

Report for Netscout

**The pivotal role of service
assurance systems in NFV/SDN
environments**

June 2015

Anil Rao and Glen Ragoonanan

1. Executive summary

Network virtualization technologies of network functions virtualization (NFV) and software-defined networking (SDN) are increasingly being prioritized by communications service providers (CSPs). CSPs agree that NFV/SDN can significantly reduce capex and opex, and there is broad agreement that the success will largely depend on the evolution of the operation support system (OSS) environment to support virtualized next-generation networks (vNGNs) – a combination of virtualized and traditional physical network assets, in which most of the core is virtualized, initially. However, most CSPs do not have clarity on the details of the capex and opex reduction. This white paper seeks to provide some insight into these costs savings from a service assurance perspective, using a supporting virtualized IMS business case.

Traditionally, service assurance operational processes are siloed and time-consuming, which increases cost overheads. With network virtualization, CSPs can move to more agile and automated trouble to resolve (T2R) and plan to provision (P2P) assurance-related processes, giving CSPs operational flexibility and efficiency. Enabling this automation will require a single “truth” open common data model that contains all network performance monitoring and management data that is granular and real time in nature. Creating and maintaining the real-time, open common data model will be important in supporting the high levels of process automation possible with network virtualization.

So far, CSPs have relied on multiple network performance data sources from network elements and management systems to drive service assurance processes; this has required significant OSS costs for the development and systems integration for capacity planning and network optimization, root-cause analysis (RCA), fault isolation and triage. Network virtualization will breed a more IP-centric infrastructure and CSPs can reduce these OSS costs by deploying virtual software IP probes that can extract crucial real-time performance intelligence at the network (traffic/packet) subscriber and application levels in IP flows. By using this multi-dimensional intelligence as the key source for the open common data model, CSPs can efficiently assure their new vNGNs.

CSPs can achieve 33% net savings by implementing a virtualized IMS (vIMS) for voice over LTE (VoLTE). Figure 1.1 provides a breakdown of the cost savings that can be achieved by using vIMS for a VoLTE service deployment and, by extension, for fixed VoIP services. Details of these costs savings are included in Section 4.2¹.

SDN use cases such as dynamic VPN or WAN configuration, self-service bandwidth on-demand (BoD), and data centre interconnect services that will provide different cost saving models, as these are more access-network-centric than IMS, which is core network. This white paper provides qualitative analysis of these SDN business cases and, as such, a detailed analysis is not in the scope of this paper.

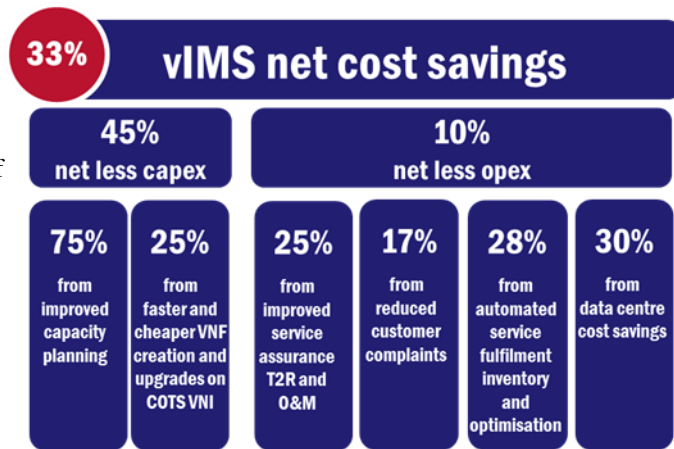


Figure 1.1: Breakdown of vIMS cost savings compared with a physical on-premise IMS for VoLTE [Source: Analysys Mason, 2015]

A better and improved service assurance OSS, which is underpinned by an open common data model using real-time and granular network performance intelligence sourced from virtualized IP probes systems, can help CSPs achieve operational flexibility and the cost optimization benefits of virtualization. This is illustrated in the vIMS business case included in this white paper.

¹ Source: Analysys Mason’s report (February 2015): [VoLTE business cases: the value of spectrum reuse, an enhanced feature set and virtualisation](#), by Glen Ragoonanan and Stephen Sale

2. Service assurance is needed to deliver operational flexibility benefits of NFV and SDN

Network virtualization is expected to bring three key benefits to the operators: service agility, cost optimization and operational flexibility. Commercial deployments of automated virtual infrastructure are still limited, but all three key aspects of network virtualization – cloud computing, NFV and SDN – are steadily advancing to commercialization. The transition to a fully virtualized telecoms network will be achieved through a process of gradual evolution. This transition is expected to take at least seven to ten years. During this period, a combination of virtualized and traditional physical network assets will co-exist and lead towards virtualized next-generation networks (vNGNs) – a combination of virtualized and traditional physical network assets, in which most of the core is initially virtualized, initially.

