HEAVY READING

White Paper

The Virtual Ascent of Software Network Intelligence

Prepared by

Jim Hodges
Senior Analyst, *Heavy Reading*
www.heavyreading.com

on behalf of

WIND RIVER

www.windriver.com

July 2013

Introduction

Although less than 12 months ago the all-IP network roadmap didn't embrace the concepts of network functions virtualization (NFV) and software-defined networking (SDN), it's now readily apparent that both will play a major definitional role going forward.

This is because, while NFV is more relevant to Layer 4-7 and SDN to Layer 2-3 networks, they are complementary and both yield similar benefits, including providing low-cost elastic scalability, enabling end-to-end service orchestration via policy control and reducing network administration costs.

It's also noteworthy that both approaches are to a great degree being shaped and driven by network operators themselves. While the telecom analyst community has often been critical of the pace and scope that network operators have adopted in implementing next-gen technologies, these operators should be given full marks for driving the creation of innovative "game changing" architectures that best meet their business and technical needs on a short-term and long-term basis.

This includes the timing of both SDN and NFV. Specifically, we mean recognizing that critical capabilities such as the ability to perform highly complex network intelligence functions can be accomplished utilizing a software-centric model. While only a short time ago this was not possible, multicore solutions from vendors such as Wind River and Intel are now mature enough to make this a commercially viable option.

Software Network Intelligence: Drivers

As noted above, software-based network intelligence is poised to dramatically change product design. This is large part because foundationally, the future mode of operation (FMO) network architecture model is more software-centric than ever before. For example, as shown in **Figure 1**, rather than relying heavily on performing network intelligence functions in hardware, software-based intelligence supports the same function purely in software, without any linkage to a specific processing platform.

Without this capability, we believe the ability to support product virtualization aligned with the NFV model and common policy and routing control advocated by SDN would not be implementable.

Figure 1: Traditional vs. Future Mode of Operation Network Architecture

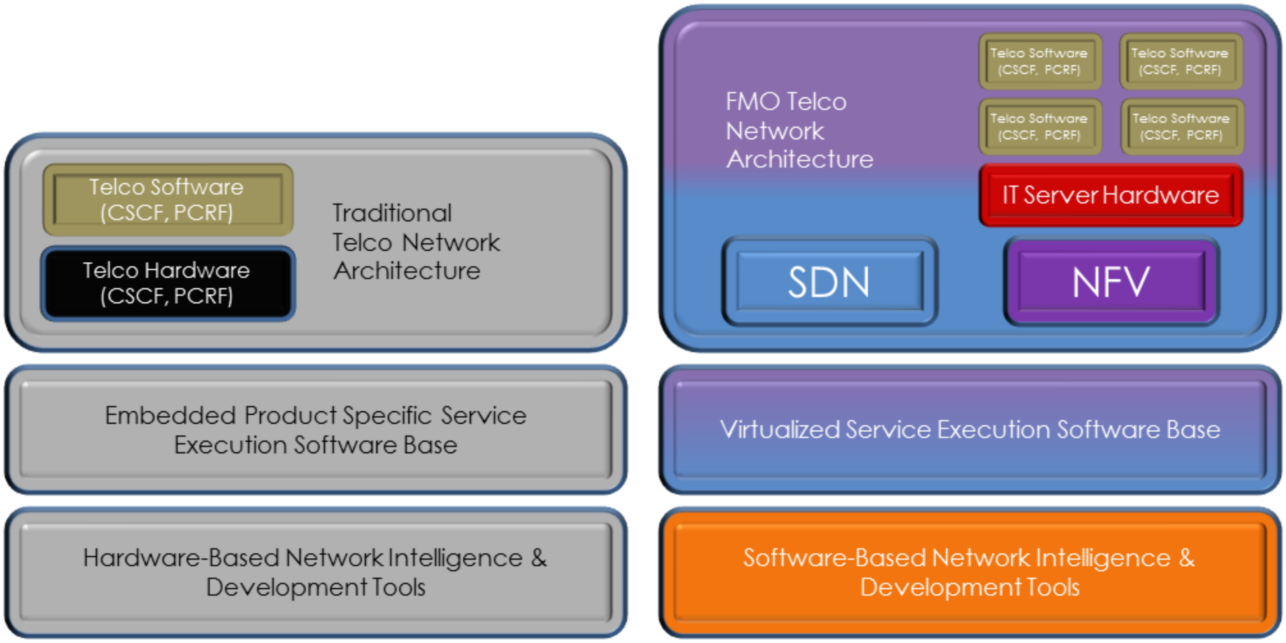
Telco Software (CSCF, PCRF)
Telco Hardware (CSCF, PCRF)
Traditional Telco Network Architecture
Embedded Product Specific Service Execution Software Base
Hardware-Based Network Intelligence & Development Tools
FMO Telco Network Architecture
Telco Software (CSCF, PCRF)
Telco Software (CSCF, PCRF)
Telco Software (CSCF, PCRF)
Telco Software (CSCF, PCRF)
IT Server Hardware
SDN
NFV
Virtualized Service Execution Software Base
Software-Based Network Intelligence & Development Tools

