White Paper
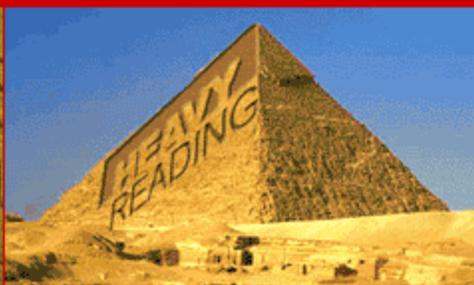
The Innovative Edge: The Rise of Cloud-Based Services

Prepared by

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

on behalf of

JUNIPER NETWORKS®

www.juniper.net

September 2012

Service Delivery on the Edge

Network service providers are wrestling with an unprecedented and seismic shift in how both personal and business services and applications are delivered within their own network and in the cloud. And as with any major market shift, levels of progress made by service providers have varied. Whereas some have embraced the new market landscape to drive further services innovation, others have not.

Operators who have not yet embraced the shift have often treated the process as an invidious exercise and continue to focus on a connectivity oriented model which is declining in value quarter by quarter. Unfortunately for these same providers, the potential for even greater revenue loss and squandering of monetization opportunities will only increase going forward as demand for video, personalized applications and cloud-based business services continues to grow.

In order to avoid this scenario, network service providers need to consider creating more progressive business strategies and monetization models. These include considering subscribers' unique preferences and needs in the service definition process, making changes to their business and service monetization models, and also addressing an often overlooked point – the importance of the underlying architecture that supports the new service offerings.

Accordingly, in this white paper we explore the impact of these factors with a focus on how the traditional architecture of service provider edge networks is undergoing an aggressive period of evolution and shifting focus from simply a point of network connectivity, to a vital services creation and innovation point.

A second, accompanying *Heavy Reading* white paper – **Implementing the Innovative Edge for Cloud-Based Services** – expands on the concepts presented in this document, defines best practices and principles to assist service providers in making the transition and provides use cases illustrating the principles in service provider deployments.

Subscriber-Driven Services

In the new world of agile applications and service offerings, being responsive to end customers' specific preferences and delivery requirements is part of the new service providers formula for success. Historically, meeting customer expectations has always been a formidable challenge for service providers. This is now an even more difficult challenge given the very the notion of a subscriber has changed as well over the past several years.

While in the past individual consumer and business subscribers have been exclusively tethered to the network services provider, today these same subscribers have in most cases forged equally powerful alliances and allegiances to OTT application providers to gain access to a variety of personal and business applications, social networking, video/multimedia, and specialized business applications.

In parallel these subscribers have become not just service consumers, but are also starting to define and shape the service innovation in part due to their ability to leverage more intelligent and programmable platforms and devices. To put in context, as of September 2012, approximately 1.3 million Android devices are now being activated per day, with a total of more than 480 million activations having taken place. Conceptually, this also applies to the applications themselves given

OTT applications are designed to be highly tunable and able to adapt dynamically to subscriber preferences. While this approach and associated characteristics are still being defined and referred to by a number of loosely defined terms including most recently bring your own device (BYOD) and cloud-based applications, we believe further stages of evolution will extend the model even further.

We believe in the future there will be powerful change in the service delivery relationship between subscribers and networks. Currently BYOD deployments include a focus on gaining a better understanding of how subscribers are using software and applications for both personal and business services. And on the other side of the equation, service providers are striving to support the enhanced services with a number of improvements at a several layers of their platforms (access, core, application and OSS/BSS).

Looking to the future, we believe the next step in the evolution will be that subscribers will not solely achieve personalized service innovation within these frameworks; they will demand input and will help define and create the very scope of the services framework to a much larger degree.

Making these changes real will involve embracing important attributes of the cloud computing paradigm in both application and service infrastructures. They will inherently need to leverage the dynamic creation and management of resources required in the cloud, as well dynamic evolution of user contexts and application features in real-time in their platforms. Making the entire service delivery platform "as responsive as the cloud" will be imperative.

This stage of evolution is an end game with virtually no service delivery limitations, since it leverages inherently elastic resources, and is both driven and *defined* by individual subscribers on a much larger scale than today. Even in this model, we believe service providers must be responsible for orchestrating the service experience, and not simply exclusively allow services to be controlled by the end user.

Evidence with service providers who have begun coupling customers' specific preferences more fully into custom-targeted "bundles" (secure cloud with SLAs, self-publishing with custom filters, etc.) has shown customers are willing to pay significant premiums above "vanilla" and "a la carte" separate offerings. As a proof point, we believe BYOD, which bridges both individual and business subscriber domains, is gaining mindshare not because it is driven by corporate policy, but rather because individual employees are requesting this capability in such great numbers that it is difficult to suppress. However, the network service provider is ultimately responsible for overall service experience and billing.