Figure 2.1 depicts CSPs' weighting for the three key drivers for network virtualization based on Analysys Mason's most recent research. CSPs trialing NFV/SDN accept that OSS automation will enable service agility and improve operational flexibility, with cost savings regarded as beneficial by-products.

vNGNs will be more dynamic than the current telecoms network environments, and poses significant service assurance challenges to CSPs, such as:

- vNGNs will not have network element persistence, which means network snapshots will be needed to track transient virtual network functions (VNF) and historical data will be needed for service assurance
- VNF fluctuations can occur as a result of network or service policy conflicts, and to resolve, and potentially avoid this, will require a different type of complex nested policy management
- advanced RCA capabilities will be required to correlate traditional element management systems (EMS) as well as network performance data from virtual infrastructure managers (VIM), VNF managers, SDN controllers and possibly even policy rules from NFV orchestrators.

As such, service assurance systems are expected to provide real-time network status and performance data to tackle these challenges and help CSPs to increase their operational flexibility by reducing:

- opex by implementing efficient and automated T2R processes for the vNGNs
- incremental capex by enabling just-in-time capacity augmentation of VNFs based on developing traffic demands.

Network virtualization will enable CSPs to not only become agile in terms of service creation and shorten the time they need to market new services, but will also allow them to transform their network into an open platform that can play an important role in the emerging digital economy value chain and offer B2B2B and B2B2C services. Service

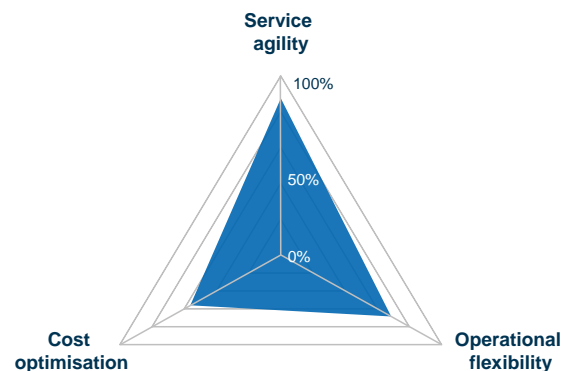


Figure 2.1: CSPs' weighting of drivers of network virtualization
[Source: Analysys Mason, 2015]

assurance systems in the virtualized environment can help CSPs make the transition to becoming digital service providers (DSPs) by enabling them to gain true end-to-end visibility of the network, service and application performance in real time, and providing the capability to proactively identify and tackle issues that will impact customer experience.

CSPS will require a new operational approach to realize the potential benefits of the dynamic vNGNs. This will be achieved by implementing highly automated end-to-end processes enabled by granular and real-time network performance intelligence.

3. Current operational processes are inefficient and will become unsustainable in virtualized environments

Telecoms networks are evolving from old circuit-switched network technology to all-IP packet-switched networks, with the network virtualization (NFV/SDN) evolution aiming to build on an IP-centric network infrastructure. To date, the IP evolution has positively influenced the economics of the CSP business, but the operational processes have not kept up with the technology, meaning that issues such as lengthy T2R resolution times and plans to provision (P2P) cycle times remain. The dynamic nature of network virtualization will make these service assurance operations unsustainable.

3.1 Traditional T2R processes are siloed and domain specific, today

A CSP's network environment typically consists of network infrastructure from multiple network equipment providers (NEPs), serving different domains (core, distribution, access), different technologies (IP, legacy TDM, SDH, WDM, etc.) and different generations (2G, 3G, 4G), each with their own dedicated EMS/NMS and lacking cross-domain correlation and RCA. To circumvent this problem, CSPs deploy independent performance monitoring (PM) and fault management (FM) systems that provide multi-vendor network equipment support. Such systems rely on vendor-specific technologies (e.g. MIB2, NetFlow, etc.), often requiring significant systems integration to gain an end-to-end network performance view. Unified, multi-domain, multi-technology network PM systems addresses this problem to some extent, but the diversity and large number of the deployed tools increases training costs, and creates the need for higher-order assurance tools, such as manager of managers. This is illustrated in Figure 3.1.

Today, Tier-1 CSPs operate hundreds of “best-of-breed” service assurance systems for monitoring and reporting on network performance, including a myriad of systems that have been obtained from a combination of NEPs, independent software vendors (ISVs), and in-house tools to support functions that are not available from a commercial off-the-shelf (COTS) software product.

As a result, the evolution of the OSS and the operational processes has been inefficient: operational and system silos have resulted in a high total cost of ownership and unsatisfactory lengthy T2R processes.

