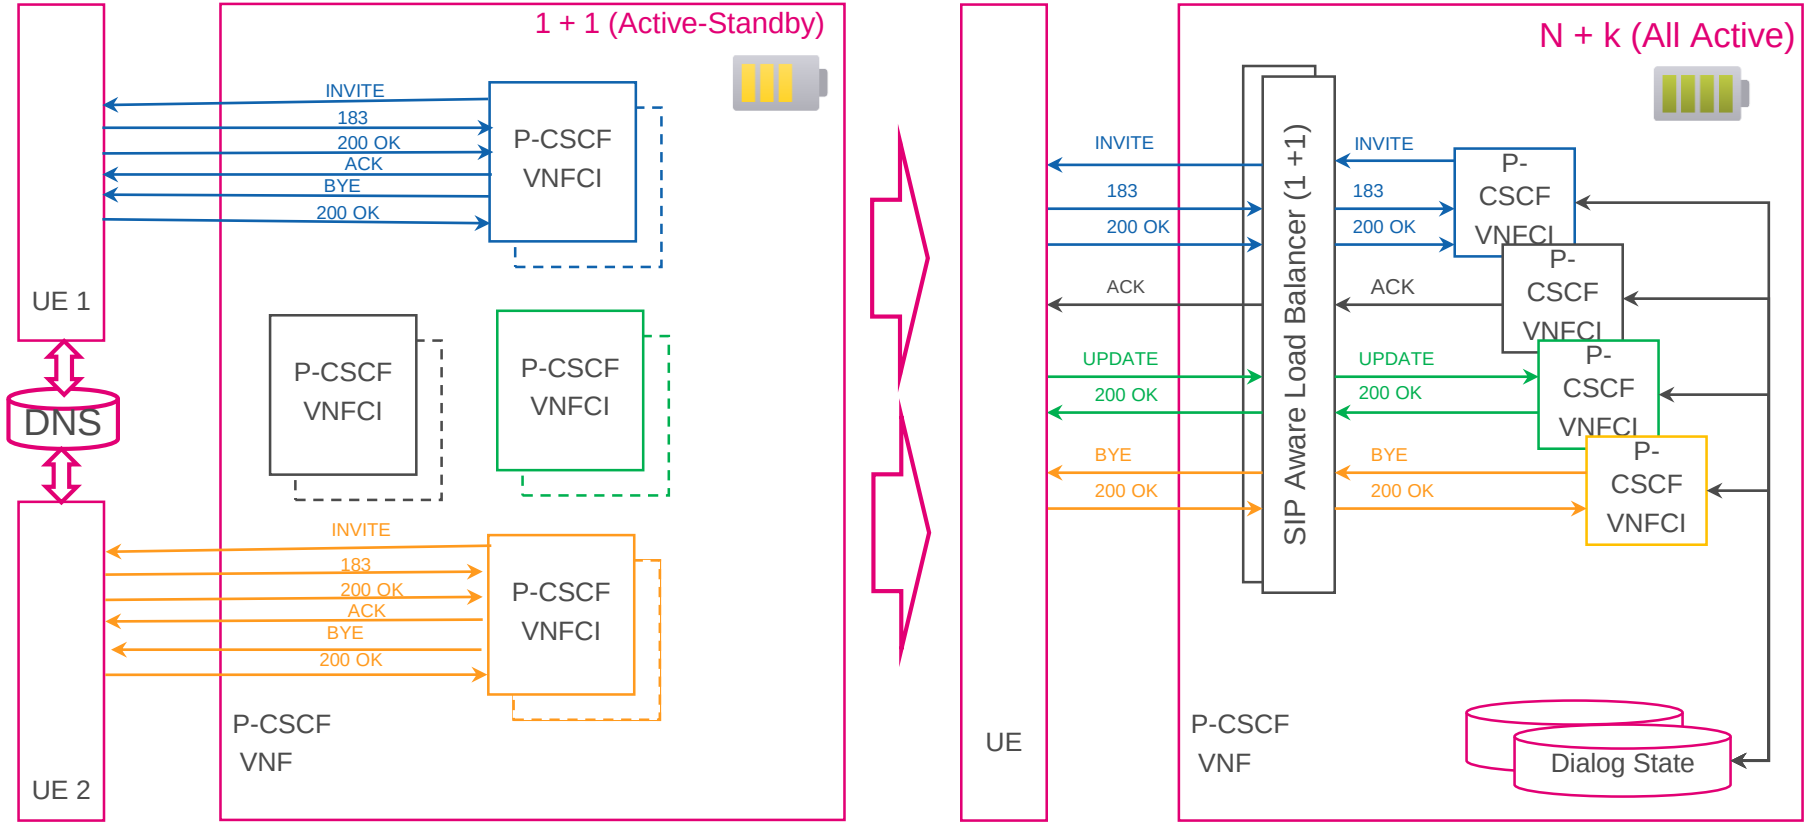

Learnings from working in layered stacks:

Layer	Enabling order	Shift Left
Application	Start in parallel where possible	Use/develop simulators**
SW platform*		Early upstreaming, pre-agreed downstreaming; as appropriate enable HostOS, VIM and orchestration
HW platform		Early Samples

* There include all SW needed to run onboarded apps in production (not just bare host OS)

** SSG does for long. There are SW tools controlling CPU cycles existing for years so if we wanted to test VNFs resilience on different core frequencies…

NIMS CO-CREATION TEAM: DATA CENTER EFFICIENCY

Scope

- NFVI HW and software infrastructure, and management of resources.
- VNF design and deployment models
- Tunable aspects of the protocol flows and resultant traffic models
- HA/GR architecture for the NFVI and VNFs
- Security architecture aspects
- Metrics and analytics tooling and policies

Process

- Each measure is verified by performance tests on the NIMS lab systems prior to adoption for the production system.
- The proposed measures will become part of the world-market features from whom vendor ever they will be provided

Work Items

- Maximize per-server throughput for media-plane VNFs
- Migration to N+k HA
- Optimization of traffic flows
- Minimizing analytics impact

Questions ?

Legal Disclaimers

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at www.intel.com

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.Notice Revision #20110804

Intel, the Intel logo, are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others

© Intel Corporation.


Deutsche
Telekom

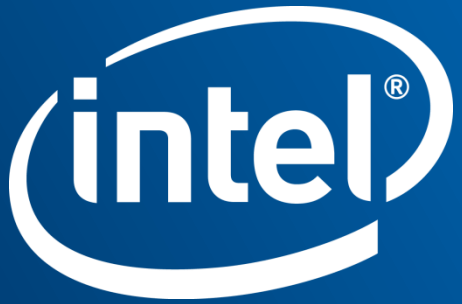
intel