Business Model Dynamics

As we have noted, traditional business models are rapidly changing and effectively becoming obsolete in response to changes in customers' application preferences, consumption patterns and time to market expectations. Therefore, we believe network operators must also innovate in creation of new more flexible business models that support an array of application consumption patterns (numbers of devices used by a subscriber, varying ranges of application features chosen, etc.) and embrace a broader supplier ecosystem in delivering the offerings to customers. This is important, since it's readily apparent that the necessary innovations will not only come from the operator's own development teams, but will be accelerated faster and made richer by finding new ways to collaborate with third parties.

While collaboration can take many forms, at a minimum we believe network operators must ultimately develop deeper and more dynamic value chains for delivering services. The most visible win/win/win outcome of this activity for network and OTT providers and consumers is an even greater degree of blending network operator applications with applications and content supplied by third parties, including independent software vendors (ISV) and over-the-top (OTT) providers. This approach inherently involves engaging aspects of "mashups" and "federation" of services that is so popular in the cloud, and can be brought to bear on a number of operator service offerings.

While this approach is perceived by some operators as more of a risk than a return scenario from a monetization point of view, we believe the model provides much greater opportunity for services innovation given consumer and business subscribers are fully engaged and seeking new ways to expand their services experience.

To meet these expectations, per **Figure 1**, it is essential operators adopt new approaches at multiple points in their service delivery infrastructures, including the network edge, to shorten the time to market for new applications and maximize opportunities to generate revenue and bottom-line profit. In an era in which applications are extremely fluid, and rely much more heavily on the cultivation of a strong partner ecosystem, the network operator can ill afford a network edge that slows down or blocks the services monetization process.

Figure 1: Extent of Subscriber Personalization vs. Application Time to Market

Source: Heavy Reading
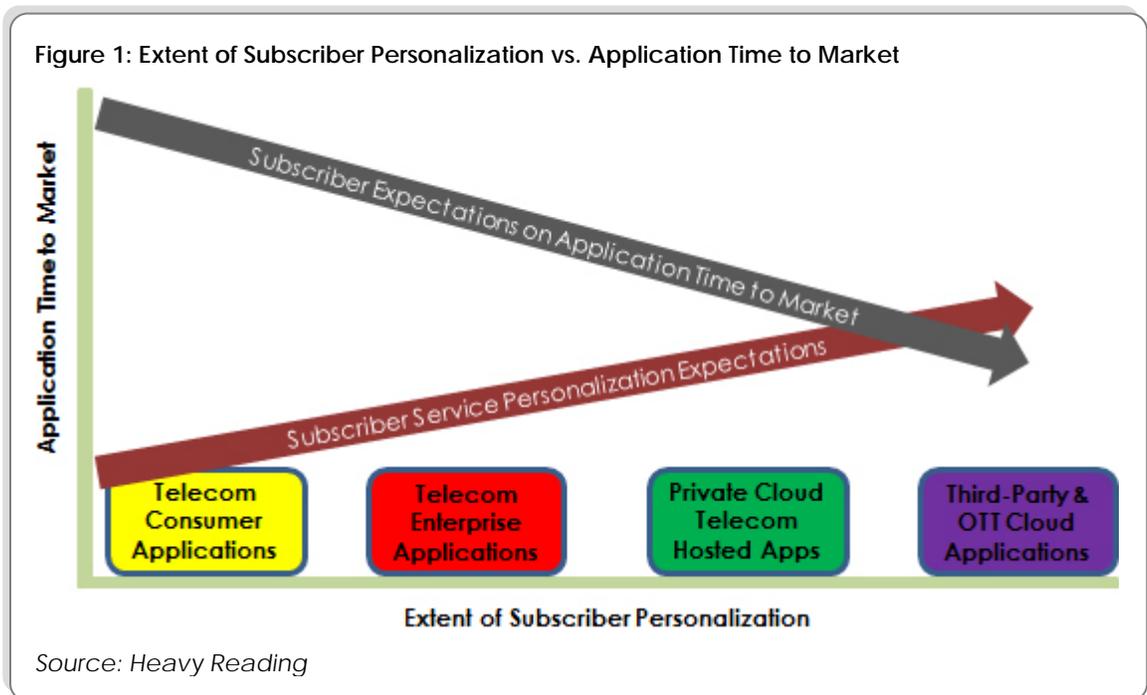
The Innovative Edge & the Rise of Cloud-Based Services

Consequently, in response to this shift in service definition and partnership models, applications are moving to a more distributed execution model to provide greater flexibility in dealing with these new demands. The best example of this is the movement of applications to the cloud which is largely accepted as the architecture that will fuel ongoing waves of Everything as a Service (XaaS) innovation. This

is a critical point since XaaS models are applicable for delivering both personal applications and enterprise business services, and therefore enables success in all of an operator's customer segments.