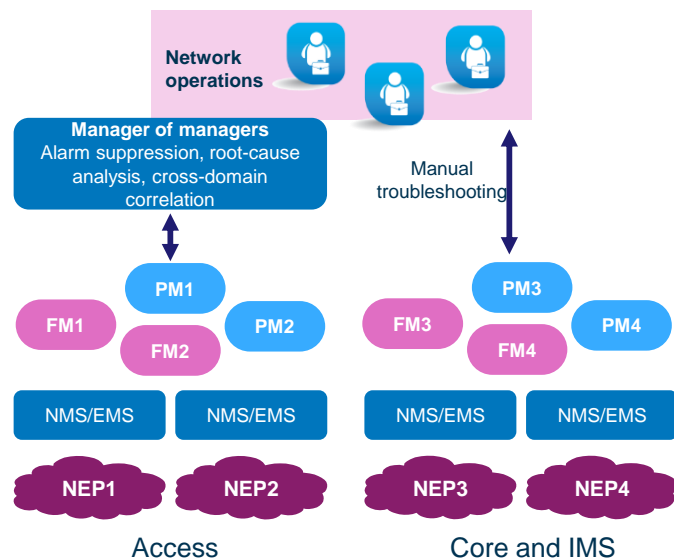


Figure 3.1: Illustration of siloed service-assurance deployment [Analysys Mason, 2015]

This has given rise to “swivel-chair management”, requiring operations personnel to use multiple screens and logons to multiple systems to conduct triage and troubleshooting. Integrating the myriad of tools tackles the inefficiency problem to some extent, but leads to high integration costs, bespoke development and scalability issues. Ultimately, CSPs incur high costs, inefficient operational processes, long T2R times and poor customer experience.

3.2 Inefficient capacity-planning processes result in unused network capacity

CSPs are under tremendous pressure to increase network capacity to support ever-increasing data consumption. But it is imperative that CSPs take steps to efficiently allocate and utilize their network capacity to improve margins. Today, there are two fundamental issues with traditional network-capacity-planning and optimization methods that inhibit maximizing network efficiency: long procurement cycles; and the use of offline network performance data as input for the planning process.

- The cycle time for capacity planning and procurement can take from three to nine months (this excludes deployment, which can take the entire cycle over a year), on average, to complete, which means CSPs have to take a cautious approach during capacity estimation and allocation. As a result, CSPs procure excess capacity to allow for this long capacity-planning and deployment cycle. Some factors that contribute to this issue are:
 - the use of traditional, manual and outdated planning tools, such as MS Excel, that takes many months to complete a planning cycle because the variables for multiple regions and technologies may need to be considered in the plan
 - the planners use different tools for each technology (3G, LTE, IP, optical, etc.) which causes further delays because of the time required for reconciliation, which is typically done manually
 - many different departments (engineering, optimization, outside plant, CTO office, procurement, CFO, etc.) are involved, which requires significant co-ordination efforts
 - with new mobile access technologies such as small-cells, LTE-A and carrier Wi-Fi coming into use, the planning process has stretched and become extremely complex, requiring significant time for meticulous analysis to find the optimal solution.
- Often network planning and optimization processes use offline network performance data as input, which renders the capacity forecasts obsolete by the time the planning process is complete. This results in an outdated capacity plan even before procurement has begun, and means that most capacity plans overestimate the capacity that is really required. This is particularly unacceptable because some contemporary service assurance OSSs offer network performance data in near real time.

Figure 3.2 illustrates some of the main inefficiencies in current CSPs' planning processes. Poor planning data leads to low confidence in capacity-planning forecasts, over-procurement of the network and inefficient capex spend, as the network is never allowed to get close to full utilization.

Analysys Mason estimates that CSPs trigger capacity augmentation processes at ~30% for core network and ~50% for all networks, to give them sufficient time to plan, procure and deploy new capacity in the network.

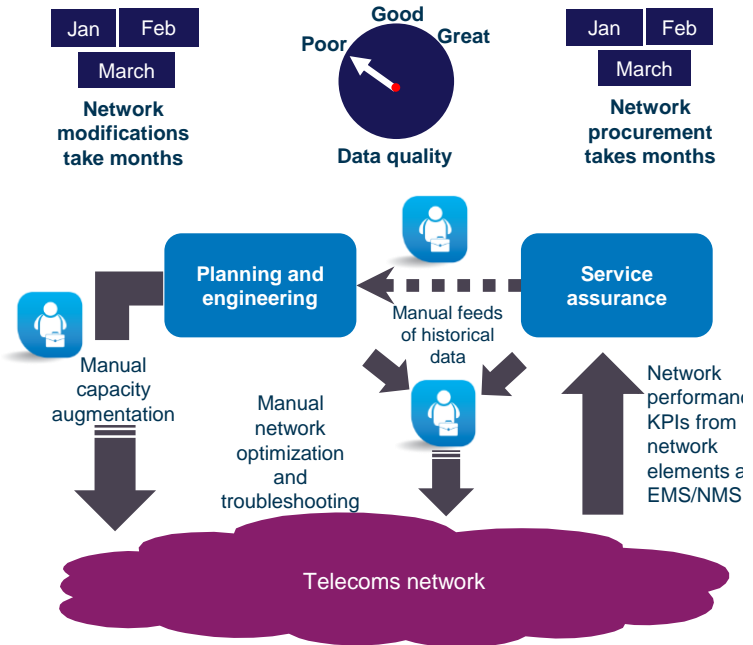


Figure 3.2: Inefficient network-planning processes
[Analysys Mason, 2015]

4. An open common data model based on granular network intelligence will increase operational efficiency in virtual networks

As the telecoms industry moves towards SDN/NFV-based network virtualization, there is a genuine need as well as an opportunity to address the operational challenges discussed in Section 3. The operational transformation is not going to be quick or easy. However, CSPs can take definitive steps towards achieving the increased operational efficiency benefits of network virtualization.