Source: Heavy Reading

One of the key tenets driving this telco FMO network architecture strategy is the realization that IP service delivery and administration is not sustainable utilizing a traditional telecom architecture model.

In order to be competitive both now and in the future, telecom operators must adopt a new model leveraging SDN and NFV that is highly scalable, application revenue positive and extensible.

Given that the telecom industry has relied on performance and design innovation from multicore vendors for major technology launches such as 3G and 4G, the evolution to software-based network intelligence and decoupling network intelligence from hardware may seem a relatively inconsequential event.

Still, as **Figure 2** illustrates, compared to a hardware based approach, software-based network intelligence delivers improvements in virtually all key network and business metrics, so we consider the impact profound.

Figure 2: Software-Based vs. Hardware-Based Network Intelligence

ATTRIBUTE	HARDWARE-BASED NETWORK INTELLIGENCE	SOFTWARE-BASED NETWORK INTELLIGENCE
Network Operation (Opex)	Higher cost: More nodes to administer.	Lower cost: Fewer nodes to administer.
Network Build & Scale (Capex)	Higher cost: Must scale hardware in both Network and IT domains.	Lower cost: Scale in IT domain – single network scale model.
Service Delivery	Longer: New service introduction is more complex and takes longer. Testing must include both software and hardware debug on a per-service basis. Hardware redesign, if necessary, would lengthen service delivery considerably. Not well suited to cloud service delivery, since the cloud assumes a hardware-agnostic service delivery model. Maintains the "status quo" service revenue model, which is often negative.	Shorter: Separation of software delivery from hardware means validation and time-to-market timeline is condensed and provides a more flexible model. Optimized for cloud service delivery and testing. Represents the greatest opportunity to maximize revenue of new services.

Source: Heavy Reading

We believe these last points related to service delivery are especially important. In talking to network operators, consistent with our view they frequently single out the need to support and monetize cloud based services as the key driver.

This was also strongly enforced in a recent *Light Reading* benchmark survey. In this survey, we asked more than 100 operators to rank the importance of a number of drivers for developing virtualization capabilities in future products.

The rankings, based on percentage of "critical" responses, were:

- Cloud-based application delivery (58 percent)
- Cloud-enabled policy control (42 percent)
- Separation of application hardware from software (39 percent)
- Enhanced access to cloud applications (33 percent)
- Portfolio consolidation (21 percent)

In any transition of this magnitude, the first question that must be tackled is: Where to start? In a telecom network intelligence context, this entails deciding which software-based network intelligence functions are necessary to support these cloud services.

While these priorities can certainly vary operator by operator, often the two key capabilities identified are policy control and DPI. This is not surprising since, as **Figure 3** shows, these capabilities represent essential underlying capabilities to support the identification and prioritization of various flow types, protocol configurations and applications.