Not surprisingly, this approach also has profound changes on service provider network and application infrastructures. And while impacts touch a number of areas, as previously noted it is most pronounced at the network edge. This is because cloud architecture, whether private or public, inherently requires a more open and flexible software architecture to support development and delivery of a wider and more flexible range services. Since the edge is a first point of contact for service execution, we view a more modular and software-intensive Innovative Edge architecture as vitally important for enabling the adoption of application delivery in the cloud.

The Innovative Edge: Architecture & Attributes

Driven by these changes, in this section we define both the architecture and the attributes an Innovative Edge network must support. Starting with architecture, we view the Innovative Edge as not a single product, but instead a more flexible and agile combination of elements that can be deployed to meet an operator's goals.

Figure 2: Innovative Edge vs. Traditional Edge
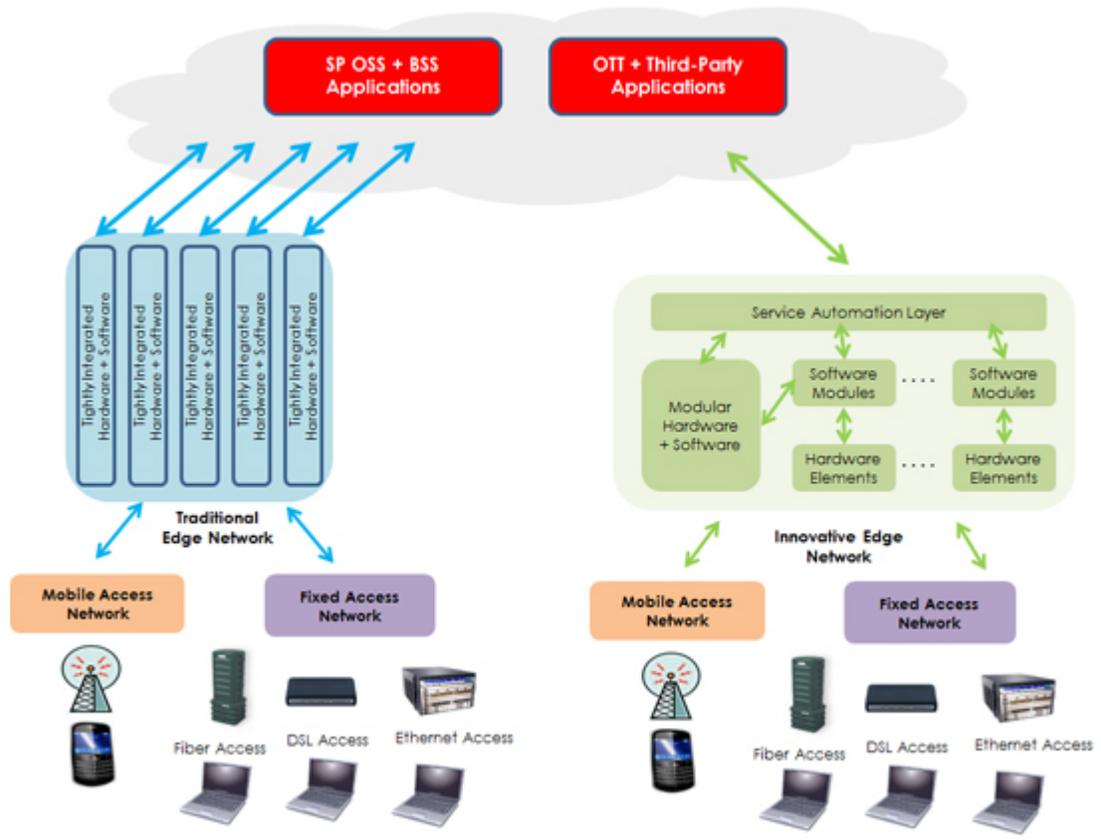
Source: Heavy Reading

In contrast to the stovepipe and "siloed" hardware and software implementations of a traditional network edge, the Innovative Edge utilizes modular software, extensible control and management planes, and modular, highly scalable hardware solutions to achieve its flexibility. This allows services to be more easily virtualized and distributed to meet specific requirements without compromising overall performance.

At a more granular level, an Innovative Edge leverages a number of visible foundation elements that a legacy Edge network either cannot support, or does only at a lower level. These include:

- Horizontal and Vertical Scalability
- Modular Software Design
- Automatic Service Creation and Deployment

Horizontal & Vertical Scalability

Scaling on the innovative edge needs to be possible along several dimensions, notably "horizontal" and "vertical." Horizontal scalability means, the edge network can "scale out" to accommodate greater numbers of customer connections at increasingly higher bandwidths, virtually without limits, so the operator isn't concerned about accommodating growth in numbers of customers and amounts of service they are consuming as the life of the investment progresses.