A common underlying issue in the scenarios discussed in Sections 3.1 and 3.2 is the lack of one single source of real-time network performance data for the T2R and P2P processes. In vNGNs, two main characteristics of the network performance intelligence will contribute to the success of the T2R and P2P processes: the real-time nature of the network data, and the level of granularity of the network data. Together, they will enable CSPs to improve the overall T2R process efficiency and defer capex with better-informed planning processes. Furthermore, granular network intelligence derived from live network traffic provides in-depth insight for making well-informed decisions during assurance and improves the network planning and process.

An open common data model based on these two characteristics and sourced using deep packet inspection (DPI) of IP packet flow with virtualized software probes systems can potentially improve operational efficiency in vNGNs by openly providing the data and network intelligence to other assurance and planning systems. This probe system can provide network, subscriber and application performance intelligence including subscriber consumption patterns, application server usage, websites visited, different network nodes and paths traversed to reach the destination, application errors and cause codes and protocols used.

Netscout's probe systems solution utilizes Adaptive Service Intelligence (ASI) technology providing real-time actionable intelligence by processing packet data locally, as close to the point of traffic capture as possible. The ASI data can potentially form an important data source for the single truth common data model, because it is both real time and granular in nature. Section 5 presents the Netscout solution in more detail.

4.1. Automated and informed T2R processes for vNGNs

By analyzing the packet-flow data, CSPs can obtain an application- and subscriber-centric view of all traffic types flowing in the network, in addition to other application characteristics such as their specific bandwidth utilization and conversations on an end-to-end basis. This application-centric view enables CSPs to generate metadata that includes information such as applications consumed by users, traffic volume per application, total number of errors received, number of transactions per server, and number of retransmissions occurring in a TCP-based network. This metadata enables network operations teams to interpret the associations and underlying cause-effect relationships between networks, applications and subscribers in order to understand the scope and the impact of service performance degradation or quality of experience (QoE) impacts on end users.

Because of the consistent granularity, similarity of data sets, and the level of depth of the information across the different dimensions of the network and application, CSPs can confidently design and implement automated T2R processes that can help unify the service assurance approach. Figure 4.1 illustrates the automation of an overall T2R process enabled by real-time network data from an open common data model feeding a unified service assurance² solution that integrates performance monitoring and fault management.

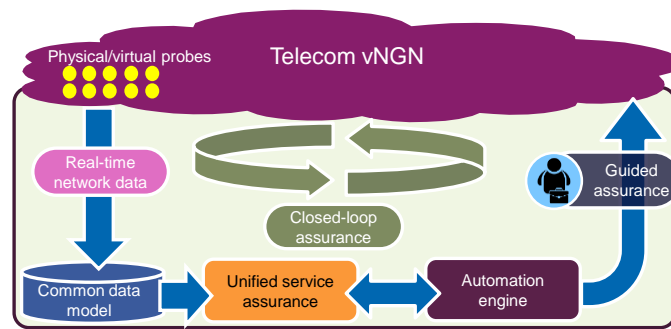


Figure 4.1: T2R automation enabled by an open common data model [Source: Analysys Mason, 2015]

The real-time nature of the pre-processing technology enables CSPs to proactively address network performance issues before impacting the customer’s QoE. To reduce troubleshooting times and increase operational efficiency in the virtualized networks, T2R process automation will be a requirement. Ultimately, CSPs will be able to achieve opex savings because the elapsed time for RCA will be significantly less labor intensive. Section 4.3 illustrates cost savings that CSPs can achieve when launching a VoLTE service over a vIMS solution.

4.2. Just-in-time capacity augmentation is possible with NFV/SDN

Another key benefit that CSPs will be able to achieve by using NFV/SDN is capex savings through more efficient network-capacity planning, which results in a deferral of network capex spend. Using real-time network and session intelligence data from the probes, CSPs can make near real-time decisions on spawning additional VNFs for capacity augmentation. This can change CSPs’ current stepwise capex spend, which overestimates capacity needs, to incremental capex in a more “pay-as-you-grow” capex model using just-in-time network capacity augmentation. Just-in-time network capacity augmentation is only possible in a virtualized network environment.

² Source: Analysys Mason’s report (April 2015): Unified service assurance: understanding the business drivers, key initiatives and benefits by Anil Rao

In a physical network environment the excess network capacity is locked to where it is provisioned in the network and cannot return to a common network resource pool; however, this is possible when NFV performs the main network function and SDN assists in the logical bandwidth capacity function in a single medium, such as fiber, copper or microwave. Just-in-time capacity augmentation allows higher network utilization before triggering a lengthy capacity procurement cycle. For IT, servers are typically run upto 80% utilization before augmentation. This can be up to 70% for CSPs with NFV/SDN (see Figure 4.4).