Figure 3: Utilizing DPI & Policy Control

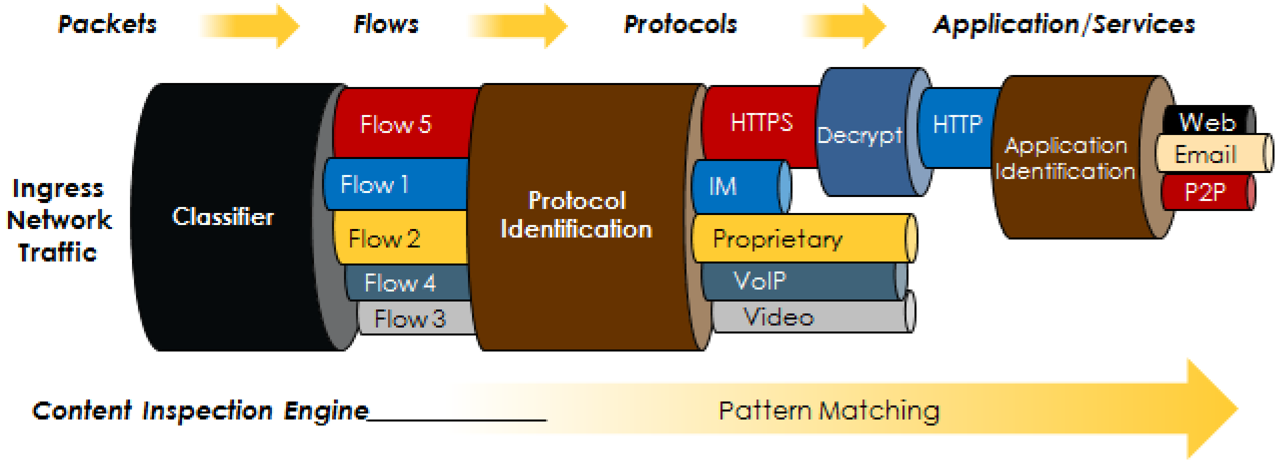

Source: Wind River

It is also pertinent to note that SDN is now introducing the concept of service chaining, which extends these capabilities by adding a layer of personalized analytics data (e.g., AAA and HSS) to enable routing of packet flows differently based on subscriber service requirements vs. the traditional static common network routing approach.

This chaining approach also enables operators to offer end users customized value-added services such as virus scanning, personalized firewall and content filtering services. This not only represents a potentially groundbreaking approach to how services are delivered in the years to come, it will inevitably result in the introduction of additional software network intelligence content inspection and pattern matching capabilities.

The next question the industry in general must address is: Which products are front-runners for virtualization? Although requirements may vary from operator to operator, we view service chaining as universally having a direct impact on a number of appliance use cases, including next-generation firewalls, DPI nodes, SBCs, policy controllers and security gateways.

This is in large part because, as illustrated in **Figure 4**, service chaining enables applications to utilize virtualized resources such as DPI and firewalls on a selective basis based on network and subscriber service requirements vs. the traditional, less efficient static routing approach.

Moreover, the benefits of leveraging SDN service chaining could be further enhanced if functions such as DPI and firewall were combined through leveraging common virtualized software-based network intelligence functionality.