Figure 3: Horizontal vs. Vertical Scalability

Source: Heavy Reading
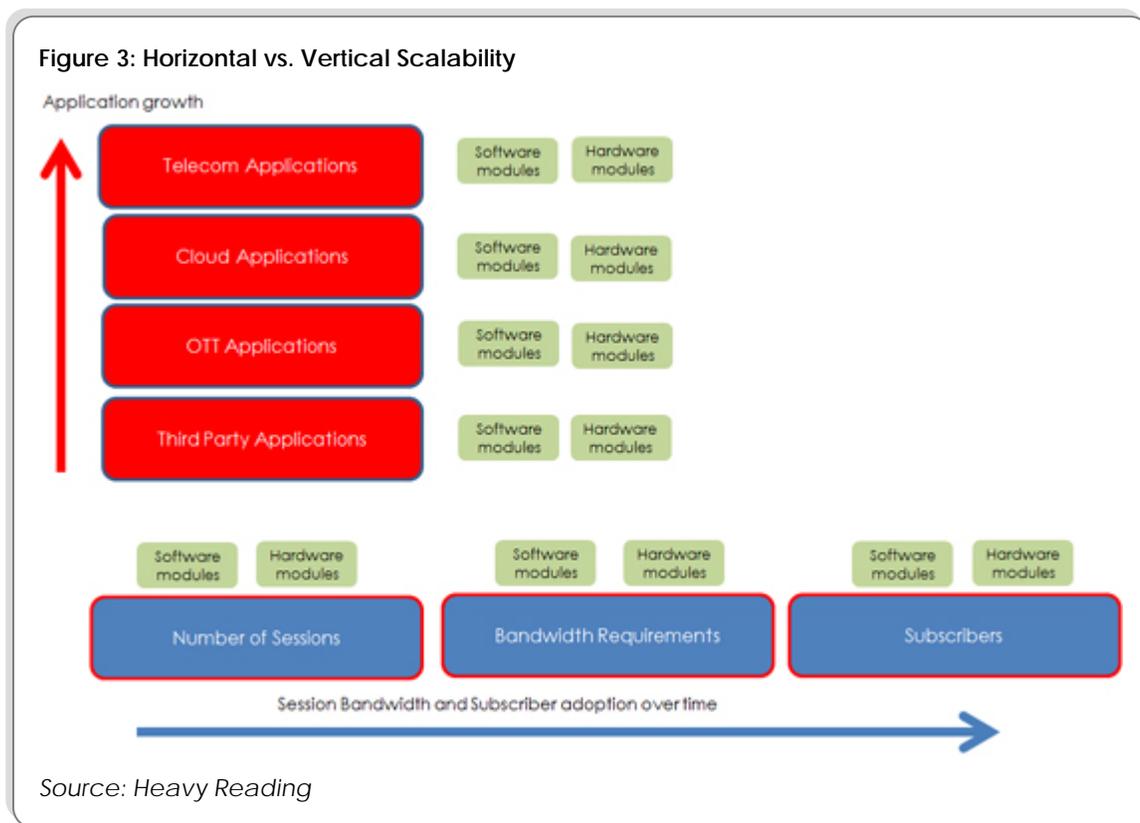
Figure 3: Horizontal vs. Vertical Scalability (Continued)

	BENEFITS	CHALLENGES
Horizontal Scalability	Faster time to market Lower total cost of ownership	Variable site sizing and bandwidth requirements Ensuring consistent performance
Vertical Scalability	Increased revenue per subscriber Application-level competitive differentiation Increase in application "stickiness"	OSS/BSS integration Service provisioning and trouble-shooting

Source: Heavy Reading

Scaling along the vertical dimension means, the operator has flexibility to "scale up" and add application and service elements to the edge – again, virtually without limits – so customer relationships can grow without worrying about the capacity of the edge to sustain the processing loads. In this way, the operator has the flexibility to add network services, customers and revenues without concerns of a major disruption in design, deployment schedules or the capex and opex allocations needed for advancing the operation.

