

Real Intelligence is Required for Virtualized Services

by Prayson Pate, Chief Technology Officer and Senior Vice President, R&D. Overture

The combination of falling rates and growing traffic is forcing Latin America's carriers and service providers to be more responsive to user demands. They need to leverage technology drive new revenues and profits, re-configuring their networks to cost-effectively handle increased voice and data traffic. Finding creative ways to trim costs, outsourcing non-core back-office tasks and finding the right partners, can make a tremendous difference by giving regional providers the opportunities to simultaneously grow both traffic and revenues throughout their networks.



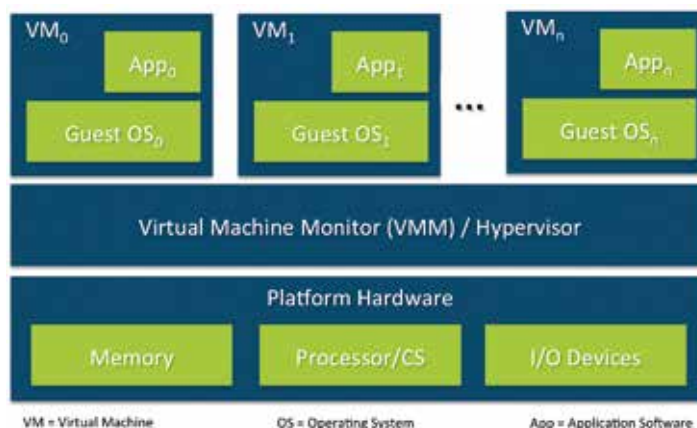
Prayson Pate, Chief Technology Officer and Senior Vice President, R&D.

Prayson co-founded Overture and brings more than 28 years of experience developing software and hardware for networking products to the Company. He began his career in 1983 at FiberLAN, moved to Bell Northern Research (Nortel) and joined Larscom in 1992. He has extensive hardware/software design and management experience, and served as Director of Engineering at Larscom before co-founding Overture. Prayson is active in standards bodies such as the MEF and IETF, and he was chosen to be the co-editor of the pseudowire emulation edge-to-edge (PWE3) requirements and architecture documents (RFCs 3916 and 3985). Prayson holds a BS in electrical engineering and computer science from Duke University and an MS in electrical and computer engineering from North Carolina State University. He holds nine patents.

Network Functions Virtualization (NFV) opens the door to radical changes in how services are delivered, with the promise of using cloud technologies, virtualized network functions and standard hardware to lower costs and drive faster times to market. These opportunities are being actively pursued by Asian communications service providers (CSPs) such as China Telecom, NTT/DOCOMO, KDDI, China Mobile, Telekom Malaysia and Telkom Indonesia, as well as by equipment suppliers such as Huawei, NEC, Hitachi and Fujitsu. However, the radical changes create a new set of challenges for the CSPs and suppliers, particularly in the area of understanding how well a service is performing, and in diagnosing the service when it is not meeting its Service Level Agreement (SLA). In addition, consumer

expectations for information on, and control of their services is high and growing. CSPs need to actively incorporate service

intelligence into their plans for new virtualized services.



Virtualization and Abstraction

Before we go deeper into collecting data and transforming it into intelligence, let us first reviewing the underlying principles of NFV. When we talk about virtualization and abstraction, what do we really mean? Some key points:

- Virtualization allows a physical resource, such as a server, to be shared by multiple software applications, each of which thinks it has access to the entire machine in the form of a virtual machine (VM). Virtualization enables statistical multiplexing of resources to lower costs, eases management of software, and facilitates resilience. However, it complicates operations if each VM is being managed by a separate entity. How do the higher-level entities gain the control that they need, without interfering with other components of the architecture?
- Abstraction allows a layer of functionality to share the essential characteristics with a higher-layer entity, but hide the details that don't matter. In particular, abstraction allows one layer to make changes to how a service or function is implemented without sharing those details with a higher layer, thus simplifying and accelerating the creation of new revenue-generating services.

Note that virtualization and abstraction may apply to all layers. Just as a layer may abstract looking up, it uses abstraction from a lower layer for delegation and gives instructions on the functions that must be provided, without concern for the details of how those functions are implemented.

Transforming Data Into Intelligence

Even without NFV, CSPs already have big challenges related to the massive amounts of data associated with today's services. Here are some particular areas that must be addressed.

Data Collection: Any useful communications service transits a variety of Network Elements (NEs), and probably traverses a range of geographical areas. Many services in Asia originate in other parts of the world, making the problem more complicated. It would be unreasonable or impossible to require that all of the NEs adhere to a common protocol and format for data collection. Instead, it makes more sense for a data management engine to collect the disparate data and process it for consumption and analysis.

However, this approach puts great demands on any such engine.

Data Storage: The network and equipment associated with a modern communications service can generate incredible amounts of data, especially in the case of some of the large Asian operators. For example, performance data may be generated for layers 1, 2 and 3, and the amount of this data may be expanded by "binning" (storage in time-sliced containers). This binning expands the amount of data by the number of bins required. Virtualized and dynamic services add another complicating factor in that the configuration may change over time, adding another dimension to the data. Finally, the AsiaPac region has half of the world's mobile subscribers, and this volume adds another factor to the problem of scale. Ideally, the required storage is minimized on the NEs and is mostly transferred to modern storage systems, where the cost of storage and processing is radically lower than on the NEs themselves.