Figure 4: SDN & Service Chaining

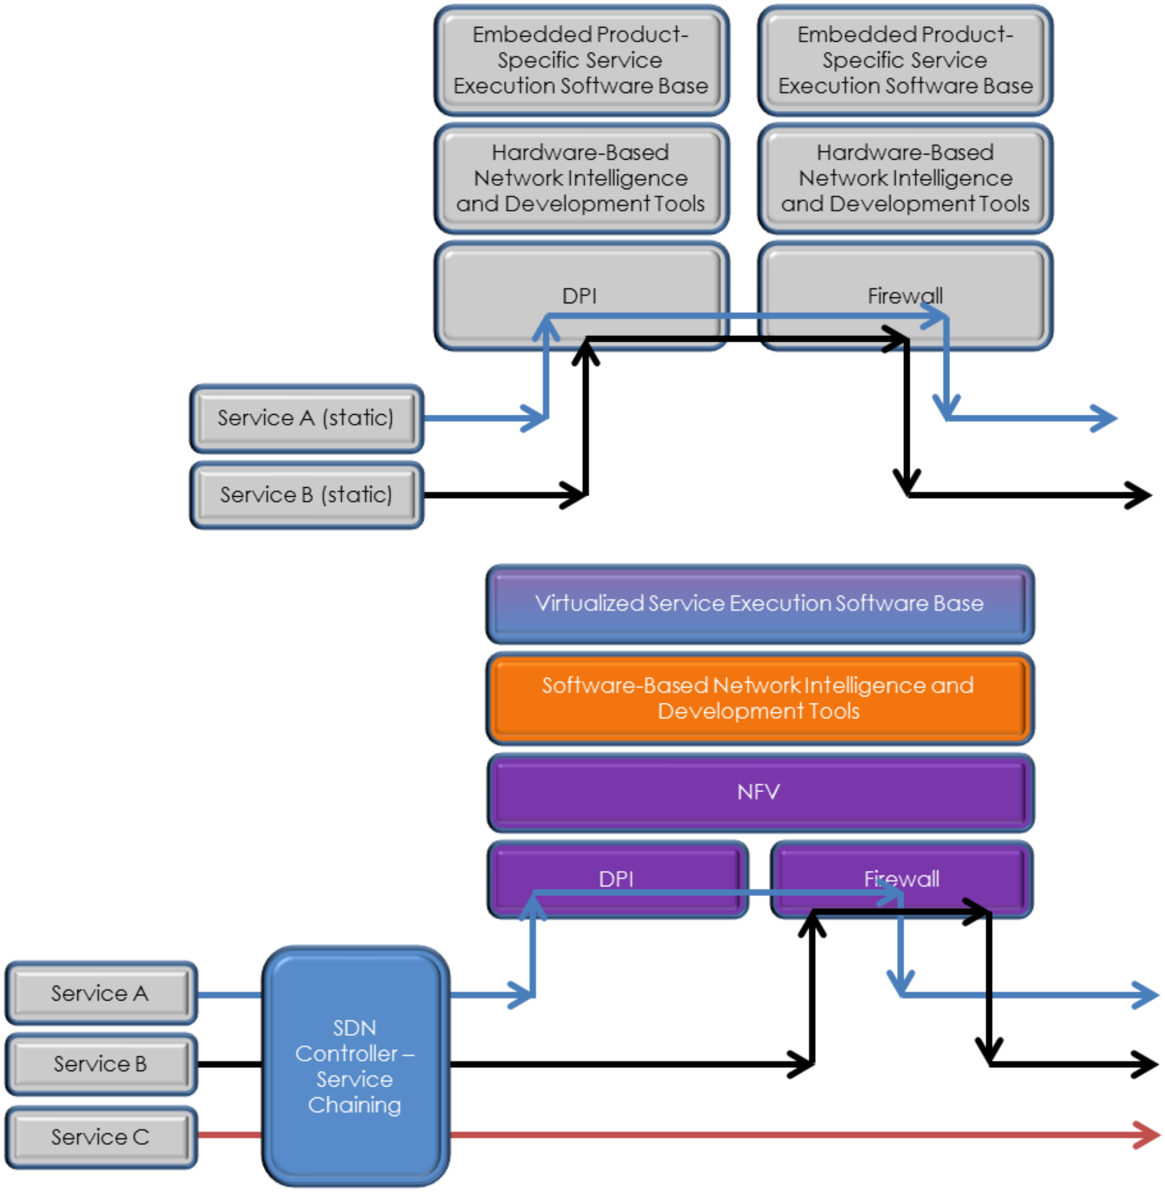

Embedded Product-Specific Service Execution Software Base
Embedded Product-Specific Service Execution Software Base
Hardware-Based Network Intelligence and Development Tools
Hardware-Based Network Intelligence and Development Tools
DPI
Firewall
Service A (static)
Service B (static)
Virtualized Service Execution Software Base
Software-Based Network Intelligence and Development Tools
NFV
DPI
Firewall
Service A
Service B
Service C
SDN Controller – Service Chaining

Source: Heavy Reading

Still, virtualization is not just limited to the nodes discussed above, but rather could be leveraged anywhere – in the mobile packet and voice core, and even the RAN. Accordingly, we view software-based network intelligence as applicable to all network equipment provider product market segments. This is consistent with the *Light Reading* benchmark in which 62 percent of network equipment provider respondents said they were currently using a common suite of network intelligence functions (27 percent) or planned to do so within 12-18 months (35 percent).

Software Network Intelligence Implementation Challenges

As we have seen, software-based network intelligence is a critical piece of the telecom architecture FMO model. And yet, while the benefits are undeniable, there are also a number of significant implementation challenges that must be addressed. Ironically, the challenges that an FMO network face are similar to those that have long impacted the traditional model: minimizing latency and achieving carrier-grade metrics.

The Impact of Latency: The introduction of new IP services will directly impact software network intelligence latency performance requirements. This is largely due to the fact that new pattern matching, application virtualization and routing will be more CPU-intensive. As well, the interaction with a hypervisor may also increase overall processing latency.

Carrier-Grade Goes Virtual: Similarly, further innovation of multicore processors and software-based network intelligence will be mandatory to maintain carrier-grade SLA metrics in a virtualized environment that product developers, subscribers and regulators alike expect.

Not surprisingly, the ability to create virtualized carrier-grade solutions is often the top issue raised by network equipment providers and telecom operators alike. This was confirmed in both the *Light Reading* benchmark, where 55 percent of all network equipment provider respondents identified this as the top concern, as well as a recent *Light Reading* Webinar, where 62 percent of the respondents chose achieving carrier-grade metrics as the greatest challenge in adopting a software-based network intelligence model for new product design.