Using the same principles deployments can start at a smaller size and achieve the same service delivery results, and also, can scale downward elegantly to optimize efficiency and avoid hardware overprovisioning in specific regional markets.
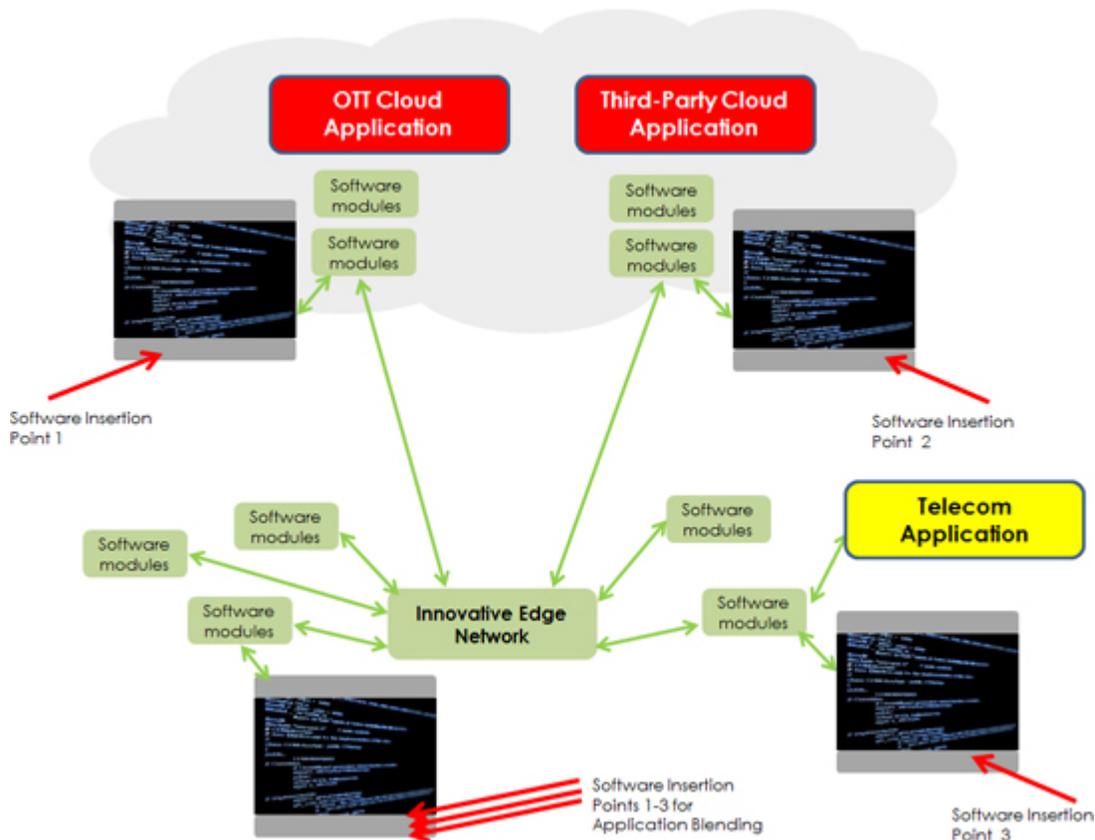
Modular Software Design

The second central element of an Innovative Edge architecture is use of open and highly modular software.

In addition to supporting a wide range of applications for vertical scalability, per **Figure 4**, software modularity supports a flexible range of insertion points allowing logic to be placed closer to customer sites and easily upgraded to support more specialized functions and new capabilities. This approach provides several benefits, including shortened service introduction timelines, customized delivery for individual subscribers and creating more agility for competitive differentiation.

One of most visible attributes of an Innovative Edge is the use of modular and well-defined software "hooks" that can be programmed and execute either in their own discrete portion of added service value, or independently of service logic in other locations for targeting of capabilities. By extending software modularity to the edge, service delivery becomes much simpler, hardware investments more leveraged, and a quantum shift in service versatility/agility and blending of applications by operators is achieved.

As illustrated in **Figure 4** below, this is especially important since moving applications to the cloud means software can run in any location without restriction (other than those inherent to the performance and functional requirements of the application). Software can run in multiple sites controlled by the network operator (for example, central offices, points of presence, regional and other data centers), as well as in other suppliers' systems running over the top and in coordination with the operator's own environment.

Figure 4: Software Insertion Strategies

OTT Cloud Application

Third-Party Cloud Application

Software modules

Software modules

Software modules

Software modules

Software Insertion Point 1

Software Insertion Point 2

Software modules

Software modules

Telecom Application

Software modules

Software modules

Innovative Edge Network

Software modules

Software modules

Software Insertion Points 1-3 for Application Blending

Software Insertion Point 3

Source: Heavy Reading

Similarly, with an increased focus on services personalization, a more dynamic software model at the edge can respond to changes in resource requirements, such as compute and storage resources for on-demand service activation, or bandwidth allocation in real time, conceptually aligned with adaptive bit rate models. The value of this is relatively intuitive and can be used to directly enhance the user's quality of experience, since the range of preferred applications varies dramatically in terms of bandwidth requirements.

From a subscriber perspective, being able to access such modular and open software environments is also desirable since it enables subscribers to manage their choices dynamically, and access capabilities that have not been available before. This is especially relevant with the advances in endpoint capabilities they have access to in both fixed and mobile services.

Automatic Service Creation & Deployment

The final differentiator between a traditional edge network and an Innovative Edge is the degree of service automation available for creating the edge configuration and the integration of software from both on-net and off-net sources. By

expanding on this dimension, operators can create and launch service offerings more rapidly, and evolve them more efficiently over time than has been possible in the past.

The customer's experience can also be enhanced by introducing a self-serve subscriber portal to allow users to adjust application settings. This vastly improves time to market, service personalization and the service provider's competitive advantage. Using a framework for automating service creation and deployment can thus become a central element in reaping greater returns for the investments made in the delivery infrastructure overall.

This approach not only shortens time to provision new services, but also minimizes the potential of provisioning errors by simplifying the process of integrating new applications into OSS/BSS systems.