Another aspect of data storage is the format used. Forcing all of the NEs to use a common format or schema for data is an untenable requirement on the NEs, and slows down innovation by requiring a priori definition of what data will be collected. A much better approach is to decouple the gathering and storage/schema of data, allowing post-collection analysis to occur as needed.

Data Analysis and Intelligence: Once the data is gathered the next step is to analyze it. This is a critical step because raw data is like the pieces of a puzzle; it has little value until assembled into a meaningful picture. In particular, analysis should extract relevant trends and correlations and present them to users and/or higher level systems in the



form of actionable intelligence. The analysis engine should also be optimized for efficient consumption of input data, and rapid and flexible output of synthesized intelligence to other systems.


Note that this processing has complexity beyond its sheer size. The sections below will detail how virtualization and abstraction introduce a measure of decoupling into the data, requiring the processing to provide correlation to put the data back together again.

Data Application: The final step is to apply the derived intelligence to improve the operation of the service, or to enable new services to drive more revenue. Examples include:

- Resource optimization (e.g. compute, connect, store) to maximize use of expensive capital. A service should be able to be provided using any resources that meet its SLA.
- Dynamic pricing of services to drive revenue of underused resources. A good example is traffic on the 4G/LTE radio network. This network may be very busy during the day, but is likely to be mostly idle at night. What if this unused capacity could be offered by the CSP at a discount for applications such as backup, which may need large amounts of bandwidth during those underused periods? Another example is Amazon's "EC2 Spot Instance" which provides a dynamic bidding system for resources.
- Advanced diagnostics for rapid remediation of network issues. One of the most difficult aspects of resolving an issue is elimination of the extraneous noise generated by a fault and the isolation of the actual failure. Judicious use of derived intelligence can help reduce the time for the fault resolution process.

Challenges of Virtualization

Applying cloud technologies such as virtualization, abstraction, and delegation have benefits that are desirable, widely discussed and well understood. They are essential to achieving scale and simplifying management, as well as in creating new services more quickly and at lower cost. However, the hiding of information introduces other complexities, and these complexities are not widely discussed or understood. Here are some areas that need to



be addressed:

Shared Hardware

The most fundamental issue with sharing abstracted hardware is informing higher layers about what hardware is being used when that information is relevant. For example, an operations group needs to know what unit to replace when a server fails, but that information may be hidden by a layer of abstraction. This situation is further complicated by the fact that a multiplicity of management groups or systems may have services that are using a shared and virtualized resource.

Policy-Based Resources

The virtualization situation is further complicated by the fact that resources required for a service may be reserved on a statistical basis so as to help lower the costs of a service. In consequence, a service will get the resources needed to meet its SLA, but these resources need not be dedicated or time-invariant. This leads to the next complication

Time-Varying Network Functions

Unlike today's fixed infrastructure, virtualized network functions may change over time in several dimensions, including the following.

- Elasticity or horizontal scaling: A big advantage of properly-constructed virtual network functions is that that can be dynamically scaled up or down to match time-varying service demands. This is a great asset in meeting the needs of a service,

but how can a service manager know what resources were assigned at any given time?

- Fault recovery or healing: Virtual network functions can take advantage of a pool of hosting resources to implement a 1:N resiliency model, where functions can be moved from a failed node to a spare node to recover from faults. This mobility complicates post-mortem analysis of outages.

- Migration: In addition to fault recovery, it may be advantageous to move a virtualized function in order to consolidate or optimize the use of compute resources. Any such migration must be auditable to ensure that the SLA was met at all times.

All of these changing attributes provide advantages in terms of cost, efficiency, ease of maintenance or resiliency. At the same time they complicate the management of the service because of the need to tie the service back to the state of the infrastructure at any given time.

What is Needed

We have discussed the needs and difficulties of data collection and analysis, as well as the benefits and challenges of virtualization. What is needed to extract data from virtualized services and transform it into intelligence to make those virtualized services meet the needs of their users?

- An efficient way to gather data from physical and virtual resources. Making use of the data currently available from NEs and from virtualization tools such as OpenStack is the first step to service intelligence. We also

need to define a common way to extract data from virtualized network functions.

- Format-independent storage of data. Requiring all data to fit a pre-defined model slows down the construction of systems and impedes innovation.

- Ability to analyze the data to provide historical and cross-source. Because the virtualized resources may be shared and may vary with time, effective analysis will require the ability correlate data, both between sources and over time.

- Efficient means to extract and apply the derived service intelligence. Finally, we must have a simple and effective way to use the service intelligence that we worked so hard to get. Even more important, we must ensure that we have enabled innovation in the application of service intelligence.

New Tools Open the Door to New Opportunities

Virtualizing services provides an opportunity for CSPs in the AsiaPac region to create new services more quickly and at a lower cost, but there are hurdles on the way. We have discussed some difficult challenges for deriving real intelligence from virtual services. Fortunately, advancements in efficient machine-machine interfaces, graph databases and web-enabled architectures provide new tools for solving these challenges. By using these tools we can both realize the benefits of virtualization and provide the service intelligence needed to meet the demands of the operator and end user. ●