Figure 4.2 illustrates the automated just-in-time capacity-augmentation process in a virtualized network. Automation of the capacity-augmentation process can reduce opex to provision services on the fly, as well as ensure capacity-planning processes have high quality data in near real time.

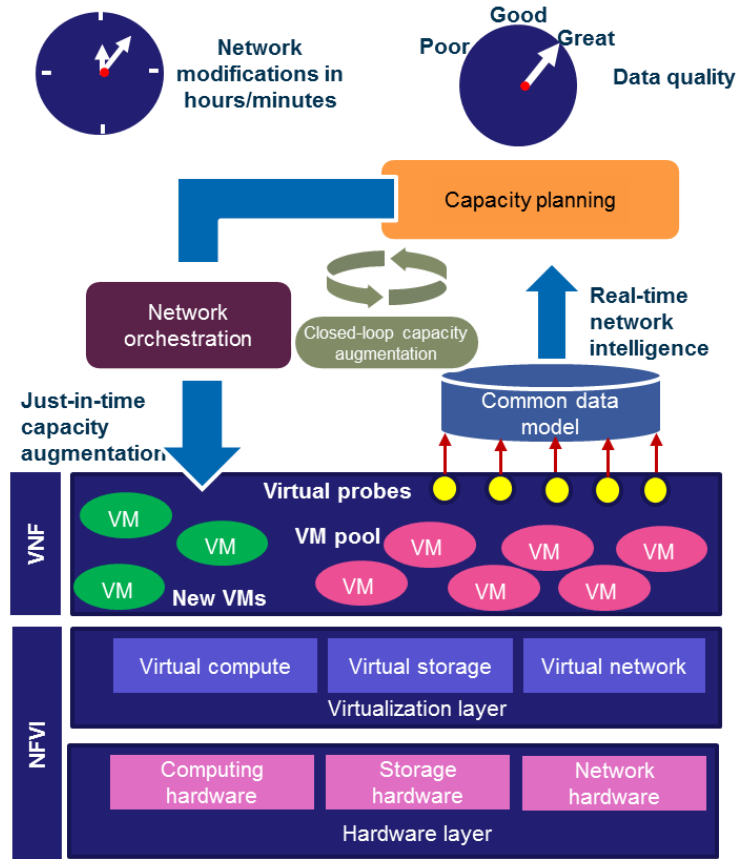


Figure 4.2: Just-in-time capacity augmentation using a common data model in a virtualized network [Source: Analysys Mason, 2015]

4.3. CSPs can achieve 33% net savings by implementing a vIMS for VoLTE; 45% in net capex saving and 10% in net opex savings

Analysys Mason has built a heuristic five-year business-case model to highlight sensitivities in the VoLTE business case with viable solution options for three generic CSP profiles. The clearest commercial benefit is from re-using spectrum, followed by virtualization cost savings with the deployment and operations of a vIMS environment for VoLTE.³

³ Source: Analysys Mason's report (February 2015): [VoLTE business cases: the value of spectrum reuse, an enhanced feature set and virtualisation](#), by Glen Ragoonanan and Stephen Sale

For this white paper, we have expanded our reference VoLTE business case and compared the base-case implementation of a physical, on-premises IR.92 IMS deployment with Single Radio Voice Call Continuity (SRVCC), with a virtualized IMS solution. To baseline this model we have excluded all advanced communications solutions such as RCS-e and WebRTC, to analyse the cost impact with moving from a physical IMS to a vIMS solution for VoLTE services.




Figure 4.3: Breakdown of vIMS cost savings compared with a physical on-premise IMS for VoLTE [Source: Analysys Mason, 2015]

Figure 4.3 provides a breakdown of the cost savings that can be achieved by using vIMS for a VoLTE service deployment and by extension fixed VoIP services, while Figure 4.4 gives a detailed explanation of the vIMS cost-saving results in Figure 4.3.

Figure 4.4: Detailed explanation of the vIMS cost-saving result in Figure 4.3 [Source: Analysys Mason, 2015]

Expected savings	Description and explanation
33% net vIMS cost savings	This is the total net overall cost savings based on lower capex and opex, as explained below compared with a physical, on-premise IMS deployment with SRVCC.
Capex savings	
45% net capex savings	Out of the 33% net vIMS cost savings, 45% will be achieved through capex savings. NFV will have similar savings as in IT with installation and configuration times reduced in deployments, upgrade and prolonged expansion cycles. Details are explained below in the capex breakdown
75% from capacity planning	75% of the capex savings will be achieved through capacity planning. Using NFV, CSPs can increase their virtual network infrastructure (VNI) utilization up to 70% (this can be up to 80% in the IT world) before adding additional infrastructure capacity. This leads to the deferral of infrastructure capex and CSPs can shift from traditional tiered, just-in-case, over-capacity (~30% over for core network and ~50% for all networks) infrastructure, to “pay-as-you grow”, just-in-time infrastructure procurement and pricing models
25% from faster implementation and upgrades	25% of the capex savings will be achieved through faster implementation and upgrades. The lower price of COTS x86 IT servers and storage will contribute to cost savings of about 10% based on IT benchmarks. Faster and cheaper implementation of VNFs (vs. physical elements) have reduced implementation and upgrade time by about 25%
Opex savings	
10% net opex savings	Out of the 33% net vIMS cost savings, 10% will be achieved through opex savings. This is a combination of two main opex-saving areas: <ul style="list-style-type: none"> • Data centre savings: savings related to power, cooling and floor-space reuse (see explanation below) • Orchestration automation: The self-healing, self-monitoring and self-expanding qualities of virtualized environments result in the automation of operational processes, which will reduce operation costs and potentially reduce the number of operations staff supporting voice platforms, thus reducing opex. This saving excludes the service-agility benefits of developing new services, as