Such an approach we believe will become of even greater importance in the future as applications become even more virtualized and unable to be controlled by traditional edge inflexible manual automated practices.

Figure 5: Automatic Service Creation & Deployment
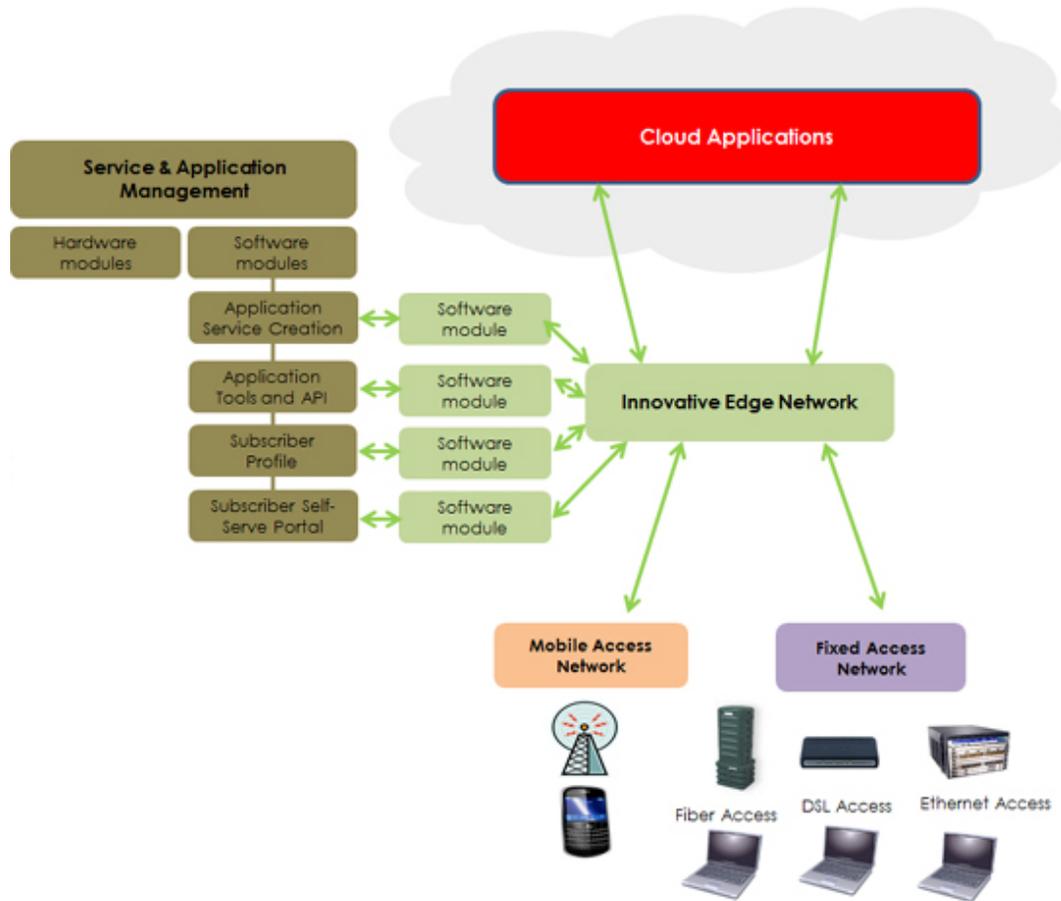
Source: Heavy Reading

As summarized in **Figure 6**, an Innovative Edge architecture represents a dramatically more flexible and extensible architectural template which delivers the agility, scale, simplicity and automation necessary for overcoming the challenges and capturing the opportunities available in the cloud-based, subscriber-centric service delivery models of the future.

Figure 6: Foundational Principle Summary

FOUNDATIONAL PRINCIPLE	TRADITIONAL EDGE NETWORK	INNOVATIVE EDGE NETWORK
Horizontal and Vertical Scalability	Limited potential. Difficult to scale upward and downward in response to OTT and cloud application demand.	Highly scalable and specifically designed to meet the dynamic demands of OTT and third-party applications.
Modular Software Design	Not supported. Assumes a tight integration of embedded hardware and software. Software insertion to support new applications is complex and problematic.	Support. Modular approach maximizes application flexibility and simplifies software insertion.
Automatic Service Creation and Deployment	Limited support due to monolithic hardware and software design.	Excellent level of support. Applications can be matched to specific hardware and software modules.

Source: Heavy Reading

# Addressing Cloud Challenges: Innovative Edge Use Cases

In this section, we review in greater depth two services use cases that illustrate how the Innovative Edge can be applied to enhance service delivery as well as monetization potential for network service providers.

Secure & Distributed Business Cloud Computing

As we have seen, deploying applications in the cloud provides an opportunity to more effectively utilize computing and software resources to enhance application delivery. Applications can be deployed with varying levels of performance and sensitivity to subscriber location.