Given the market implications, software developers are now aggressively bringing to market a first wave of new products that address these challenges. A recent example is Wind River's Linux based Open Virtualization Profile (OVP), which can run virtual applications with extremely low latency (less than 5 microseconds) (See Appendix A).

Conclusion

The emergence of NFV and SDN over the past 9-12 months has once again placed network operators at a crossroads of sorts. And to chart a way forward, they will be forced to make vital service and network strategy decisions that will impact customers, as well as relationships with network equipment provider and ecosystem partners.

However, this time around their path is clearer, given they are in many respects defining the route that they will take via participation in SDN and NFV industry working groups. And without question, these operators will embrace the power of software-based network intelligence on a much greater scale for a number of reasons. First, as we have documented, it represents a lower cost and more flexible service delivery model; and secondly, it is aligned with the software-centric vision of both SDN and NFV.

Appendix A: Wind River Network Virtualization Profile Architecture

In this appendix, we provide a summary overview of Wind River's Open Virtualization Profile Architecture (OVP).

As illustrated in **Figure 5**, Wind River's OVP implementation is based on an adaptive software-based virtualization layer. Since this approach employs open-source interfaces, it permits designers to create applications that are hardware-agnostic and can be implemented using templates and test tools defined by open-source-focused initiatives such as the Yocto Project.

Performance-wise, this open-source focus means that OVP supports a number of highly desirable capabilities, including the lowest possible latency profile, high availability and 99.999 uptime, as well as optimized virtual switching.

Figure 5: Wind River Open Virtualization Profile Architecture

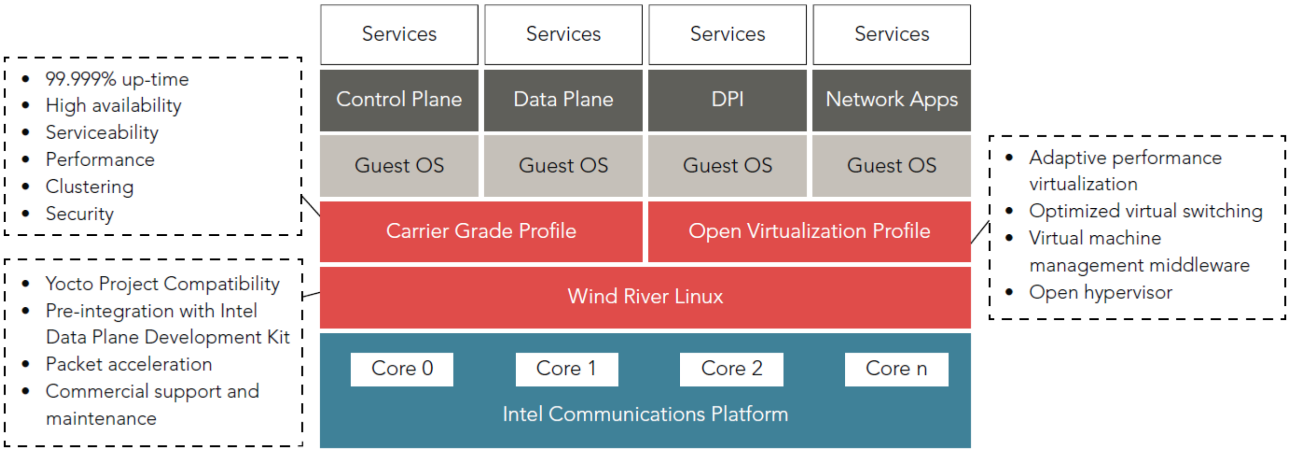

Source: Wind River

Appendix B: About the Author

Jim Hodges
Senior Analyst, *Heavy Reading*

Jim Hodges has worked in telecommunications for more than 20 years, with experience in both marketing and technology roles. His primary areas of research coverage at *Heavy Reading* include softswitch, IMS and application server architectures, signaling protocols, NFV, Subscriber Data Management, application delivery and managed services.

Hodges joined *Heavy Reading* after nine years at Nortel Networks, where he tracked the VoIP and application server market landscape, most recently as a senior marketing manager. Other activities at Nortel included definition of media gateway network architectures and development of Wireless Intelligent Network (WIN) standards. Additional industry experience was gained with Bell Canada, where Hodges performed IN and SS7 planning, numbering administration and definition of regulatory-based interconnection models.

Hodges is based in Ottawa and can be reached at hodges@heavyreading.com.