	<p>these benefits are less certain than other savings in the model.</p> <p>However, there will be opex additions as well in a vIMS solution that does not exist in a physical IMS solution: increased software licenses and maintenance from VNFs, VIM, OSS changes and orchestrators</p> <p>For vIMS, there are more software licenses and maintenance costs for the virtualized network functions (VNF, e.g. CSCF, MGW, MRF), infrastructure manager (VIM) and NFV orchestrators. These software costs increase the opex of the vIMS solution. These additional software components were factored as a capex addition as well; hence the 45% is a net capex saving</p>
25% from improved service assurance	<p>25% of the opex savings will be achieved from improved service assurance.</p> <p>With a single correlated near real-time network analytic, automated closed-loop fault and performance actions can be enforced for the accurate and efficient up-/down-scaling and reassignment of workloads in minutes without manual intervention. Swivel-chair service assurance can be reduced with a unified service assurance view. Most importantly, near real-time analytics of all service assurance data can reduce T2R process time and provide accurate RCA recommendations in less time</p>
17% from reduced complaints	<p>17% of the opex savings will be achieved from reduced customer complaints.</p> <p>Customer complaints will largely be reduced due to the self-healing capabilities of the NFV platform. To a lesser extent, improved service assurance process efficiency can lead to proactive announcements to call-centre staff to reduce the number of escalated customer complaints</p>
28% from improved service fulfillment	<p>28% of the opex savings will be achieved from improved service fulfillment.</p> <p>NFV automation benefits are concentrated in the service fulfillment inventory and activation and optimization processes. For vIMS, the benefits will largely be focused on inventory. The vIMS VNFs scale up/down and are moved across workload, so the inventory will need to be dynamic and tracked, and will also play a vital part in the capacity-planning processes and the optimization of the operational processes</p>
30% from data-centre cost savings	<p>30% of the opex savings will be achieved from data centre cost savings.</p> <p>As in IT, virtualization provides power, cooling and floor savings, many of which may not be immediately recognized by CSPs. However, floor-space savings reduce the need for data-centre expansion, while the power and cooling reduction will be reflected in reduced energy bills</p>

The following are the key vIMS five-year cost-model assumptions:

- Circuit-switched fallback (CSFB) is in place in all GSM mobile CSPs with LTE deployments to provide an initial voice service to customers with LTE handsets. This cost model does not model non-GSM (CDMA or WiMAX) CSPs.
- CSPs will implement a new IMS because it is simpler and faster than manipulating an established IMS. In practice, many CSPs are deploying multiple separate IMSs to support different services, such as fixed VoIP services.
- Affordable VoLTE-capable smartphones will be available and will not inhibit VoLTE subscriber migration.
- LTE RAN coverage and optimization costs are excluded because LTE roll-out plans for data are assumed to be ongoing and budgeted independently of VoLTE.
- The initial deployment cost is tiered for a minimum capacity of ~500 000 subscribers, which includes 25% excess capacity, which is how CSPs typical size the core network. Expansion and upgrades have lower prices as we assume the initial IMS platform is scalable and thus the price per subscriber decreases over the five-year period.

- We forecast NFV technology maturity in 2016, but operators are testing vIMS today. As such, we have kept the vIMS cost savings as a constant percentage, based on extrapolated IT data, as there is insufficient real-world data at present in telecoms.
- CSPs implementing vIMS have virtualization skills and software in their IT domain and can re-use enterprise virtualization software licences.

Figure 4.4 provides details of areas that CSPs could focus on to ensure they capture these benefits. Specifically for capacity planning and T2R service assurance process benefits, CSPs should start with an open common, complete unified data model as the foundation with near real-time analytics applied, and integration with NFV orchestrators and VIMs enabled.

SDN use cases will present different cost-saving models compared with vIMS – a core network VNF

In 2014, CSPs opex was 68% (~80% of total annual cost) of annual revenue and capex was 16% (~20% of total annual costs). Costs are typically skewed to expensive access network challenges (namely field engineering team (FET) manual processes) and not core network services such as IMS. As such, we have identified the following SDN use case services that will curb CSP high access network costs:

- dynamic VPN or WAN configuration for enterprise customers
- self-service bandwidth on-demand (BoD)
- data centre interconnect services (BoD, bandwidth calendaring, load balancing, etc.).