In this way, an operator's menu of services may include offerings that benefit from resources located at the edge closest to the customer, and require the flexibility to create and secure the distributed service. Offerings that some operators are working on in this regard include auxiliary computing and storage services, interactive endpoint monitoring and data collection applications.

Figure 7: Innovative Edge & Enterprise Cloud Applications
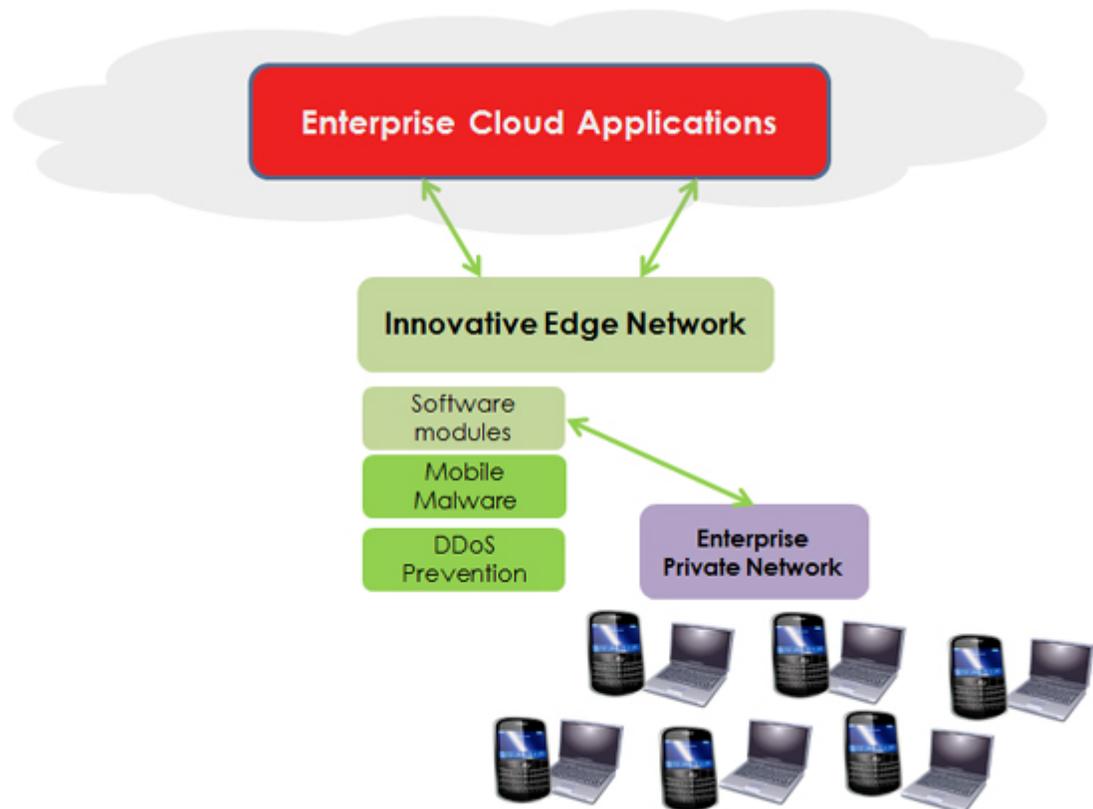
Source: Heavy Reading

Figure 7: Innovative Edge & Enterprise Cloud Applications (Continued)

CAPABILITY	TRADITIONAL EDGE NETWORK	INNOVATIVE EDGE NETWORK
DDoS Protection	Use of a standalone dedicated appliance limits flexibility, scalability and increases costs.	Leveraging in-house or third-party DDoS modular software and hardware as a virtual appliance maximizes scalability, while simplifying network monitoring and threat assessment
Malware Protection	Also requires the use of a standalone product(s) to detect, and repair malware. This approach is limits flexibility and scalability.	Leverages a modular approach to enhance flexibility and scalability.

Source: Heavy Reading

Leveraging this approach, service providers can introduce and monetize additional "value add" applications and optional security for them that have not been possible to deliver in the past. Applications can be differentiated by responsiveness to end customer needs, as well as incremental protection afforded by being delivered within a secure and protected service. In some ways, this is a natural evolution for both operator and subscriber, because service providers have a well-established market position of providing secure private networks for a variety of services.

A starting point for adding value can be protection against common threats such as distributed-denial-of-service (DDoS) attacks. Beyond this, the market potential is much greater. For example, additional "layers" of threat protection can be supplied as "add-ons" to the base. And the environment can be expanded over time by enabling additional on-demand offerings such as burst computing and auxiliary storage allocations for latency sensitive applications.
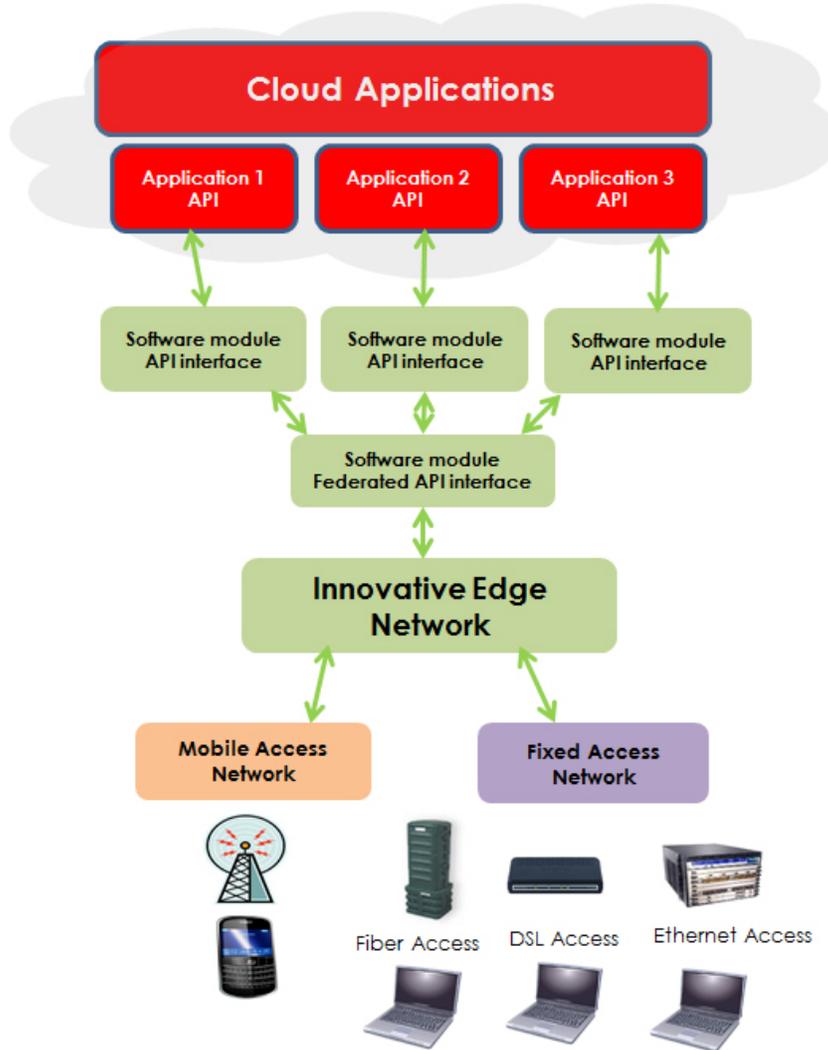
The underlying goal is to provide enterprises with a more extensive suite of tools and resources to offer additional applications and value. This approach allows enterprise customers to forego deploying specialized security appliances on site, simplifying the introduction of more extensive security measures, and reducing time to market for new applications.

Cloud-Based Applications – Federated Services & Open APIs

A key pillar in the value proposition of the cloud is that it delivers a richer ecosystem of business partners to benefit from via bridging of applications using service federation and a more open API approach. However, the introduction of a service federation model also requires that service providers' networks be able to interface to the specific APIs and possess the intelligence to aggregate them together to allow subscribers to combine the applications and deliver a truly personalized service. The network edge must also be modular and agile enough to achieve the same outcome.

Specifically, per **Figure 8**, this means the network edge must support some form of federation capabilities to allow users to access all services vs. simply relying on a single access point per application approach which is more expensive and less flexible given application location can shift in the cloud.

Figure 8: Application & API Federation

CAPABILITY	TRADITIONAL EDGE NETWORK	INNOVATIVE EDGE NETWORK
Federated Services	Limited ability to enhance federated services given the focus is on transport and access to services	Strong level of support. The Innovative Edge was designed with the assumption that services could run in a distributed and federated environment aligned with a cloud model.
Open APIs	Limited support. The traditional edge embedded design was developed before services personalization and application virtualization became key industry thrusts.	Extensive support. Modular software ensures a suite of APIs can be supported.

Source: Heavy Reading

Conclusion & Summary

The strong advance of OTT and Cloud based applications confirms that subscribers will embrace any architecture that fosters service innovation. And while network service providers harbor few if any illusions that meeting subscriber requirements will simplify over time, we believe that operators in practice must face the challenge head-on and develop more innovative and ecosystem rich application strategies.

Therefore, to aid in this transition, we further believe providers must adopt a more modular and agile architecture at the network edge capable of meeting both current and future diverse challenges associated with delivering and monetizing advanced services.

For additional information on this topic, see **Implementing the Innovative Edge for Cloud-Based Services**. This accompanying *Heavy Reading* white paper defines network service provider best practices for ensuring a successful Innovative Edge implementation based on the foundational principles discussed in this document. In addition, the white paper presents case studies of early adopter service providers who are leveraging this architecture to enable them to more effectively compete and innovate.

Appendix A: 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, SDM and analytics, signaling protocols such as Diameter, IPV6 migration, environmental initiatives 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.