For the above services to be complete will require on-premise virtualized customer premises equipment (vCPE) that will scale and auto-configure depending on the network dynamics or customer's self-service request.

Analysys Mason has not yet modelled these services using network virtualization, but qualitatively it will reap higher service assurance opex benefits for CSPs than vIMS from actively monitoring and producing remedial feedback to SDN controllers and NFV orchestrators. The service fulfillment opex savings will continue to be higher than service assurance as the automated activation element of these SDN-enabled self-service services. However, service fulfillment costs will continue to be higher for non-active lines that require manual intervention for port configuration and activation (plugging in the cable in the right port). But this will not be the case for service assurance workforce management, as multiple network repairs can be better coordinated and scheduled thanks the self-healing and high-availability capabilities of NFV and SDN technologies.

5. Netscout Adaptive Service Intelligence (ASI) technology

Netscout’s distributed solution architecture – utilizing its ASI technology at the edges along with the centralized service assurance platform nGeniusONE – addresses the demands of a virtualized infrastructure. Netscout has extended its ASI technology to run in a virtual machine on multiple hypervisors, including KVM and VMware. The ASI software running in a virtual machine (VM) provides full visibility into both “north–south” and “east–west” traffic from all tenant VMs. With the virtual ASI solution, all the features of the ASI technology are available in a virtualized environment. From the nGeniusONE perspective, data from a virtual ASI looks the same as data from a physical Netscout probe. This allows CSPs to manage both their physical and virtual network infrastructures in an integrated manner.

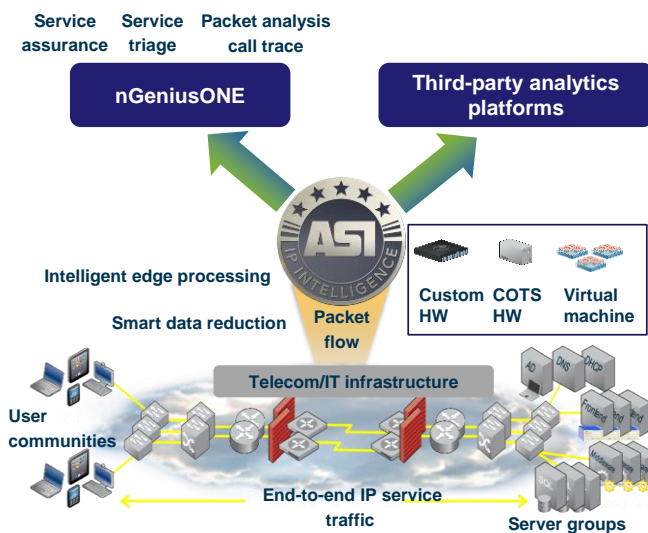


Figure 5.1: Netscout’s ASI technology
[Source: Netscout, 2015]

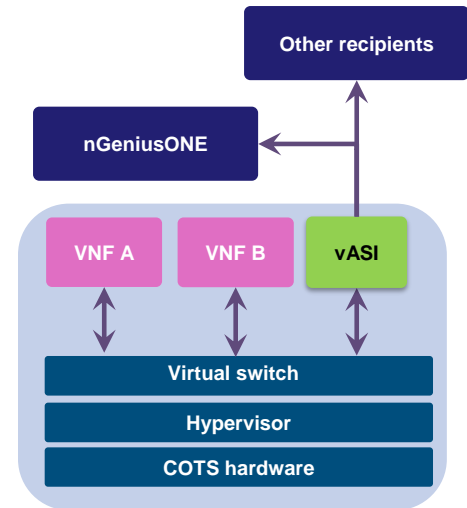


Figure 5.2: Virtual ASI
[Source: Netscout, 2015]

Figure 5.1 and Figure 5.2 illustrate the key capabilities of Netscout’s virtual ASI solution and how they can help CSPs address the challenges of operating a vNGN using ASI as the single data source for creating an open common data model.

Figure 5.3 describes the main capabilities of Netscout’s virtual ASI solution.

Figure 5.3: Netscout virtual ASI [Source: Netscout, 2015]

Capability	Description
Real-time intelligence	The Netscout solution alleviates the need for middleware components which typically introduce additional latency. As the packet data is processed closer to the data source, operations teams can obtain the metadata and performance metrics required for analyzing application and network performance in real time. Furthermore, ASI data can work with policy control and orchestration systems to support real-time changes into the network
Granular and uniform data	ASI-generated metadata includes metrics related to the network performance, application and servers, providing complete visibility into the various layers of the OSI stack and end-user experience. ASI also generates session records along with payload metadata related to the specific application session/transaction. The generated metadata, performance metrics, session records and transaction level details can form the key inputs for the ASI common data model, and can be exported to the Netscout nGeniusONE platform and other analytics platforms for further analysis
Efficient data reduction and injection	In a virtual environment, packets flowing through the hypervisor and vSwitch are likely to see lower throughput due to the overhead introduced by these additional software layers. A virtual Tap technology that passively collects and sends raw data upstream to a physical probe can overburden the hypervisor. Netscout’s virtual ASI technology reduces the volume of data sent upstream to the application layer, allowing the solution to operate at a higher degree of efficiency and scale without compromising its real-time nature
High performance	The performance of virtual ASI can be enhanced by avoiding the hypervisor overhead and adopting SR-IOV or PCI Passthrough technologies. These technologies allow the VM hosting the virtual ASI to bypass the hypervisor and cut through directly to the server NIC card
Ubiquitous deployment	The virtual ASI can scale down to run in small virtualized appliances, e.g. in virtual CPEs, as well as scale up to run in high-end virtualized blade servers in data centres, enabling CSPs to deploy the technology across the entire network. The reduction in complexity compared to current component-based solutions can lead to significant operational savings
Multi-tenancy	A major benefit of embracing virtualization is to support both network services and cloud services on an open common data-centre platform, enabling true multi-tenancy, but at the same time creating the need for management segregation. The Netscout solution can support this type of multi-tenant environment by creating different monitoring port groups on the same hypervisor
Real-time resource management	Using the real-time application metadata and performance metrics from the virtual ASI, CSPs can trigger just-in-time capacity augmentation and assist automating P2P processes providing opex and capex deferral

6. Conclusion and recommendations

Network virtualization evolution led by NFV and SDN technologies promise significant business benefits; prime among these are operational flexibility, service agility and cost optimization. Many major CSPs are embarking on trials and pilots in an effort to prove and operationalize the technology. However, to achieve the business benefits, CSPs accept that they must adopt a new operational approach enabled by automation of operational processes such as T2R and P2P. Existing OSS systems and processes will need to evolve to become more real time in nature to suit the dynamic nature of the virtualized network.

A single truth open common data model based on granular and real-time network performance data will be an important enabler to accelerate the move to automated service assurance operations in virtualized networks. Multi-dimensional network intelligence data sourced from DPI-based virtual software probes can potentially be the key input for the open common data model, which can aid in improving operations flexibility and efficiency.

The following are the key benefits that CSPs stand to gain from the open common data model created using virtualized software probes to collect real-time network performance data:

- improve triage, mean time to resolution times and RCA in the overall T2R process by using granular real-time network intelligence sourced using network traffic flow data from a single uniform data source instead of multiple NEP specific data sources
- reduce the cycle time of the network planning processes by implementing just-in-time capacity augmentation, moving the network capex model to more of a “pay-as-you-grow” model
- alleviate systems integration of multiple ISV solutions for advanced root-cause analysis and troubleshooting
- granularity of data including network performance data, subscriber data, and application and session level intelligence
- virtual probes are scalable and can coexist with the VNFs in the same VNI; they can also exist alongside existing probe solutions.

About the authors



Anil Rao (Senior Analyst) is a member of Analysys Mason's Telecoms Software research team, and is the lead analyst for the Service Assurance programme, focusing on producing market share, forecast and research collateral for the programme. He has published research on IP probes, real-time network analytics and the importance of service assurance in reducing churn and improving customer experience. Anil holds a BEng in Computer Science from the University of Mysore and an MBA from Lancaster University Management School, UK.



Glen Ragoonanan (Principal Analyst) is the lead analyst for Analysys Mason's Infrastructure Solutions, Service Delivery Platforms and Software-Controlled Networking research programmes. He joined Analysys Mason in 2008 and has worked as a consultant on projects on next-generation IT and telecoms networks, systems and technologies for incumbents, new entrants, private companies, regulators and public-sector clients. His primary areas of specialization include operations and business support systems (OSS/BSS) solution architecture and integration for business process re-engineering, business process optimization, business continuity planning, procurement and outsourcing operations and strategies. Before joining Analysys Mason, Glen worked for Fujitsu, designing, delivering and managing integrated solutions. Glen is a Chartered Engineer and project management professional with an MSc from Coventry University.

Analysys Mason does not endorse any of the vendor's products or services discussed in this whitepaper.

Published by Analysys Mason Limited • Bush House • North West Wing • Aldwych • London • WC2B 4PJ • UK
Tel: +44 (0)20 7395 9000 • Fax: +44 (0)20 7395 9001 • Email: research@analysismason.com • www.analysismason.com/research
Registered in England No. 5177472

© Analysys Mason Limited 2015

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means – electronic, mechanical, photocopying, recording or otherwise – without the prior written permission of the publisher.

Figures and projections contained in this report are based on publicly available information only and are produced by the Research Division of Analysys Mason Limited independently of any client-specific work within Analysys Mason Limited. The opinions expressed are those of the stated authors only.

Analysys Mason Limited recognises that many terms appearing in this report are proprietary; all such trademarks are acknowledged and every effort has been made to indicate them by the normal UK publishing practice of capitalisation. However, the presence of a term, in whatever form, does not affect its legal status as a trademark.

Analysys Mason Limited maintains that all reasonable care and skill have been used in the compilation of this publication. However, Analysys Mason Limited shall not be under any liability for loss or damage (including consequential loss) whatsoever or howsoever arising as a result of the use of this publication by the customer, his servants, agents or any third party.